arXiv:2611.23068v1 [cs.AI] 17 Nov 2026

# The Resolution Hypothesis: Agent Evolution's Key Is Identifying Improvements

**Yuanhao Zeng**[1,2,3,*], **Shaocheng Lan**[1,2,*], **Ao Lu**[1,2,*], **Kuang Yan**[2], **Heng Lin**[2], **Xinang Xu**[2], **Yuan Luo**[4], **Yueting Wang**[1], **Yexin Li**[3], **Ruo Shan**[2], **Kan Ren**[1,†]

[1]ShanghaiTech University [2]ByteDance, China [3]BIGAI [4]SUSTech

[*]Equal contribution, [†]Corresponding author

## Abstract

Agent evolution follows a simple loop: a proposer writes candidate artifacts (prompts, skills, memories) and selection accepts those that score higher. This loop has shown promise, yet practitioners face a puzzle: methods trade wins across papers, and evolution often fails to bootstrap on certain tasks. We trace this puzzle to the *Resolution Hypothesis*: agent evolution stalls not because methods fail to propose better artifacts, but because noisy evaluations cannot reliably identify them. Across seven evolution methods and six models (over 10M LLM calls), we find support for the hypothesis at both levels where evaluation scores are used. At the candidate level (choosing artifacts), helpful candidate artifacts exist, but selection is no better than random at telling them from harmful ones, so errors accumulate across rounds. At the method level (ranking methods), typical margins of a few percentage points are smaller than the difference between two independent runs of the same method, so published rankings often rest on luck. Together, these findings indicate that agent evolution is often bottlenecked by reliably identifying improvements. This calls for a shift in how we develop and evaluate self-improving agents, from designing ever more candidate improvement systems to making true improvements detectable and accumulable.

**Correspondence:** zengyh2025@shanghaitech.edu.cn, renkan@shanghaitech.edu.cn
**Project Page:** https://lineshogan.github.io/resolution-hypothesis/

## 1 Introduction

Agent evolution improves frozen LLM agents by refining persistent external artifacts such as prompts, skills, memories, and entire agent harnesses. It typically follows a loop as shown in Figure 1b: a proposer writes candidate artifacts, the resulting agent is evaluated on a task set, and selection accepts the candidate artifacts that score higher. Evolution methods of this family have shown promise across benchmarks and models [24, 27, 37], making the prospect of systems that continuously improve their own workflows increasingly plausible.

Yet the picture is far from settled. Studies often reach conflicting conclusions: the same methods flip between effective and ineffective across benchmarks, and their rankings trade places across papers. On Terminal-Bench [22] and other coding-agent benchmarks, evolution methods have repeatedly been reported to stall [7, 17, 29], while on other benchmarks each new method claims to outperform its predecessors [18, 24, 31, 43]. Whether evolution helps depends on the setting in ways the community cannot yet predict.

This instability erodes confidence in this direction in both academia and industry, and rewards whichever method happened to be luckier at evaluation time. Yet unreliable evaluation affects more than the reported

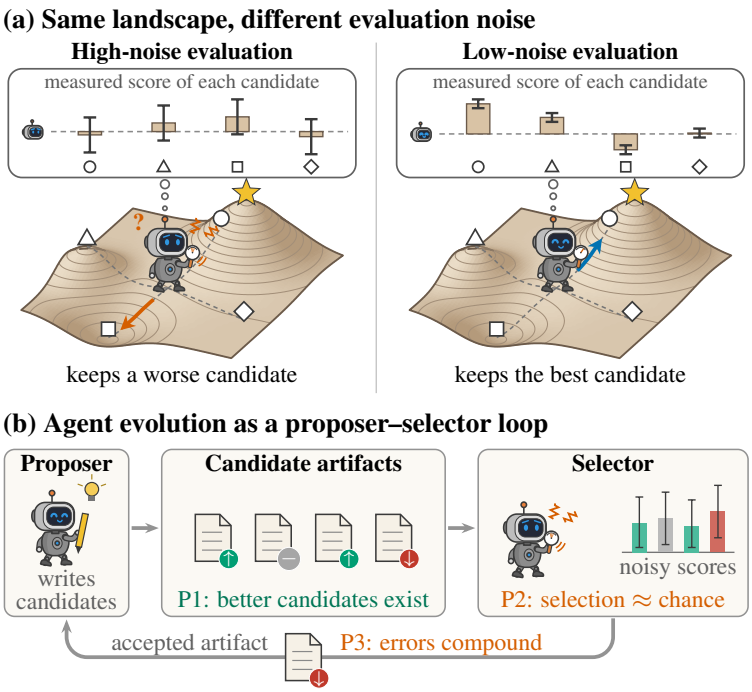



**Figure 1** (a) Four candidate moves on the same landscape, scored under high and low evaluation noise (thought bubble: score relative to the current artifact). Under high noise, a worse candidate artifact scores highest. (b) Agent evolution as a loop: a proposer writes candidate artifacts, and a selector accepts the artifacts with the highest noisy score. P1–P3 are the predictions we test in Section 5.

rankings of evolution methods. At the candidate level, the same kind of evaluation decides which artifacts to carry forward to the next round. As a result, the evolution methods themselves become unreliable.

We argue that this evaluation rests on an untested premise that its scores can identify the improvements evolution brings. We find that the premise fails at both levels where the scores are used. At the method level, the margins that separate methods in papers are smaller than the difference between two independent runs of the same method: on Terminal-Bench, current common practice (Appendix D.3) can only identify the truly best of four text-artifact methods 33% of the time, barely above a random guess (Section 4). At the candidate level, the selector picks the better of two candidate artifacts only about half the time, no better than a coin flip (Section 5.2). Because each round accepts the artifact that happened to score highest, the scores go up but the artifacts do not get better (Section 5.3).

Both failures trace to a single constraint, which we state as a hypothesis.

**The Resolution Hypothesis.** *The bottleneck of agent evolution is identifying which artifacts bring improvement, not proposing them.*

Why is identifying improvements so hard? Because model weights are frozen, an artifact helps only by changing the model's behavior at inference time, while LLM trajectories also vary each time the agent attempts a task. As a result, separating a true behavioral change from evaluation noise is difficult. We define the *resolution* of an evaluation pipeline as the smallest true improvement it can reliably distinguish from noise. Once the true improvement falls below this resolution, candidate selection degrades to chance, so little improvement accumulates during evolution (Figure 1a). Proposal determines how large the improvements on offer are; resolution determines how small an improvement can be and still be reliably identified.

We test this hypothesis at a much larger scale than current practice. Whereas prior work typically relies on a single evolution run (Appendix D.3), we run each method 8 times per model and benchmark and evaluate each

resulting artifact up to 16 times, totaling 1,030 independent evolution runs and over 2.1 million test trajectories. In total, the study covers seven agent evolution methods over six models and four agent benchmarks, plus nine further benchmarks on which we re-examine previously reported gains; to our knowledge, it is the largest reproducibility study of agent evolution to date. Our findings are:

- **Method level (Section 4).** Typical method margins are smaller than what a single run can identify. The large historically reported gains are the clearest instance of our hypothesis: on those benchmarks, evolution succeeds by exploiting evaluation shortcuts that lift nearly every task at once, producing gains large enough to exceed the resolution of the evaluation, and stalls once the evaluation is repaired (Section 4.2).
- **Candidate level (Section 5).** We derive three predictions from the hypothesis. All three are supported by our experiments on Terminal-Bench:
  - **P1** Candidate pools contain artifacts that genuinely improve task performance: proposal is not the bottleneck (Section 5.1).
  - **P2** Common selection mechanisms—whether based on execution scores or LLM judgments—perform near chance at identifying the better of two candidate artifacts (Section 5.2).
  - **P3** Per-round selection errors compound across rounds, producing the large differences between independent runs of the same method that flip rankings (Section 5.3).

Finally, we explore practical ways to address low evaluation resolution at both the candidate and method levels, from lower-variance behavioral checks to separate evaluations of proposers and selectors (Section 6).

In conclusion, these findings suggest a different perspective on agent evolution. The proposer determines what improvements are available, but evaluation resolution determines which improvements can be retained and accumulated. When this resolution is insufficient, true improvements are lost in noise, while large, easily detectable changes can dominate the evolution process even when they exploit evaluation shortcuts. Progress requires either enlarging the improvements being proposed or making smaller improvements reliably detectable. More broadly, developing higher-resolution feedback may be as important as developing more powerful evolution algorithms.

All reproducible materials and trajectory data will be released upon the acceptance of this paper.

## 2 Related Work

Agent evolution methods share a common propose-evaluate-select loop but vary in what they evolve: prompts [2], skills [24, 37], memories [27, 43], or executable code [18, 42]. Several analyses document the fragility of this paradigm, e.g., sensitivity to run-to-run variance and task order [39], sparse feedback-dependent gains [19], and skill-induced failures or contamination from uncontrolled accumulation [7, 17, 29]. These findings point to noisy selection as a shared failure mode, connecting to best-arm identification [12], uncertainty-aware RL evaluation [1], and the false-commit analysis of PACE [30]. Our work isolates the mechanism: we separate evaluation noise from run-to-run variance and use controlled pools to test proposal quality and selection reliability independently. Appendix A gives a full discussion.

## 3 Variance Decomposition and Experimental Setup

To determine whether benchmark scores can identify improvements, we first need a formal model that separates the stochasticity of the evolution algorithm from the noise of the evaluation instrument. This section defines the variance decomposition and outlines our experiments.

### 3.1 Two Sources of Score Variance

Consider an evolution method evaluated on a benchmark. Let artifact $A$ denote the persistent external state produced by the method (a prompt, a skill library, a memory playbook, or an agent harness). Evaluating artifact $A$ once yields an observed score $\hat{Y}(A)$. We define the true quality of artifact $A$ as $Y(A) = \mathbb{E}[\hat{Y} \mid A]$, its expected score over repeated evaluations. A single observed score decomposes as $\hat{Y}(A) = Y(A) + \epsilon$, where

$\epsilon$ is zero-mean evaluation noise ($\mathbb{E}[\epsilon \mid A] = 0$). Across independent evolution runs, the produced artifacts vary; let $\tau^2 = \mathrm{Var}(Y(A))$ denote this run-to-run variance in true quality. By the law of total variance, the variance of an observed score decomposes as:

$$\mathrm{Var}(\hat{Y}) = \underbrace{\mathbb{E}_A\Big[\mathrm{Var}(\hat{Y} \mid A)\Big]}_{\sigma^2} + \underbrace{\mathrm{Var}_A\Big(\mathbb{E}[\hat{Y} \mid A]\Big)}_{\tau^2}. \tag{1}$$

The first term $\sigma^2$ averages over evaluation noise within each artifact; the second $\tau^2$ measures how much artifact quality itself varies across runs. Standard practice often reports a single run evaluated once (Table 8), conflating both sources. To identify true improvements, each source must be bounded separately. Crucially, the resolution depends on the question. Identifying the gain of a single fixed artifact over a reference is limited by $\sigma^2$ alone, while identifying a difference between two methods is limited by $(\tau^2 + \sigma^2)$, because each run of a method draws a new artifact.

### 3.2 Experiments

We study seven evolution methods in two families. Four *text-artifact methods* evolve text given to the agent: prompts (GEPA [2]), skills (SkillOpt-Lite [31], Trace2Skill [24]), and memory playbooks (ReasoningBank [27]). Three *harness methods* evolve the agent's scaffolding (HarnessLens [36], Meta-Harness [15], AHE [18]). We use six agent models: GPT-5.4 nano [25], GPT-5.6 Luna and Terra [26], DeepSeek-V4-Flash [6], Kimi-K3 [13], and GLM-5.2 [9, 41]. Model coverage follows cost: the cheaper experiments use all six models and the more expensive ones a subset. The experiments fall into four groups:

- **Variance measurement** (Section 4.1). All seven methods on ALFWorld [32], SpreadsheetBench [20], OfficeQA [5], and Terminal-Bench v2.1 [22], each evolved and evaluated repeatedly to estimate $\tau$ and $\sigma$.
- **Gain audit** (Section 4.2). The largest reported gains on twelve benchmarks, the first three above and nine from the methods' original papers, re-examined under the released and repaired evaluations; on four of them, we also re-evolve the methods under the repaired evaluations.
- **Candidate pools** (Sections 5.1 and 5.2). Without loss of generality, we use text-based skills to test proposal quality and selection reliability in isolation, constructing two pools of $n$=100 skills tailored to each question (Appendix F.1): a **Proposal Pool** (Section 5.1) generated on 10 Terminal-Bench v2.1 training tasks and evaluated on 10 disjoint test tasks (16 rollouts per pair) to measure out-of-sample proposal gains, and a **Selection Pool** (Section 5.2) combining candidates from Terminal-Bench and DeepSWE trajectories across 11 difficulty-stratified Terminal-Bench tasks (8 to 15 rollouts per pair) to benchmark selection rules and trajectory judges.
- **Trajectory re-evaluation** (Section 5.3). Every intermediate artifact of eight SkillOpt-Lite and eight GEPA evolution runs on Terminal-Bench, re-evaluated eight times per task.

Appendix C gives the setup shared by all experiments and the coverage of each group (Table 5); each group's protocol and full results are in its own appendix, listed in the table. In total, the experiments consumed over 10M LLM calls and 400K CPU core hours.

## 4 Resolution at the Method Level

With both variance components measured, we can test whether standard evaluation practice can identify the method differences these papers report.

### 4.1 Single Reports Cannot Rank Methods

For each method, model, and benchmark, we run evolution 8 times independently (16 times for the text-artifact methods on GPT-5.4 nano) and evaluate each resulting artifact 16 times on the held-out test split (8 times on Terminal-Bench and for harness methods). The unevolved baseline is evaluated equally often. We call a method or the unevolved baseline on a given model and benchmark a *configuration*; in total, these experiments comprise 150 configurations, 1,030 independent evolution runs, and over 2.1 million test trajectories. We

estimate $\sigma$ from how much repeated evaluations of the same artifact differ, and $\tau$ from how much the mean scores of the artifacts from different runs differ (Appendix B.2).

*Most method differences are below one single-run standard deviation.* The standard deviation of the evaluation noise, $\sigma$, ranges from 1.0 to 3.7 percentage points (pp) on SpreadsheetBench, OfficeQA, and ALFWorld, and reaches 4.5 pp on Terminal-Bench, while run-to-run standard deviation $\tau$ ranges from 0 to 12.3 pp. Together, a method evolved once and evaluated once has a typical score standard deviation of 1.7 to 4.1 pp, the median over the configurations of each benchmark (values for every configuration: Table 6). Most method differences are smaller than this: for over 60% of pairs of configurations on the same model and benchmark, the gap between their all-run means falls below one standard deviation of a single-run comparison (Figure 2a). The result holds method by method: for every method, 53% to 79% of its comparisons fall below one single-run standard deviation (Appendix D.1).

*Reported margins fall in the same range.* Across 171 benchmark–model results in the main tables of nine agent-evolution papers, the proposed method leads its strongest competitor by a median of 3.4 pp, with an interquartile range of 1.2 to 6.6 pp. None of these results reports a variance estimate, and only 6 average over more than one evolution run. Appendix D.3 lists the surveyed results.

*Identifying such margins takes many runs.* Identifying a true margin $\Delta$ at 80% power with $\alpha = 0.05$ requires

$$N \geq 2 \cdot \left( \frac{z_{1-\alpha/2} + z_{1-\beta}}{\Delta} \right)^2 \cdot (\tau^2 + \sigma^2) \tag{2}$$

independent runs per method. At the median reported margin of 3.4 pp, this means about 5 runs on SpreadsheetBench, 10 on OfficeQA, 19 on ALFWorld, and 23 on Terminal-Bench (181 on Terminal-Bench at the lower-quartile margin) (Figure 2b). A single run cannot rank methods whose true difference is this small: a results table with one run per method picks the best of the four text-artifact methods only 51% to 73% of the time on SpreadsheetBench, OfficeQA, and ALFWorld, and only 33% on Terminal-Bench, where evaluation noise is largest; a random guess gets 25%. Appendix D.2 gives the full distribution and the ranking simulation.

## 4.2 Large Reported Gains Come from Evaluation Shortcuts

Section 4.1 shows that typical method differences are too small for current evaluation to identify. Some benchmarks, however, show gains of 10 to 40 pp, far above the resolution of a single run, both in the original papers and in our own runs, for example on SpreadsheetBench and LiveMathematicianBench (Appendix E.1). This suggests that evolution can find large improvements after all. The Resolution Hypothesis predicts exactly this pattern: evolution can accumulate an improvement only if the true improvement exceeds the resolution. An evaluation shortcut can produce such an improvement: a single behavior that the evaluation rewards on most tasks of a benchmark at once yields a gain large enough for even a noisy evaluation to identify. We test this prediction by auditing the largest reported gains on twelve benchmarks with three models, using two tests (design in Appendix E).

*One sentence reproduces the gain.* Evolved artifacts contain many instructions, so a large gain could come from many improvements or from a single one. To find out, we audit, instruction by instruction, the artifacts released with GEPA and SkillOpt-Lite and those we evolved ourselves (an LLM agent flags candidate instructions, and we check each by hand), and append to the unevolved baseline only the one rule that appears most responsible for the gain. This single sentence nearly matches what full evolution achieved (Figure 3a): across the 27 model–benchmark pairs with a reported gain above 3 pp, it recovers a median of 90% of the gain. The rule typically takes an evaluation shortcut: it adapts to a grader that credits only one answer format, or to a version of the benchmark, released with a method, that departs from the official evaluation (Table 9). Appendix E.2 discusses the exceptions.

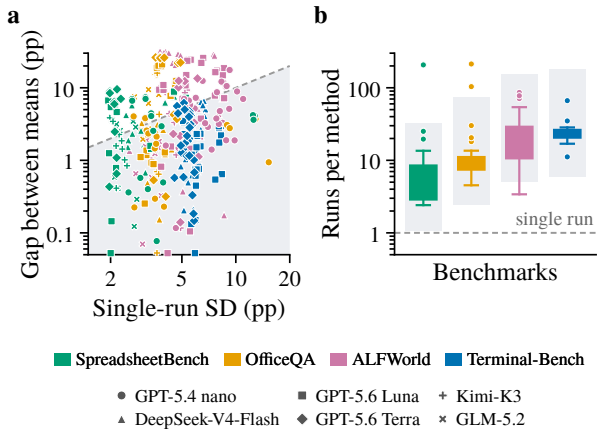



**Figure 2** (a) Gap between all-run means of same-benchmark configuration pairs vs. the standard deviation of a single-run comparison; shaded: gap below one standard deviation. (b) Runs per method to identify the median literature margin (3.4 pp) at 80% power; gray bar: interquartile range 1.2–6.6 pp. Per-method values: Table 6. *Typical method differences are within single-run resolution.*

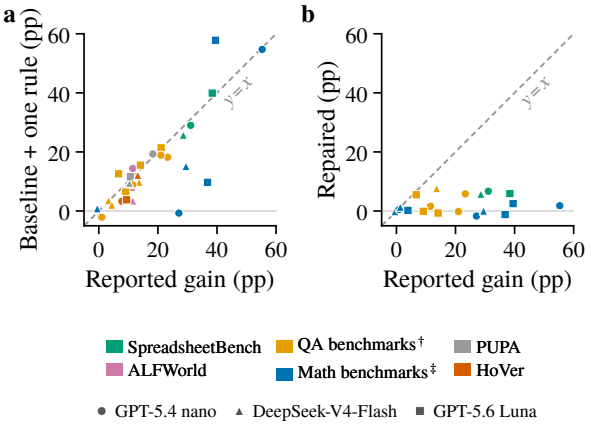



**Figure 3** Gain audit; each point is one model–benchmark pair. (a) Gain of the unevolved baseline plus one rule vs. the reported gain; dashed line: $y{=}x$. (b) Gain of the same artifacts under the repaired evaluation. Benchmark groups and per-pair values: Appendix E.2. *The largest reported gains are reproduced by one sentence and vanish after repair.*

*Repairing the evaluation removes it.* We then repair each evaluation that admits a shortcut, restoring the official evaluation or grading answers by meaning rather than surface form, and re-score the same artifacts. The median gain falls from 17.6 to 0.8 pp (Figure 3b), and on several benchmarks the repaired baseline leaves little headroom at all. Re-running the text-artifact methods under the repaired evaluations does not bring the gain back (Appendix E.3).

This controlled contrast supports our argument: evolution succeeds when an artifact delivers an improvement large enough to exceed the resolution, and fails when only small, task-specific improvements remain. Two questions follow: do such small improvements exist at all, and if so, why does selection fail to accept them?

## 5 Resolution at the Candidate Level

Progress within an evolution run depends on two conditions: proposals must offer true improvements, and selection must identify them. We examine both conditions on frozen candidate pools (Sections 5.1 and 5.2), then follow how selection errors behave across rounds (Section 5.3).

### 5.1 Helpful Artifacts Exist

**Testing P1.** To observe what a proposer produces before selection filters candidates out, we construct the **Proposal Pool**: $n{=}100$ candidate skills generated offline by a SkillOpt-Lite-based proposer from trajectories on 10 Terminal-Bench v2.1 training tasks, and evaluated across 10 disjoint test tasks with 16 independent rollouts per skill–task pair (16,000 rollouts total; protocol in Appendix F.1).

*More candidates yield real gains.* If a proposer writes $K$ candidate skills and we keep the strongest one for each task, does task performance genuinely improve as $K$ grows? The solid curves in Figure 4a track this quantity, *Quality@K*, defined as the expected success rate of the best among $K$ randomly drawn skills. Under a 2-hour execution budget (blue solid curve), which allows multi-step solutions to run to completion, Quality@$K$ climbs steadily from 49.2% at $K{=}1$ to 61.8% at $K{=}100$. Under Terminal-Bench's shorter original timeouts (charcoal solid curve), overall success starts lower (39.7% at $K{=}1$) because slow trajectories are cut short, yet Quality@$K$ still rises to 51.6% at $K{=}100$.

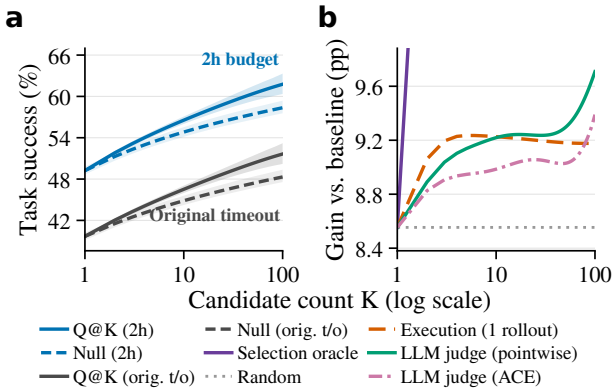



**Figure 4** Headroom in the pool (a) and what selectors capture of it (b) on Terminal-Bench v2.1. (a) Quality@$K$ (solid) vs. matched equal-quality null (dashed) on the 10 disjoint test tasks ($n$=100 skills, $\kappa_{\max}$=100 EB), contrasting the 2-hour budget (blue) and original timeouts (charcoal); shading: 95% intervals. (b) Selected-artifact gain on the 11-task cross-source selection pool (Appendix F.3); Random draws an artifact, and the oracle is clipped to the selector range.

| Selection | Pairwise acc. (%) | False accept. (%) |
|---|---|---|
| Random | $50.0 \pm 0.0$ | $30.5 \pm 14.0$ |
| Execution (1 rollout) | $49.7 \pm 1.9$ | $62.1 \pm 42.9$ |
| SkillOpt-Lite$^\dagger$ | $48.6 \pm 2.6$ | $59.4 \pm 46.2$ |
| GEPA$^\dagger$ | $48.6 \pm 2.6$ | $31.6 \pm 14.5$ |
| LLM judge (pointwise) | $45.4 \pm 4.5$ | $39.9 \pm 17.5$ |
| LLM judge (pairwise) | $49.0 \pm 2.2$ | $57.6 \pm 23.3$ |
| LLM judge (ACE) | $46.9 \pm 3.5$ | $39.0 \pm 15.6$ |

**Table 1** Selection mechanisms evaluated at $K$=$n$ on the fixed candidate pool. **Pairwise acc.**: probability that the selector picks the higher-quality artifact in a head-to-head comparison; 50% is chance. **False accept.**: among artifacts accepted as baseline replacements, the fraction that are actually worse. Values are posterior means ± half-width of 95% credible intervals. $^\dagger$Uses two visible rollouts per artifact; all others use one. *Neither execution-based selection nor LLM judges can reliably identify helpful artifacts among regressions.*

Because taking the maximum over $K$ noisy evaluations would curve upward even if all 100 skills were identical, each solid curve in Figure 4a is paired with a dashed *equal-quality null*—the curve produced when all 100 candidates share the same true success rate and pass through the same empirical-Bayes estimator (Appendix F.2). The vertical gap between the solid and dashed curves isolates genuine skill differences from sampling luck. In both timeout regimes, the observed curve pulls clearly above its dashed null by $K$=5 and maintains a +3.3 to +3.4 pp lead at $K$=100 (with intermediate timeout cutoffs in Appendix F.1). Generating more candidates thus yields genuinely stronger skills, with gains comparable to the post-repair method margins in Section 4.2.

*Shareable artifacts exist.* While Quality@$K$ allows each task to pick its own best skill, a single skill can also improve performance across all 10 test tasks at once. Averaged over all 160 test rollouts per candidate, individual skills in the Proposal Pool span more than 15 percentage points: under the 2-hour budget, the strongest shared skill (`artifact-063`) reaches 57.5% compared with the 49.2% pool mean (+8.3 pp) and 41.9% for the weakest skill, with a similar +8.4 pp lead under the original timeout (48.1% vs. 39.7%). Appendix I.1 inspects representative shared and task-specific skills from our candidate pools to show what these improvements look like in practice.

Pools of this kind are inconsistent with the claim that LLMs cannot propose effective artifacts. Because proposal alone does not explain why evolution stalls, we next isolate selection in Section 5.2.

## 5.2 Selection Cannot Identify Them

**Testing P2.** Because candidate pools contain helpful skills (Section 5.1) yet evolution fails to accumulate them, we next isolate the selection step. We construct the **Selection Pool**: $n$=100 skills generated from trajectories on Terminal-Bench and DeepSWE, evaluated across 11 Terminal-Bench tasks spanning five difficulty levels (Appendix F.1). Each selector observes only the first rollout (two rollouts for native SkillOpt-Lite and GEPA selection), while true skill quality is measured on the remaining held-out rollouts (Appendix F.3). We test whether standard selectors choose better skills more often than chance.

*Under realistic budgets, selection operates near chance.* Greedy accept/reject on execution scores, if effective, would accept improvements and reject regressions with high reliability. Table 1 and the rank correlation

below show the opposite on all three metrics. In head-to-head comparisons, execution-based selection prefers the better artifact only 49.7% of the time, statistically indistinguishable from the 50% achieved by random selection. Regressions make up 30.5% of the pool but 62.1% of accepted candidates, so selection does not filter out regressions; the wide interval leaves open whether it is worse than random acceptance. The rank correlation between a single execution score and posterior artifact quality is 0.054, essentially zero. The native parent-selection operators of SkillOpt-Lite and GEPA, which observe two rollouts per artifact rather than one, achieve 48.6% pairwise accuracy. Every execution-based selector in this study is statistically indistinguishable from random choice.

*No selector benefits from more candidates.* Figure 4b plots the gain over baseline of the selected skill as the candidate budget $K$ grows from 1 to 100, averaged across the 11 tasks (per-task curves in Appendix F.4). While the *Selection oracle* (purple curve at the top) climbs steadily with $K$ as better skills enter the sample, every practical selector remains nearly flat around the random-selection baseline (gray dashed line). Even after examining all 100 candidates, no selector recovers the quality that an oracle already reaches within $K < 5$ candidates.

*Model-based judges do not restore identification.* Model-based judges are a natural candidate for restoring identification, because they read the full trajectory rather than a single terminal outcome [43, 45]. So one might hope they could tell helpful artifacts from harmful ones. Table 1 shows that this additional evidence does not help. No judge exhibits reliable pairwise accuracy. This may be because the judge has limited ability, while its reading noise is independent of evaluation noise. The two therefore accumulate rather than cancel out, introducing additional noise into the evaluation process.

The consistency of this failure across such different selectors supports P2: under realistic budgets, standard selection mechanisms cannot identify helpful artifacts. Appendix F.3 examines whether repeated validation helps.

## 5.3 Noise Accumulates Across Rounds

**Testing P3.** Section 5.2 showed that individual selection steps discriminate at chance level. An evolution run applies such a step in every round, so the relevant quantity is the cumulative effect of per-round errors over rounds.

*In-loop gains reflect noise exploitation.* Near-chance selection implies that most of the gain a run records in its own validation scores, its in-loop gain, should disappear under careful re-evaluation. This is what the re-evaluation shows. We independently re-evaluate every artifact produced by eight evolution runs of SkillOpt-Lite and eight of GEPA on the same Terminal-Bench validation set, with eight rollouts per artifact–task pair where evolution used one (protocol in Appendix G). In-loop gains grow to 6.9 to 9.0 pp over the course of evolution, while re-evaluated gains remain between +0.4 and −3.2 pp (Figure 5). The mechanism follows from the acceptance rule, which compares a candidate's noisy score against the current artifact's and therefore accepts lucky draws from the upper tail of the noise distribution round after round [3, 30]. In-loop scores improve under this process; agent capability does not.

*Run-to-run variance grows over rounds.* Per-round selection errors also drive independent runs apart. All runs begin from the same artifact, so $\tau_0^2 = 0$, and each round of noisy selection is one more step of an unguided random walk over artifact space. By the final round, the corrected run-to-run variance $\hat{\tau}_t^2$ reaches 4.2 to 24.5 squared pp across method–model groups (Figure 6; estimation details in Appendix G). Independent runs of one method now differ by as much as the margins that separate methods in the literature, which is the instability documented in Section 4. These results support P3: per-round selection errors compound across rounds and produce the run-to-run variance behind the ranking instability between methods.

In summary, the variance that prevents a single selection step from identifying improvement is the same variance that destabilizes rankings between methods, compounded over rounds.

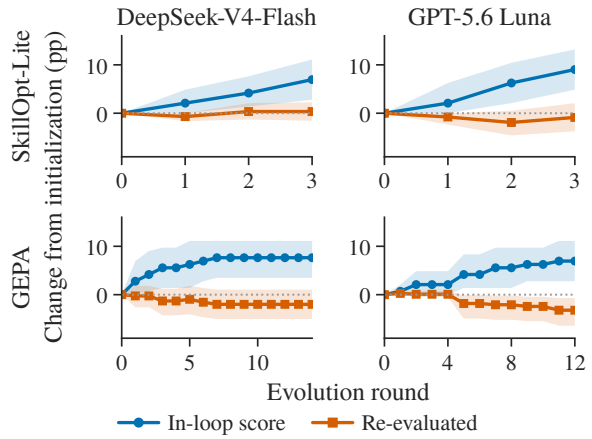



**Figure 5** In-loop vs. independently re-evaluated validation gains across rounds (relative to initialization; shading: 95% bootstrap intervals over eight runs). *In-loop gains reflect noise exploitation, not true improvement.*

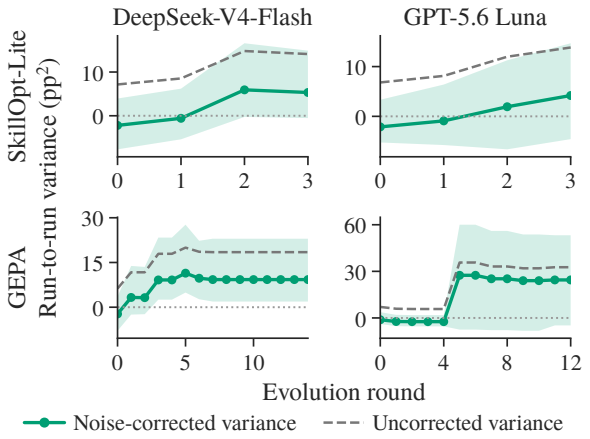



**Figure 6** Run-to-run variance $\hat{\tau}_t^2$ over rounds ($\tau_0^2 = 0$; dashed: raw; solid: noise-corrected; shading: 95% bootstrap intervals over eight runs). *Run-to-run variance grows over rounds.*

## 6 Towards Better Evaluation Resolution

Because low evaluation resolution limits both candidate selection (Section 5) and method comparison (Section 4), we examine two practical ways to detect improvements that terminal pass/fail scores miss: finer-grained behavioral checks at the candidate level, and separate evaluations of proposers and selectors at the method level.

*Candidate level: finer-grained feedback.* Terminal pass/fail rewards may obscure improvements in specific agent behaviors. A promising direction is to supplement them with programmatic checks of intermediate behavior, such as unmasked failure diagnostics on verification commands or whether the agent runs a test after editing a file. Such checks do not fix selection by themselves, since an artifact that triggers more intermediate checks need not solve more tasks. They are however cheap diagnostics: in preliminary analysis (Appendix H), unmasked-diagnostic check rates exhibit 13×–57× lower evaluation noise than terminal pass/fail verdicts (Table 13), and post-edit verification rates expose behavioral shifts that terminal verdicts miss.

*Method level: component-wise evaluation.* Rather than comparing evolution methods only through noisy end-to-end scores, we encourage evaluating their components separately. For proposers, *Quality@K* curves measure the quality of the best candidate among $K$ proposals (Section 5.1). For selectors, *selected-artifact quality* curves measure the quality recovered from a fixed candidate pool as the budget grows (Section 5.2). These diagnostics do not replace end-to-end evaluation, but they split one comparison that falls below the resolution into two that exceed it, revealing whether stagnation lies in proposal or in selection even when method-level margins are below the resolution of a single run. At the reporting level, publishing both evaluation noise ($\hat{\sigma}$) and run-to-run variance ($\hat{\tau}$) alongside point scores lets readers judge directly whether a claimed margin exceeds the evaluation's resolution.

## 7 Conclusion

We propose a hypothesis that agent evolution is bottlenecked not by proposing helpful artifacts but by identifying them. We provide evidence for this at both levels where evaluation scores are used: the margins separating published methods are comparable to the fluctuation of a single run, so single-run rankings are largely determined by noise; helpful candidate artifacts exist, but selection identifies them no better than chance, and these errors accumulate into run-to-run variance. The largest reported gains are consistent with this limit from the other side—they exploit evaluation shortcuts and vanish once the evaluation is repaired.

As a promising direction for future research, we suggest making evaluations resolution-aware: less noisy intermediate process metrics, component-wise diagnostics for proposers and selectors, and variance reporting that lets readers judge whether a claimed margin exceeds the resolution.

## Limitations

*Frozen weights only.* We study agent evolution where model weights remain frozen and evolution operates only on external state (prompts, memories, skills, harnesses). Methods that update model weights via reinforcement learning operate under distinct gradient dynamics and may face different resolution constraints.

*Candidate pool scope.* The candidate-level evidence is evaluated on Terminal-Bench v2.1 only: the Proposal Pool in Section 5.1 evaluates $n$=100 skills proposed from 10 Terminal-Bench training tasks across 10 disjoint test tasks, the Selection Pool in Section 5.2 evaluates $n$=100 skills proposed from Terminal-Bench and DeepSWE trajectories across 11 tasks, and the trajectory re-evaluation in Section 5.3 covers two methods and two models. Whether the conclusions hold for larger pools, different generation strategies, or other benchmark domains remains an open question.

*Model strength and resolution.* Our grid includes models up to Kimi-K3, among the strongest available at the time of writing. We cannot rule out that future, substantially stronger models will exhibit lower evaluation noise and thereby better resolution. In our grid, evaluation noise shrinks for stronger models mainly where their success rate approaches the ceiling, as on ALFWorld (Appendix D.1), which is also where little headroom remains (Appendix E.4); on Terminal-Bench, where headroom is large, its median stays at 3.5 to 4.0 pp for all three models.

*Grid coverage.* Seven methods, six models, and four agent benchmarks represent a broad but not exhaustive survey. Methods that operate on fundamentally different artifact representations, or benchmarks with substantially different task structures, may exhibit different variance profiles.

*Beyond agent artifacts.* Program evolution, which evolves the program that solves a task and scores it deterministically, rather than the agent's prompts, memories, skills, or harness code, has shown steadily accumulating gains in domains such as GPU kernel optimization [4], ARC-AGI-2 [23], and Atari policy search [14]. This paradigm structurally reduces evaluation noise through cheap parallel re-evaluation of deterministic code and more accessible program-environment diagnostics. Our analysis covers the current generation of text-artifact and harness evolution methods. Whether resolution similarly constrains program evolution is therefore an open question that our framework can address but our experiments do not.

## Use of AI Assistants

We used AI coding assistants to write experiment and analysis code. All outputs were reviewed by the authors, who take full responsibility for them. In the gain audit, an LLM agent flags candidate instructions, each of which we check by hand (Section 4.2, Appendix E.1).

## References


[1] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34:29304–29320, 2021.

[2] Lakshya A Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziems, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alexandros G. Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. GEPA: Reflective prompt evolution can outperform reinforcement learning, 2025. URL https://arxiv.org/abs/2507.19457.

[3] Gavin C. Cawley and Nicola L. C. Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(70):2079–2107, 2010. URL https://jmlr.org/papers/v11/cawley10a.html.

[4] Kaiming Cheng, Laura Wang, Jack Khuu, Mark Saroufim, Wenyuan Chi, Jiannan Wang, and Joe Isaacson. KernelAgent: Hardware-guided GPU kernel optimization via multi-agent orchestration. PyTorch Blog, 2026. URL https://pytorch.org/blog/kernelagent-hardware-guided-gpu-kernel-optimization-via-multi-agent-orchestration/. Accessed: 2026-10-09.

[5] Databricks. OfficeQA: A grounded reasoning benchmark. Official dataset card, 2025. URL https://huggingface.co/datasets/databricks/officeqa. Documents from Hugging Face revision 763a836; splits as released with SkillOpt-Lite (commit 4cb4eee); official grader from https://github.com/databricks/officeqa, commit 7b9a3c1.

[6] DeepSeek-AI. DeepSeek-V4: Towards highly efficient million-token context intelligence, 2026. URL https://arxiv.org/abs/2606.19348.

[7] Gen Dong, Yanjie Gao, Liqun Li, Tianyin Xu, Yu Hua, and Fan Yang. Agent skills can be harmful: An empirical study of skill-induced failures in LLM agents, 2026. URL https://arxiv.org/abs/2608.11888.

[8] Matthew Dunn, Levent Sagun, Mike Higgins, V. Ugur Guney, Volkan Cirik, and Kyunghyun Cho. SearchQA: A new Q&A dataset augmented with context from a search engine, 2017. URL https://arxiv.org/abs/1704.05179.

[9] GLM-5-Team. GLM-5: from vibe coding to agentic engineering, 2026. URL https://arxiv.org/abs/2602.15763.

[10] Linyang He, Qiyao Yu, Hanze Dong, Baohao Liao, Xinxing Xu, Micah Goldblum, Jiang Bian, and Nima Mesgarani. LiveMathematicianBench: A live benchmark for research-level mathematical reasoning with proof sketches, 2026. URL https://arxiv.org/abs/2604.01754.

[11] Yichen Jiang, Shikha Bordia, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal. HoVer: A dataset for many-hop fact extraction and claim verification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3441–3460, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.309. URL https://aclanthology.org/2020.findings-emnlp.309/.

[12] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1–42, 2016. URL https://arxiv.org/abs/1407.4443.

[13] Kimi Team. Kimi K3: Open frontier intelligence. Technical report, Moonshot AI, 2026. URL https://github.com/MoonshotAI/Kimi-K3/blob/main/k3_tech_report.pdf.

[14] Zhiyi Kuang, Ryan Rong, YuCheng Yuan, and Allen Nie. Learning game-playing agents with generative code optimization. In *ICML 2025 Workshop on Programmatic Representations for Agent Learning*, 2025. URL https://arxiv.org/abs/2508.19506.

[15] Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-Harness: End-to-end optimization of model harnesses, 2026. URL https://arxiv.org/abs/2603.28052.

[16] Siyan Li, Vethavikashini Chithrra Raghuram, Omar Khattab, Julia Hirschberg, and Zhou Yu. PAPILLON: Privacy preservation from Internet-based and local language model ensembles. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3371–3390, Albuquerque, New Mexico, 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.naacl-long.173. URL https://aclanthology.org/2025.naacl-long.173/.

[17] Xiangyi Li, Yimin Liu, Wenbo Chen, Bingran You, Zonglin Di, Yifeng He, Shenghan Zheng, Kyoung Whan Choe, Jiankai Sun, Shuyi Wang, et al. SkillsBench: Benchmarking how well agent skills work across diverse tasks, 2026. URL https://arxiv.org/abs/2602.12670.

[18] Jiahang Lin, Shichun Liu, Chengjun Pan, Lizhi Lin, Shihan Dou, Zhiheng Xi, Xuanjing Huang, Hang Yan, Zhenhua Han, Tao Gui, and Yu-Gang Jiang. Agentic harness engineering: Observability-driven automatic evolution of coding-agent harnesses, 2026. URL https://arxiv.org/abs/2604.25850.

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

[20] Zeyao Ma, Bohan Zhang, Jing Zhang, Jifan Yu, Xiaokang Zhang, Xiaohan Zhang, Sijia Luo, Xi Wang, and Jie Tang. SpreadsheetBench: Towards challenging real world spreadsheet manipulation, 2024. URL https://arxiv.org/abs/2406.14991. arXiv:2406.14991.

[21] Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. DocVQA: A dataset for VQA on document images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2200–2209, 2021. URL https://openaccess.thecvf.com/content/WACV2021/html/Mathew_DocVQA_A_Dataset_for_VQA_on_Document_Images_WACV_2021_paper.html.

[22] Mike A. Merrill, Alexander G. Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Yeon Shin, Thomas Walshe, E. Kelly Buchanan, Junhong Shen, Guanghao Ye, Haowei Lin, Jason Poulos, Maoyu Wang, Marianna Nezhurina, Jenia Jitsev, Di Lu, Orfeas Menis Mastromichalakis, Zhiwei Xu, Zizhao Chen, Yue Liu, Robert Zhang, Leon Liangyu Chen, Anurag Kashyap, Jan-Lucas Uslu, Jeffrey Li, Jianbo Wu, Minghao Yan, Song Bian, Vedang Sharma, Ke Sun, Steven Dillmann, Akshay Anand, Andrew Lanpouthakoun, Bardia Koopah, Changran Hu, Etash Guha, Gabriel H. S. Dreiman, Jiacheng Zhu, Karl Krauth, Li Zhong, Niklas Muennighoff, Robert Amanfu, Shangyin Tan, Shreyas Pimpalgaonkar, Tushar Aggarwal, Xiangning Lin, Xin Lan, Xuandong Zhao, Yiqing Liang, Yuanli Wang, Zilong Wang, Changzhi Zhou, David Heineman, Hange Liu, Harsh Trivedi, John Yang, Junhong Lin, Manish Shetty, Michael Yang, Nabil Omi, Negin Raoof, Shanda Li, Terry Yue Zhuo, Wuwei Lin, Yiwei Dai, Yuxin Wang, Wenhao Chai, Shang Zhou, Dariush Wahdany, Ziyu She, Jiaming Hu, Zhikang Dong, Yuxuan Zhu, Sasha Cui, Ahson Saiyed, Arinbjörn Kolbeinsson, Jesse Hu, Christopher Michael Rytting, Ryan Marten, Yixin Wang, Alex Dimakis, Andy Konwinski, and Ludwig Schmidt. Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces, 2026. URL https://arxiv.org/abs/2601.11868.

[23] Daniel Mewes, Nayana Bannur, and Catherine Kim. Beating ARC-AGI-2 with code evolution. Imbue Blog, 2026. URL https://imbue.com/blog/2026-02-27-arc-agi-2-evolution. Accessed: 2026-10-09.

[24] Jingwei Ni, Yihao Liu, Xinpeng Liu, Yutao Sun, Mengyu Zhou, Pengyu Cheng, Dexin Wang, Erchao Zhao, Xiaoxi Jiang, and Guanjun Jiang. Trace2Skill: Distill trajectory-local lessons into transferable agent skills, 2026. URL https://arxiv.org/abs/2603.25158.

[25] OpenAI. Introducing GPT-5.4 mini and nano, 2026. URL https://openai.com/index/introducing-gpt-5-4-mini-and-nano/. Accessed: 2026-10-09.

[26] OpenAI. GPT-5.6: Frontier intelligence that scales with your ambition, 2026. URL https://openai.com/index/gpt-5-6/. Accessed: 2026-10-09.

[27] Siru Ouyang, Jun Yan, I-Hung Hsu, Yanfei Chen, Ke Jiang, Zifeng Wang, Rujun Han, Long Le, Samira Daruki, Xiangru Tang, Vishy Tirumalashetty, George Lee, Mahsan Rofouei, Hangfei Lin, Jiawei Han, Chen-Yu Lee, and Tomas Pfister. ReasoningBank: Scaling agent self-evolving with reasoning memory. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=jL7fwchScm.

[28] Valentina Pyatkin, Saumya Malik, Victoria Graf, Hamish Ivison, Shengyi Huang, Pradeep Dasigi, Nathan Lambert, and Hannaneh Hajishirzi. Generalizing verifiable instruction following. In *Advances in Neural Information Processing Systems (Datasets and Benchmarks Track)*, volume 38, 2025. URL https://arxiv.org/abs/2507.02833.

[29] Linfang Shang, Ming Xu, Yiding Sun, Tianle Xia, Lingxiang Hu, Lan Xu, and Ning Zheng. When self-evolution backfires: Pre-commit gating against skill contamination in LLM agents, 2026. URL https://arxiv.org/abs/2608.05810.

[30] Zayx Shawn. PACE: Anytime-valid acceptance tests for self-evolving agents. *arXiv preprint arXiv:2606.08106*, 2026.

[31] Yifei Shen, Bo Li, and Xinjie Zhang. SkillOpt-Lite: Better and faster agent self-evolution via one line of vibe, 2026. URL https://arxiv.org/abs/2607.03451. Code: https://github.com/EvolvingLMMs-Lab/SkillOpt-Lite, commit 4cb4eee.

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

[33] Wenxiao Wang, Priyatham Kattakinda, and Soheil Feizi. Do agent optimizers compound? a continual-learning evaluation on terminal-bench 2.0, 2026. URL https://arxiv.org/abs/2607.14004.

[34] Yike Wang, Huaisheng Zhu, Zhengyu Hu, Yige Yuan, Zhengyu Chen, Shakti Senthil, Hannaneh Hajishirzi, Yulia Tsvetkov, Noah A. Smith, Pradeep Dasigi, and Teng Xiao. Rethinking the evaluation of harness evolution for agents, 2026. URL https://arxiv.org/abs/2607.12227.

[35] Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Benjamin Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, Shubh-Agrawal, Sandeep Singh Sandha, Siddartha Venkat Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, and Micah Goldblum. LiveBench: A challenging, contamination-limited LLM benchmark. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=sKYHBTAxVa.

[36] Jinghan Xu, Yikai Zhang, Aili Chen, Weiyuan Li, Jiaqing Liang, and Deqing Yang. Verify smarter, evolve further: Efficient harness evolution through behavior-aware verification, 2026. URL https://arxiv.org/abs/2608.27311.

[37] Yifan Yang, Ziyang Gong, Weiquan Huang, Qihao Yang, Ziwei Zhou, Zisu Huang, Yan Li, Xuemei Gao, Qi Dai, Bei Liu, Kai Qiu, Yuqing Yang, Dongdong Chen, Xue Yang, and Chong Luo. SkillOpt: Executive strategy for self-evolving agent skills, 2026. URL https://arxiv.org/abs/2605.23904.

[38] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium, 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1259. URL https://aclanthology.org/D18-1259/.

[39] Qinyuan Ye, Yu Li, Yada Pruksachatkun, Jiaxin Zhang, and Chien-Sheng Wu. On the fragility of self-improving agents: Variance, task order, and underspecification, 2026. URL https://arxiv.org/abs/2608.18066v2. arXiv:2608.18066, version 2, September 6, 2026.

[40] Zhaochen Yu, Yingcheng Wu, Zhenfei Yin, Kaiyuan Chen, Zhe Zhao, Mengdi Wang, Shuicheng Yan, and Ling Yang. Recursive experiential-working memory evolution for long-horizon agent harnesses, 2026. URL https://arxiv.org/abs/2608.24876.

[41] Z.ai. GLM-5.2. Model card, 2026. URL https://huggingface.co/zai-org/GLM-5.2.

[42] Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin gödel machine: Open-ended evolution of self-improving agents, 2025. URL https://arxiv.org/abs/2505.22954. arXiv:2505.22954.

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

[44] Congjie Zheng, Chuanyi Xue, Bin Liang, Jun Yang, and Changshui Zhang. SEAGym: An evaluation environment for self-evolving LLM agents, 2026. URL https://arxiv.org/abs/2606.17546.

[45] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL https://openreview.net/forum?id=uccHPGDlao.

# A Extended Related Work

*Methods that evolve persistent artifacts.* Recent efforts to improve frozen LLM agents without updating model weights organize along two dimensions: the macro-structure of the evolution loop, and the object being evolved. Along the structural dimension, methods include *sequential refinement* of a single skill [37]; *set-based evolution* of prompt populations or agent archives [2, 42]; and *memory- and playbook-based evolution* [27, 40, 43]. Along the object dimension, methods evolve *non-executable text artifacts*, including skills distilled from trajectories [24], or *executable code artifacts*, including tools and harness components [18, 42]. Despite their structural variety, all such methods share the same loop: a proposer produces candidate artifacts, an evaluator scores the resulting agent, and a selector accepts higher-scoring candidates. Our analysis applies to this shared loop, and the constraint we identify holds across representations.

*Analyses of why evolving fails.* A parallel line of work documents the fragility of evolving agents and the unreliability of their benchmarks. Ye et al. [39] identify sensitivity to run variance and task order; Wang et al. [34] emphasize matched-budget baselines and held-out evaluation; and Liu et al. [19] find sparse, feedback-dependent gains across evolution rounds. SkillsBench measures the effects of curated skills across tasks [17], while subsequent analyses examine skill-induced failures [7] and contamination from uncontrolled skill accumulation [29]. SEAGym and continual-learning evaluations further distinguish intermediate gains from transfer, regression, and sustained improvement [33, 44].

*Statistical reliability of selection.* Noisy candidate selection relates to best-arm identification [12] and overfitting of model-selection criteria [3]. Uncertainty-aware evaluation is also established in deep reinforcement learning [1]. PACE directly studies false commits under repeated noisy comparisons and proposes an anytime-valid acceptance test [30]. Our study separates evaluation noise from variation across evolution runs and uses candidate pools to examine proposal quality and selection reliability separately.

# B Variance Decomposition

## B.1 Derivation of the Variance Decomposition

An evolution experiment is a two-stage sampling procedure. First, one run of evolution produces an artifact $A$, drawn from the distribution over artifacts induced by the method's stochasticity. Second, the artifact is evaluated, producing an observed score $\hat{Y}$ whose conditional distribution given $A$ reflects the stochasticity of agent execution, with true quality $Y(A) = \mathbb{E}[\hat{Y} \mid A]$.

Write $\hat{Y} = g(A, U)$, where $U$ collects all evaluation-side randomness (instance sampling, decoding stochasticity, environment randomness), independent of $A$. By the law of total variance (Eq. 1),

$$\mathrm{Var}(\hat{Y}) = \underbrace{\mathbb{E}_A\Big[\mathrm{Var}(\hat{Y} \mid A)\Big]}_{\sigma^2} + \underbrace{\mathrm{Var}_A\Big[\mathbb{E}(\hat{Y} \mid A)\Big]}_{\tau^2}. \tag{3}$$

The first term, $\sigma^2$, is the average across artifacts of the variance of evaluating a *fixed* artifact: the variance of the evaluation noise. The second term, $\tau^2$, is the variance across evolution runs of the *true* quality $Y(A)$ of the produced artifact: the run-to-run variance. The decomposition is exact and exhaustive: any fluctuation of an observed score must come either from evaluating a given artifact differently or from drawing a different artifact.

## B.2 Estimators for the Two Variance Components

For each configuration (method, benchmark, and model) we run evolution $R$ times independently and evaluate each resulting artifact $M$ times on the test set, giving observed scores $\hat{Y}_{rj}$, $r = 1, \ldots, R$, $j = 1, \ldots, M$.

*Evaluation noise.* For a fixed artifact $A_r$, the variance of the evaluation noise, $\mathrm{Var}(\hat{Y} \mid A_r)$, is estimated by the sample variance

$$s_r^2 = \frac{1}{M-1} \sum_{j=1}^{M} \left( \hat{Y}_{rj} - \bar{Y}_r \right)^2, \tag{4}$$

where $\bar{Y}_r = \frac{1}{M} \sum_{j=1}^{M} \hat{Y}_{rj}$ is the mean of its $M$ evaluations, and the evaluation-noise component is estimated by the mean within-artifact variance, $\hat{\sigma}^2 = \frac{1}{R} \sum_r s_r^2$; we report $\hat{\sigma} = (\hat{\sigma}^2)^{1/2}$.

*Run-to-run variance.* The artifact means $\bar{Y}_r$ have variance $\tau^2 + \sigma^2/M$, where $\tau^2 = \mathrm{Var}_A[\mathbb{E}(\hat{Y} \mid A)]$. The moment estimator of a random-effects analysis of variance removes the evaluation noise that remains in each mean:

$$\hat{\tau}^2 = \frac{1}{R-1} \sum_r \left( \bar{Y}_r - \bar{Y}_{\cdot\cdot} \right)^2 - \frac{\hat{\sigma}^2}{M}. \tag{5}$$

Both $\hat{\sigma}^2$ and $\hat{\tau}^2$ are unbiased. A negative $\hat{\tau}^2$ means that the artifact means differ no more than evaluation noise alone would make them differ; we then report $\hat{\tau} = 0$, which happens for 9 of the 117 method configurations.

*Design per configuration.* Each method configuration has $R = 8$ independent evolution runs ($R = 16$ for the text-artifact methods on GPT-5.4 nano). Every artifact of a configuration is evaluated the same number of times: $M = 16$ for the text-artifact methods on SpreadsheetBench, OfficeQA, and ALFWorld, and $M = 8$ on Terminal-Bench and for the harness methods, where a single trajectory is substantially more expensive. Estimators are computed within each configuration, so the unbiasedness above holds on all benchmarks; a smaller $M$ only widens the uncertainty of $\hat{\sigma}$ and $\hat{\tau}$. The required-runs formula in Eq. 2 depends on $\tau^2 + \sigma^2$, the variance of a single run evaluated once, and is therefore unaffected by the choice of $M$.

# C Experimental Setup Shared Across Experiments

This appendix describes the models, methods, and benchmarks that the experiment groups of Section 3.2 share. Table 5 lists which of them each group uses and the appendix that describes its protocol.

## C.1 Models and Inference

Six models span the current capability spectrum: GPT-5.4 nano [25], GPT-5.6 Luna and Terra [26], DeepSeek-V4-Flash-0731 [6], Kimi-K3 [13], and GLM-5.2 [9, 41]. Every role of a method, whether solving tasks, reflecting on failures, or rewriting the artifact, uses the same model, so no method benefits from a stronger model writing its artifacts. On ALFWorld, SpreadsheetBench, and OfficeQA, inference uses medium reasoning effort and at most 16,384 completion tokens per call. On Terminal-Bench, reasoning effort is left at each provider's default.

## C.2 Methods and Evolution Budgets

Seven methods span two artifact families. *Text-artifact methods* evolve prompts, skill documents, or memory banks: SkillOpt-Lite [31], a lightweight variant of SkillOpt [37], GEPA [2], Trace2Skill [24], and ReasoningBank [27]. *Harness methods* evolve the agent's own scaffolding: HarnessLens [36] edits its instructions, skills, and tool descriptions, and Meta-Harness [15] and AHE [18] edit its code. All methods run their official implementations, and each method has the same budget on every model. Table 2 summarizes what each method evolves, what its proposer writes, and how it decides what to accept; every validation score is one rollout per task. On Terminal-Bench, the splits have 18 training and 18 validation tasks.

*SkillOpt-Lite.* On ALFWorld, SpreadsheetBench, and OfficeQA, SkillOpt-Lite follows the closed-loop example of its official release and runs at most 10 rounds. Each round proposes one edit from a rollout of 40 training tasks and accepts it if one pass over the validation split beats the current skill by at least 2 points; an edit within 2 points is kept, but the reference score is not updated. The loop stops early after five consecutive rejected or unchanged rounds. On Terminal-Bench, it runs 3 rounds with batches of 12 training tasks and accepts an edit only if it raises the validation score.

| Method | Evolves | Proposer writes | Selection |
|---|---|---|---|
| SkillOpt-Lite | skill document | 1 edit per round, $\leq$10 rounds | accept if the validation score is $\geq$2 points above the current skill |
| GEPA | skill prompt | full rewrite, 1 per iteration | 3-task training gate; parent sampled from the Pareto frontier on validation |
| Trace2Skill | skill document | 1 patch per trajectory, merged | none (all patches kept) |
| ReasoningBank | memory bank | $\leq$3 memories per task | none (all memories kept) |
| HarnessLens | instructions, skills, tool descriptions | 1 edit per cycle | paired check on training tasks; best confirmed candidate on validation |
| Meta-Harness | harness code | new program, 1 per round | best training score of 5 candidates |
| AHE | full workspace | code edits, 1 per round | best validation score of 5 versions |

**Table 2** What each method evolves, what its LLM proposer writes, and how it decides what to accept, as run on ALFWorld, SpreadsheetBench, and OfficeQA. Every validation score is one rollout per task.

*GEPA.* GEPA accepts a child prompt only if it beats its parent on a minibatch of 3 training tasks, scores accepted children once on the validation split, and samples parents from the per-task Pareto frontier, without merge. Its budgets are 1,210, 829, and 939 metric calls on ALFWorld, SpreadsheetBench, and OfficeQA, and at most 128 on Terminal-Bench.

*Trace2Skill and ReasoningBank.* Trace2Skill rolls out the training split once, writes one patch per trajectory, and merges all patches hierarchically into the skill. ReasoningBank streams once through the training split, extracts up to three memories per task, and freezes its bank before testing. On ALFWorld, SpreadsheetBench, and OfficeQA, neither method has score-based selection; in the loop of Figure 1b, keeping every proposal is the degenerate selector that accepts every candidate, and like a selector driven by noisy scores, it accepts regressions along with improvements. On Terminal-Bench, Trace2Skill rolls out each training task three times, builds three candidate skills, and keeps the one that scores best on the training split.

*Harness methods.* Meta-Harness and AHE run 5 rounds, each rolling out the current harness once on the full training split. Meta-Harness returns the candidate with the highest training score; AHE scores each of its five versions once on the validation split and returns the best. HarnessLens has no fixed number of rounds: it compares each candidate with the current harness on 5–6 training tasks, under a cap of five passes over the training split, scores only confirmed candidates once on the validation split, and returns the initial harness if none is confirmed.

## C.3 Benchmarks

Table 3 summarizes the splits, scoring, and agents. All evolution uses the training and validation splits only; all reported scores are on the held-out test split unless stated otherwise.

*ALFWorld.* We halve the released split by stratified sampling (seed 42). The environment runs in process, and an episode succeeds when the agent completes the household goal within 50 steps.

*SpreadsheetBench.* We use the Verified-400 subset. One test task has no test case and scores 0 for every artifact, so scores are over all 281 test tasks and gains are unaffected. The agent writes code that edits the workbook, run in a sandbox. The released grader reads formula cells without recalculating them, so every

| Benchmark | Train/val/test | Score | Agent | Limit per task |
|---|---|---|---|---|
| ALFWorld | 100/70/67 | success | ReAct | 50 steps |
| SpreadsheetBench | 80/39/281 | exact match | code generation | 30 turns, 600 s |
| OfficeQA | 49/49/148 | exact match | tool use, offline | 24 tool calls |
| Terminal-Bench v2.1 | 18/18/53 | task tests | Mini-SWE-agent | task timeout |

**Table 3** Benchmark splits, scoring, and agents. Scores are binary per task. Text-artifact methods and the baseline run the agent in the table; harness methods evolve the agent itself.

| **Training (18 tasks)** | **Validation (18 tasks)** |
|---|---|
| cancel-async-tasks | chess-best-move |
| circuit-fibsqrt | custom-memory-heap-crash |
| count-dataset-tokens | dna-assembly |
| filter-js-from-html | dna-insert |
| fix-git | extract-elf |
| install-windows-3.11 | feal-linear-cryptanalysis |
| kv-store-grpc | fix-code-vulnerability |
| log-summary-date-ranges | llm-inference-batching-scheduler |
| modernize-scientific-stack | mcmc-sampling-stan |
| mteb-retrieve | path-tracing-reverse |
| overfull-hbox | portfolio-optimization |
| password-recovery | pypi-server |
| polyglot-c-py | pytorch-model-cli |
| polyglot-rust-c | qemu-startup |
| sanitize-git-repo | query-optimize |
| schemelike-metacircular-eval | reshard-c4-data |
| torch-tensor-parallelism | sam-cell-seg |
| video-processing | tune-mjcf |

**Table 4** Terminal-Bench v2.1 training and validation splits for method-level evolution (18 tasks each, sorted alphabetically); the remaining 53 tasks form the held-out test split.

formula scores as empty (Appendix E.1); for the variance measurement we recalculate both the predicted and the reference workbook with LibreOffice before comparison.

*OfficeQA.* Questions are answered from 697 parsed U.S. Treasury Bulletins. The agent has no web access and searches the documents with file-system tools; the parsed pages of each question's source documents are included in the prompt. Answers are graded by normalized exact match, the grader of the SkillOpt-Lite release.

*Terminal-Bench v2.1.* The benchmark contains 89 tasks. For the method-level evolution and trajectory-re-evaluation experiments (Sections 4 and 5.3), we partition the 89 tasks into 18 training, 18 validation, and 53 held-out test tasks (Table 4), stratified by task difficulty so that all three splits span a balanced mix of easy, medium, and hard tasks. Each task runs offline in its own sandbox with the official task image and timeout, and is scored by its own task-specific tests. The baseline and the text-artifact methods run Mini-SWE-agent 2.4.6; harness methods evolve their own Mini-SWE-agent workspace, starting from version 1.14.4.

For the candidate-level experiments in Section 5, we construct two pools of $n$=100 skills (Appendix F.1). The **Proposal Pool** (Section 5.1) generates 100 skills offline from execution trajectories on 10 Terminal-Bench v2.1 training tasks and evaluates them across 10 disjoint test tasks to measure out-of-sample proposal gains. The **Selection Pool** (Section 5.2) generates 100 skills from trajectories on six Terminal-Bench tasks (filter-js-from-html, mteb-retrieve, count-dataset-tokens, kv-store-grpc, modernize-scientific-stack, log-summary-date-ranges) and three DeepSWE tasks (anko-typed-variable-bindings, abs-module-cache-flags, ytt-jsonpath-query-api), and evaluates them across 11 Terminal-Bench tasks spanning five baseline difficulty

| Experiment | Benchmarks | Methods | Models | Details |
|---|---|---|---|---|
| Variance measurement | ALF, SSB, OQA | text-artifact (4) | all six | App. D |
| | ALF, SSB, OQA | harness (3) | nano, DeepSeek, Luna | |
| | Terminal-Bench | all seven | DeepSeek, Luna, Terra | |
| Gain audit | 12 (App. E.1) | released and our artifacts | nano, DeepSeek, Luna | App. E |
| Candidate pools | Terminal-Bench (Proposal Pool / Selection Pool) | skills (100 per pool) | DeepSeek | App. F |
| Trajectory re-evaluation | Terminal-Bench | SkillOpt-Lite, GEPA | DeepSeek, Luna | App. G |

**Table 5** Benchmarks, methods, and models of each experiment group, and the appendix with its protocol and full results. ALF: ALFWorld; SSB: SpreadsheetBench; OQA: OfficeQA. Models: "all six" comprises GPT-5.4 nano, DeepSeek-V4-Flash, GPT-5.6 Luna, GPT-5.6 Terra, Kimi-K3, and GLM-5.2.

levels: the six above plus five held out from skill generation (torch-tensor-parallelism, video-processing, cancel-async-tasks, sanitize-git-repo, fix-git).

### C.4 Coverage of Each Experiment Group

Model coverage follows cost. The text-artifact methods on ALFWorld, SpreadsheetBench, and OfficeQA, the cheapest setting, use all six models. Two settings are far more expensive than the rest. Harness methods write and debug code at every evolution step, and one harness evolution run consumes a median of 26M tokens. On Terminal-Bench, every task is a long interactive episode in its own container: one evolution run consumes about 43M tokens, and evaluating each resulting artifact takes another 424 containerized episodes (53 test tasks $\times$ 8 repetitions). We therefore run these settings on a subset of the models.

## D Variance Measurement

### D.1 Construction of Figure 2 and Full Results

*Configurations and their noise.* A configuration is a method, or the unevolved baseline, on one model and benchmark. For each method configuration we evolve $R$ artifacts independently and evaluate each $M$ times on the test split ($R \times M = 8 \times 16$; $16 \times 16$ for GPT-5.4 nano; $8 \times 8$ on Terminal-Bench, 53 test tasks, and for harness methods), and the unevolved baseline is evaluated $M$ times. Harness methods run their own agent workspace, so their baseline is the initial harness in that workspace, evaluated at least 8 times; on SpreadsheetBench and OfficeQA they are graded with the repaired graders of Appendix E. We report the mean over all runs and evaluations and the estimates $\hat{\tau}$ and $\hat{\sigma}$ of Appendix B.2, with $\hat{\tau} = 0$ where the moment estimate is negative. The standard deviation of a single run, one evolution evaluated once, is $s = (\tau^2 + \sigma^2)^{1/2}$; for the baseline, which is not evolved, $s = \sigma$. Table 6 lists these quantities for every configuration.

*Figure 2a.* Each point is a pair of configurations on the same model and benchmark: the horizontal axis is the standard deviation of a single-run comparison, $(s_1^2 + s_2^2)^{1/2}$, and the vertical axis is the absolute difference between their means. Two configurations are paired only if the same grader scores them, and a baseline only with the methods that run in its agent workspace, so that a difference does not mix the method with the grader or the agent. Harness methods are therefore paired with each other and with their initial harness, and, on ALFWorld and Terminal-Bench, where the grader is shared, also with the text-artifact methods. This leaves 341 pairs, 217 of which fall below the diagonal. The share is similar for pairs of two methods (144 of 224) and for pairs of a method and a baseline (73 of 117), and for the pairs involving each method: 54 of 101 for GEPA, 61 of 101 for ReasoningBank, 63 of 101 for SkillOpt-Lite, 68 of 101 for Trace2Skill, 38 of 57 for AHE, 40 of 57 for Meta-Harness, and 37 of 47 for HarnessLens.

*Figure 2b.* For each method configuration, the number of runs per method needed to identify a margin $\Delta$ at 80% power with $\alpha = 0.05$ is $N = \lceil 2(z_{1-\alpha/2} + z_{1-\beta})^2 s^2/\Delta^2 \rceil$ (Appendix D.2), evaluated at the median literature margin $\Delta = 3.4$ pp. Each box summarizes the method configurations of one benchmark; the gray bar spans $N$ at the median single-run variance of the benchmark for $\Delta$ between 1.2 and 6.6 pp, the interquartile range of reported margins (Appendix D.3).

| | | | Text-artifact methods | | | | | | | | | | | | Harness methods | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | Baseline | | SkillOpt-Lite | | GEPA | | Trace2Skill | | ReasoningBank | | Initial harness | | HarnessLens | | Meta-Harness | | AHE | |
| Benchmark | Model | Runs×evals | Mean | $\sigma$ | Mean | $\tau/\sigma$ | Mean | $\tau/\sigma$ | Mean | $\tau/\sigma$ | Mean | $\tau/\sigma$ | Mean | $\sigma$ | Mean | $\tau/\sigma$ | Mean | $\tau/\sigma$ | Mean | $\tau/\sigma$ |
| SpreadsheetBench | nano | 16×16 | 57.7 | 1.8 | 64.3 | 1.3/1.5 | 64.7 | 2.1/1.5 | 61.7 | 2.3/1.7 | 57.3 | 0.6/1.7 | 58.4 | 2.3 | 58.9 | 2.5/2.0 | 59.0 | 0.8/1.6 | 54.8 | 12.2/1.9 |
| | DeepSeek | 8×16 | 66.9 | 1.9 | 72.6 | 0.9/1.6 | 71.0 | 1.8/1.7 | 70.3 | 0.9/1.7 | 68.4 | 0.0/1.6 | 69.4 | 1.3 | † | | 75.7 | 3.4/1.7 | 77.0 | 4.0/1.6 |
| | Luna | 8×16 | 69.4 | 1.6 | 74.0 | 0.7/1.2 | 75.3 | 0.4/1.3 | 68.6 | 0.5/1.3 | 68.6 | 0.1/1.4 | 70.6 | 1.4 | 72.1 | 2.6/1.5 | 70.7 | 0.7/1.3 | 79.4 | 4.1/1.4 |
| | Terra | 8×16 | 69.2 | 1.4 | 75.2 | 0.8/1.3 | 77.7 | 0.7/1.2 | 68.2 | 0.7/1.5 | 68.8 | 0.4/1.3 | – | | – | | – | | – | |
| | Kimi | 8×16 | 73.6 | 1.6 | 75.5 | 1.1/1.3 | 77.1 | 0.5/1.3 | 74.6 | 1.1/1.4 | 73.2 | 0.1/1.5 | – | | – | | – | | – | |
| | GLM | 8×16 | 69.8 | 1.5 | 73.3 | 0.4/1.3 | 72.7 | 0.3/1.5 | 69.9 | 1.5/1.7 | 70.2 | 0.0/1.6 | – | | – | | – | | – | |
| OfficeQA | nano | 16×16 | 12.5 | 1.6 | 10.2 | 2.4/1.6 | 10.0 | 2.2/1.7 | 8.5 | 1.7/1.3 | 12.3 | 0.7/2.1 | 18.8 | 3.0 | 18.5 | 0.2/1.8 | 22.6 | 12.3/2.2 | 21.6 | 8.4/2.4 |
| | DeepSeek | 8×16 | 55.3 | 2.0 | 58.5 | 0.7/2.1 | 58.4 | 2.2/2.1 | 57.0 | 2.5/1.9 | 55.8 | 1.1/2.2 | 64.4 | 1.8 | † | | 67.2 | 1.6/1.8 | 68.1 | 1.8/1.8 |
| | Luna | 8×16 | 35.5 | 2.6 | 56.6 | 1.8/2.2 | 55.3 | 1.5/2.4 | 44.1 | 2.8/2.3 | 36.4 | 0.8/2.2 | 62.8 | 2.0 | 61.6 | 0.0/2.1 | 63.0 | 1.1/2.1 | 63.9 | 1.6/1.8 |
| | Terra | 8×16 | 36.7 | 3.0 | 62.6 | 1.7/2.0 | 63.2 | 1.5/1.7 | 58.9 | 3.4/1.9 | 37.0 | 1.7/2.1 | – | | – | | – | | – | |
| | Kimi | 8×16 | 49.5 | 2.5 | 64.7 | 1.8/1.9 | 64.7 | 1.7/1.9 | 65.4 | 1.7/2.0 | 45.2 | 0.9/2.6 | – | | – | | – | | – | |
| | GLM | 8×16 | 52.7 | 1.7 | 60.2 | 1.3/1.7 | 58.8 | 2.0/1.9 | 54.3 | 4.2/2.1 | 50.7 | 1.0/2.3 | – | | – | | – | | – | |
| ALFWorld | nano | 16×16 | 48.4 | 2.4 | 55.4 | 3.7/2.7 | 60.6 | 7.4/2.9 | 51.7 | 5.7/2.7 | 43.0 | 3.6/3.0 | 50.4 | 2.7 | 53.5 | 7.1/2.5 | 48.8 | 0.7/2.9 | 50.3 | 2.4/2.9 |
| | DeepSeek | 8×16 | 65.0 | 3.2 | 69.4 | 6.6/3.3 | 71.3 | 3.2/3.2 | 76.5 | 2.9/2.9 | 69.2 | 3.0/3.3 | 66.8 | 3.7 | 66.4 | 0.5/3.4 | 97.1 | 1.3/1.5 | 98.5 | 1.2/1.0 |
| | Luna | 8×16 | 68.2 | 3.1 | 69.7 | 1.8/2.5 | 79.1 | 1.5/2.0 | 72.0 | 3.4/2.8 | 68.3 | 3.9/2.4 | 70.5 | 2.6 | 69.7 | 1.2/2.5 | 86.9 | 6.7/2.7 | 94.1 | 4.2/2.0 |
| | Terra | 8×16 | 74.6 | 3.2 | 74.8 | 2.5/2.8 | 87.1 | 3.0/2.0 | 73.6 | 4.3/2.5 | 76.9 | 1.7/2.6 | – | | – | | – | | – | |
| | Kimi | 8×16 | 80.6 | 2.0 | 85.0 | 2.5/2.8 | 90.3 | 2.3/2.0 | 81.8 | 3.6/2.1 | 81.7 | 1.5/2.1 | – | | – | | – | | – | |
| | GLM | 8×16 | 82.6 | 1.6 | 84.5 | 2.1/1.8 | 84.2 | 2.0/1.8 | 83.3 | 1.5/1.7 | 82.7 | 1.6/2.0 | – | | – | | – | | – | |
| Terminal-Bench | DeepSeek | 8×8 | 72.4 | 3.0 | 73.3 | 0.0/3.9 | 73.5 | 1.5/4.0 | 72.7 | 0.0/4.1 | 71.9 | 0.0/4.5 | 71.0 | 4.4 | † | | 69.6 | 1.5/3.6 | 66.9 | 3.3/3.8 |
| | Luna | 8×8 | 54.0 | 3.2 | 57.4 | 1.0/3.9 | 54.6 | 6.0/3.6 | 57.5 | 0.0/4.0 | 56.3 | 1.3/3.9 | 56.0 | 4.1 | 56.8 | 1.3/4.2 | 57.5 | 1.4/4.1 | 55.2 | 1.5/4.3 |
| | Terra | 8×8 | 70.3 | 4.2 | 71.5 | 0.2/3.5 | 69.6 | 1.4/3.8 | 70.9 | 0.0/2.9 | 69.3 | 0.0/4.0 | 64.7 | 3.6 | 65.3 | 1.7/3.3 | 66.5 | 1.5/3.5 | 69.5 | 1.7/3.8 |

**Table 6** Variance measurement for every configuration, in pp, with the quantities and estimators defined in Appendix D.1. Runs×evals: independent evolution runs per method and test evaluations per artifact. –: not run (coverage in Table 5). †: every run returned the initial harness, so this is not a method configuration.

| | True gap $\Delta$ (pp) | | |
|---|---|---|---|
| Benchmark | 1.2 | 3.4 | 6.6 |
| SpreadsheetBench | 33 | 5 | 2 |
| OfficeQA | 75 | 10 | 3 |
| ALFWorld | 152 | 19 | 6 |
| Terminal-Bench v2.1 | 181 | 23 | 6 |

**Table 7** Independent runs per method required to identify a gap at 80% power ($\alpha = 0.05$). The three gaps are the lower quartile, median, and upper quartile of margins reported in nine agent-evolution papers.

*Distribution of a reported score.* Figure 7 shows, for each method, model, and benchmark, the distribution of the score that a single run would report, one evolution evaluated once, obtained by resampling the observed artifacts and their repeated evaluations. The distributions of different methods overlap for most models and benchmarks, and many are multimodal because the artifacts of independent runs differ in quality.

## D.2 Repetitions Required for Method-Level Comparisons

The number of complete evolution repetitions required to identify a gap $\Delta$ between two methods with confidence $1-\alpha$ and power $1-\beta$ follows from the two variance components of Section 3. A reported difference between two methods, each run $N$ times and evaluated $M$ times per artifact, has variance $2(\tau^2+\sigma^2/M)/N$, so

$$N \approx \frac{2\,(z_{1-\alpha/2}+z_{1-\beta})^2\,(\tau^2+\sigma^2/M)}{\Delta^2}. \tag{6}$$

Table 7 instantiates this expression with $\alpha = 0.05$, $1-\beta = 0.8$, and $M = 1$, using the median single-run variance over the configurations of each benchmark, at the quartiles of the margins reported in the literature (Appendix D.3). Re-evaluation shrinks only $\sigma^2$, which is the larger component on three of the four benchmarks: with many evaluations per artifact, the median margin would need about 1, 5, 9, and 3 runs instead of 5, 10, 19, and 23, so reducing evaluation noise directly reduces the cost of identifying a margin.

*Ranking agreement of a single run.* Section 4.1 also reports how often a single run ranks the four text-artifact methods correctly. For each model and benchmark, we draw for every method one evolution run uniformly at random and one evaluation of its artifact, rank the methods by these scores, and compare the result with the ranking by their means over all runs and evaluations. Over 20,000 draws per model and benchmark, we

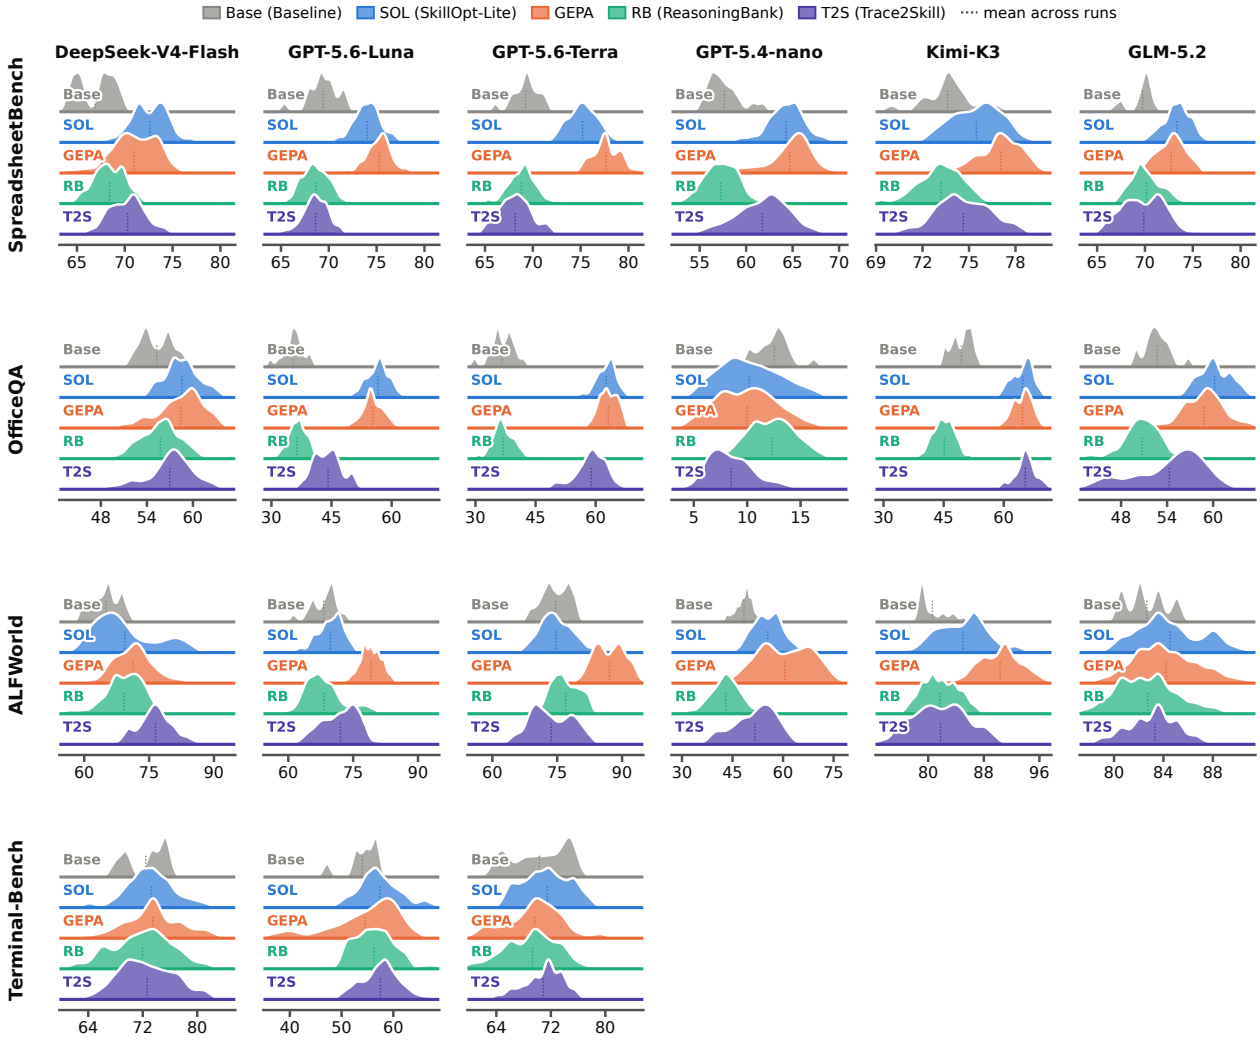



**Figure 7** Distribution of the score a single run would report (Appendix D.1) for the unevolved baseline (Base) and the four text-artifact methods (SOL: SkillOpt-Lite, GEPA: GEPA, RB: ReasoningBank, T2S: Trace2Skill), from $10^5$ resamples of the observed artifacts and evaluations. Dotted line: mean over all runs and evaluations. Ridges share one height; horizontal ranges differ between panels.

record how often the top-ranked method and the full ranking agree with those of the means, and average over models. The top-ranked method agrees 67%, 51%, 73%, and 33% of the time on SpreadsheetBench, OfficeQA, ALFWorld, and Terminal-Bench, and the full ranking 33%, 29%, 24%, and 8% of the time; a random ranking of four methods gets the top method right 25% and the full ranking $1/24 \approx 4\%$ of the time.

## D.3 Margins Reported in Agent-Evolution Papers

To place our noise estimates against the differences that papers claim, we surveyed the results reported in agent-evolution papers.

*Papers.* We include the papers introducing the seven methods we evaluate (GEPA [2], Trace2Skill [24], ReasoningBank [27], SkillOpt-Lite [31], Meta-Harness [15], HarnessLens [36], and AHE [18]), together with ACE [43] and SkillOpt [37], two methods that these papers build on or compare against. From each paper we take the main results tables, listed in Table 8; ablations, analyses, and appendix tables are not included.

| Paper | Tables | Columns | Margin | Runs |
|---|---|---|---|---|
| GEPA | 1, 2 | 12 | 2.5 | not stated |
| ACE | 1, 2 | 12 | 10.4 | not stated |
| ReasoningBank | 1–3 | 41 | 3.6 | not stated |
| Trace2Skill | 5 | 6 | 7.6 | mean of 3 |
| SkillOpt | 1 | 52 | 3.0 | not stated |
| SkillOpt-Lite | 2 | 30 | 2.0 | not stated |
| Meta-Harness | 2 | 3 | 3.8 | not stated |
| HarnessLens | 1 | 12 | 0.8 | 1 |
| AHE | 1 | 3 | 0.0 | 1 |
| All | | 171 | 3.4 | |

**Table 8** Survey of reported margins (Appendix D.3). Columns: benchmark–model columns with an optimization competitor. Margin: median, in pp, of the proposed method's score minus its strongest competitor's. Runs: evolution runs behind each reported score, as stated in the paper; "not stated" marks a single reported score, which we treat as a single run.

*What we record.* Each column of a results table, that is, one benchmark evaluated with one model in one setting, gives one record. For each record we take, from the table, the score of the proposed method, the score of its strongest competitor in the same column, and the score of the unoptimized baseline, and, from the text and captions, the number of evolution runs behind the reported score and whether an error bar or variance estimate is given. A competitor is another optimization method (prompt, skill, memory, or harness optimization); zero-shot prompting, human-written prompts or harnesses, and fixed retrieval baselines are not competitors. We use accuracy-type metrics only (accuracy, success rate, pass@1, task-goal completion), expressed in percentage points, and skip columns that average other columns of the same table, so that no result is counted twice. This gives 178 columns, 171 of which have an optimization competitor. The margin of a column is the score of the proposed method minus that of its strongest competitor; the gain is the score of the proposed method minus that of the unoptimized baseline.

*Results.* Across the 171 columns, the median margin is 3.4 pp (interquartile range 1.2–6.6; quartiles by linear interpolation), and the median gain over the unoptimized baseline is 9.3 pp. The median margin is stable across subsets. Taking the median within each paper first gives 3.0 pp (interquartile range 2.0–3.8 across papers); excluding Mind2Web, the benchmark with the most columns, gives 3.5; and restricting the survey to the 52 columns on the four benchmarks we measure (SpreadsheetBench, OfficeQA, ALFWorld, and Terminal-Bench) also gives 3.5, so Section 4.1 uses the full-survey quartiles. None of the 178 columns reports an error bar or a variance estimate. Only the six Trace2Skill columns average over several evolution runs (three seeds); 20 columns state a single run, and the remaining 152 do not state the number of runs and report a single score. Competitor scores also depend on who runs them: in four of the papers, the strongest competitor scores at or below the unoptimized baseline on some columns (15 in total).

# E Gain Audit

## E.1 Audit Design

*Benchmarks and artifacts.* The audit covers three of the four benchmarks of our variance grid (all but Terminal-Bench) and the nine further benchmarks on which the original papers of the audited methods report optimized artifacts: HotpotQA [38], PUPA [16], HoVer [11], IFBench [28], AIME 2025, and LiveBench-Math [35] from GEPA [2], and LiveMathematicianBench [10], SearchQA [8], and DocVQA [21] from the skills released with SkillOpt-Lite [31]. On the grid benchmarks the reported gain is that of our own evolved artifacts (the best method, averaged over evolution runs); elsewhere it is the gain of the released artifact when transferred to our models, which is typically smaller than the gain reported for the model it was optimized on. All tests use GPT-5.4 nano, DeepSeek-V4-Flash, and GPT-5.6 Luna; gains are paired by task and averaged over at least two evaluation repetitions (four or more except on SearchQA). DeepSeek-V4-Flash does not

| Benchmark | Shortcut in the released evaluation | Repair |
|---|---|---|
| SpreadsheetBench | Formula cells are read without recalculation, so every formula scores as empty | Recalculate with LibreOffice before comparison, as the official evaluation does |
| AIME 2025 | The parser accepts only a bare integer | Equivalence check after removing LaTeX wrappers |
| LiveMathematicianBench | A "stronger result" option is correct whenever it appears | Score only questions without that option |
| DocVQA | Near-exact string match (ANLS $\geq$ 0.999) | Semantic grading |
| HotpotQA | Exact match rejects correct answers written differently | Semantic grading |
| OfficeQA | The version released with SkillOpt-Lite drops the official instruction to answer with a bare value | Restore the official instruction and reward |
| SearchQA | Exact match rejects correct answers written differently | Semantic grading |
| LiveBench-Math | Format-specific matching on AMPS problems | Mathematical-equivalence grading |
| PUPA, ALFWorld, HoVer, IFBench | No shortcut found | – |

**Table 9** Shortcuts in the released evaluations and their repairs.

accept image input, so DocVQA uses the other two models. For GPT-5.4 nano on OfficeQA we did not run the repaired evaluation: its reported gain, 1.0 point, is within noise.

*Reported and reproduced gains.* The large gains are not specific to our setup. SkillOpt-Lite reports gains of 36 to 45 pp on SpreadsheetBench for the four GPT-5 models it evaluates (+36.3 pp on GPT-5.4 nano) and 20 to 37 pp on LiveMathematicianBench [31]. Our runs show gains of the same size: 29 to 38 pp on SpreadsheetBench with the artifacts evolved in our grid, and 27 to 37 pp on LiveMathematicianBench with the skills released with SkillOpt-Lite (Table 11).

*The one-rule test.* For each benchmark, an LLM agent first went through the artifacts, those released with GEPA and SkillOpt-Lite and those evolved in our grid, instruction by instruction and flagged the instructions that could explain the gain; we then checked each flagged instruction by hand against the artifacts and the benchmark's evaluation, and wrote the one that appeared most responsible as a single rule before running the test. The rule was not revised after seeing results. The rule is appended to the unevolved baseline, which is otherwise unchanged. Because the rule is chosen after inspecting the artifacts, the test shows that one sentence is *sufficient* to obtain the gain, not that the artifacts contain nothing else. Table 10 lists the rules.

*Repairs.* A repair either restores the benchmark's official evaluation where the method's port departed from it, or replaces surface-form matching with grading by meaning (Table 9). Semantic grading counts an answer as correct if it matches exactly or if two LLM judges (DeepSeek-V4-Flash and GPT-5.6 Luna), shown only the question, the reference answer, and the extracted answer, both accept it. We found no shortcut in the released evaluations of ALFWorld, PUPA, HoVer, and IFBench, so they have no repaired results.

## E.2 Per-Benchmark Results

*Reading the figures.* In Figure 3, each point is one model–benchmark pair and the horizontal axis is the gain under the evaluation released with the method; QA benchmarks (†) are OfficeQA, HotpotQA, SearchQA, and DocVQA, and math benchmarks (‡) are AIME 2025, LiveMathematicianBench, and LiveBench-Math. Benchmarks with no shortcut found have no repaired result and appear only in panel (a), except IFBench, which has no one-rule result either. Figure 8 expands Figure 3 into one bar per model and benchmark,

| Benchmark | Rule appended to the baseline |
|---|---|
| SpreadsheetBench | Write computed literal values into the answer cells, never Excel formula strings: compute the result in Python and write the resulting number or text. |
| AIME 2025 | Provide final answers exactly in the requested format: single integer, fraction in lowest terms, or exact closed-form expressions without additional commentary or formatting. |
| LiveMathematicianBench | If one option states that one of the remaining options is correct but a stronger result can be proven, choose that option. |
| DocVQA | Inside `<answer>`, output only the minimal span exactly as printed in the document, with no full sentence, field label, unit, qualifier, or trailing period. |
| HotpotQA | The answer is scored by exact match against a short reference answer. Output only the minimal answer span, with no full sentence, explanation or qualifiers. |
| OfficeQA | Put only the bare value inside `<answer>`: no units, currency signs, %, or explanatory words, unless the question explicitly asks for them. Keep the requested decimal places. |
| SearchQA | Inside the answer tag, write only the shortest span that exactly answers the question: no parenthetical clarifications, translations, or acronym expansions, no quotes, and no trailing period. |
| PUPA | Never copy user-specific details into the external request: replace them with generic descriptions or placeholders, while keeping the task fully specified. |
| ALFWorld | Search systematically: first look at every open surface in numbered order, then open closed receptacles; never revisit an empty location. You can hold only one object: never pick up anything that is not your target. For 'two X' tasks, place the first one at the destination before searching for the second. |
| HoVer | Target the Wikipedia articles that have not been retrieved yet: name the specific entities whose own pages are still missing, instead of restating the claim. |

**Table 10** Rules used in the one-rule test, abridged; the full text is released with our code. The ALFWorld rule is a short search procedure rather than a statement about evaluation. LiveBench-Math and IFBench have no rule (Appendix E.2).

and Table 11 lists the scores behind both figures. Where the evaluation admits a shortcut, the reported gain (gray) is matched by the one-rule baseline (black tick) and shrinks to near zero under the repaired evaluation (green). The dotted line marks the headroom left after repair, one hundred minus the repaired baseline: on several benchmarks the reported gain exceeds this headroom, so it cannot reflect a gain in the capability the benchmark intends to measure. The one-rule test does not reproduce every gain; the paragraphs below discuss the exceptions.

*Gains that do not come from a shortcut.* On HoVer, whose gains come from multi-hop retrieval strategies, one sentence recovers only 43% and 41% of the gain for GPT-5.4 nano and GPT-5.6 Luna. For DeepSeek-V4-Flash, a single retrieval rule, to name the entities whose pages are still missing, recovers 91%, so for this model the retrieval improvement itself fits in one sentence. On HotpotQA, a gain of 5 to 8 pp survives the repair; it also comes from better retrieval (Appendix E.4).

*Rules that the model does not follow.* More often, the gain does come from a shortcut but the model does not act on a single sentence. On ALFWorld with DeepSeek-V4-Flash, the rule recovers only 30% of the gain because the model does not follow the search procedure, while the longer evolved artifact does change its behavior. On LiveMathematicianBench, with the rule the models choose the stronger-result option on 4%, 54%, and 27% of the questions where it is correct, against 67%, 86%, and 92% with the evolved skill, so the rule recovers 0% to 51% of the gain, while the repaired gain stays near zero for both. The skill differs from the rule in four ways: it gives a reason, that the concrete options are "almost always deliberately understated"; it makes the stronger-result option the default unless a concrete option can be argued to be sharp; it lists the wordings under which that option appears; and it repeats the instruction just before the final answer. Evolution thus optimizes not only which shortcut to take but also how to state it so that the model takes it.

*Rules that score differently from the evolved artifact.* On AIME 2025 with GPT-5.6 Luna, the rule scores higher than the evolved artifact (69.3 versus 51.0): both adapt to a parser that credits only a bare integer, but

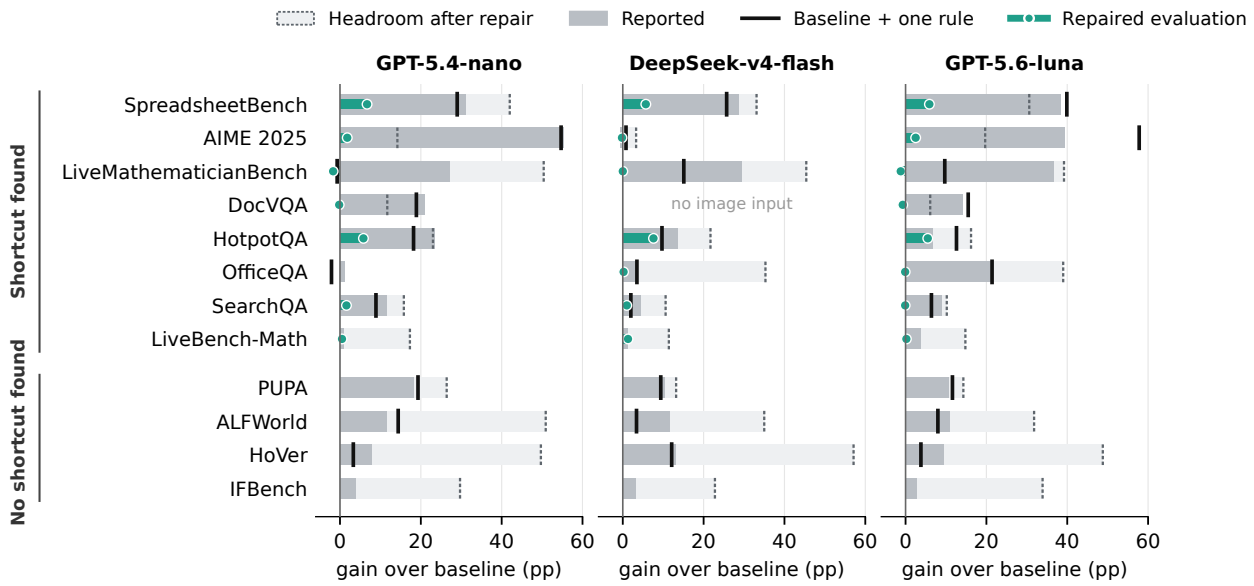



**Figure 8** Per-benchmark gain audit. Gray: gain reported under the released evaluation. Black tick: unevolved baseline plus one rule. Green: the same artifacts under the repaired evaluation. Light bar with dotted end: headroom after repair (100 minus the repaired baseline). SpreadsheetBench gains and the rule use the released grader; its repaired values use LibreOffice recalculation.

the rule removes more of the answers the parser fails to read. On SearchQA the rule recovers 45% to 78% of the gain, but under semantic grading it lowers the score of GPT-5.4 nano by 2.0 pp relative to the unevolved baseline: semantic grading credits hedged answers such as "George Calvert (the 1st Lord Baltimore)", whereas a model told to write only the shortest span sometimes keeps the wrong part.

*Benchmarks without a one-rule result.* IFBench has neither a one-rule result nor a repaired result, for two reasons. Its released evaluation checks each instruction with code, and grading the same outputs more strictly, on the model's original response only, makes the gains larger rather than smaller, so we found no shortcut to repair and no requirement of the evaluation that a single rule could state. Its gains are also small, 2.7 to 3.8 pp, comparable to the single-run standard deviation measured in Section 4.1, so the share of such a gain recovered by one rule could not be interpreted. We keep IFBench in the figure as an example of the small improvements that remain when a gain does not come from a shortcut. LiveBench-Math has no one-rule result because its gains are too small: under the released evaluation they are 0.9, 1.3, and 3.9 pp, and only the last has a 95% confidence interval that excludes zero, so the share recovered by one rule could not be interpreted.

## E.3 Re-evolving Under Repaired Evaluations

Re-scoring released artifacts leaves open whether evolution, given a repaired evaluation from the start, would find gains of its own. We therefore re-run evolution with the repaired evaluation used for training feedback, validation, and testing (Table 12). On OfficeQA and SpreadsheetBench, the re-evolved text-artifact methods gain at most 1.9 pp over the repaired baseline, except for SkillOpt-Lite with DeepSeek-V4-Flash on SpreadsheetBench (+5.8 pp), whose artifacts avoid spreadsheet functions the recalculation engine does not support. On these two benchmarks we re-evolve GEPA and SkillOpt-Lite only with DeepSeek-V4-Flash and GPT-5.6 Luna, which already show that the large gains under the released evaluations do not return. On LiveMathematicianBench, the leaked option can only be balanced during evolution, so that it is correct half the time where it appears (Appendix E.5); this halves rather than removes the shortcut. On the full balanced test set, DeepSeek-V4-Flash gains 5.2 pp with GEPA and 3.8 pp with SkillOpt-Lite, but the gains come entirely from questions that contain the option, and on the questions without it no configuration

| | | | Released evaluation | | | Baseline + one rule | | Repaired evaluation | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Benchmark | Model | Items×reps | Base | Evolved | Gain | Score | Gain | Base | Evolved | Gain |
| SpreadsheetBench | nano | 281×16 | 33.5 | 64.6 | **+31.1** | 62.5 | **+29.0** | 58.0 | 64.7 | **+6.7** |
| | DeepSeek | | 43.4 | 72.0 | **+28.6** | 69.1 | **+25.7** | 66.9 | 72.6 | **+5.7** |
| | Luna | | 35.3 | 73.7 | **+38.4** | 75.2 | **+39.9** | 69.4 | 75.3 | **+5.9** |
| AIME 2025 | nano | 30×4 | 32.0 | 87.3 | **+55.3** | 86.7 | **+54.7** | 85.8 | 87.7 | **+1.8** |
| | DeepSeek | | 96.5 | 96.0 | −0.5 | 97.3 | +0.8 | 96.7 | 96.5 | −0.2 |
| | Luna | | 11.5 | 51.0 | **+39.5** | 69.3 | **+57.8** | 80.3 | 82.8 | +2.5 |
| LiveMathematicianBench | nano | 106×4 | 29.0 | 56.1 | **+27.1** | 28.3 | −0.7 | 49.6 | 47.9 | −1.7 |
| | DeepSeek | | 38.9 | 68.4 | **+29.5** | 54.0 | **+15.1** | 54.6 | 54.6 | +0.0 |
| | Luna | | 37.0 | 73.8 | **+36.8** | 46.7 | **+9.7** | 60.8 | 59.6 | −1.2 |
| DocVQA | nano | 374×4 | 61.0 | 82.0 | **+21.0** | 79.9 | **+18.9** | 88.3 | 88.1 | −0.2 |
| | DeepSeek | | *not run: no image input* | | | | | | | |
| | Luna | | 73.1 | 87.2 | **+14.1** | 88.6 | **+15.5** | 93.9 | 93.2 | −0.7 |
| HotpotQA | nano | 300×4 | 40.2 | 63.6 | **+23.3** | 58.4 | **+18.2** | 77.0 | 82.8 | **+5.8** |
| | DeepSeek | | 57.4 | 71.0 | **+13.6** | 67.1 | **+9.7** | 78.3 | 85.9 | **+7.6** |
| | Luna | | 51.2 | 57.9 | **+6.7** | 63.8 | **+12.6** | 83.8 | 89.2 | **+5.5** |
| OfficeQA | nano | 148×16 | 11.4 | 12.3 | +1.0 | 9.2 | **−2.1** | – | – | – |
| | DeepSeek | | 55.3 | 58.5 | **+3.3** | 58.8 | **+3.5** | 64.7 | 64.9 | +0.2 |
| | Luna | | 35.5 | 56.6 | **+21.1** | 56.9 | **+21.4** | 61.0 | 60.9 | −0.1 |
| SearchQA | nano | 1400×2 | 52.0 | 63.5 | **+11.5** | 60.9 | **+8.9** | 84.2 | 85.8 | **+1.6** |
| | DeepSeek | | 69.3 | 73.7 | **+4.4** | 71.3 | **+2.0** | 89.4 | 90.4 | **+1.0** |
| | Luna | | 63.9 | 73.0 | **+9.1** | 70.2 | **+6.4** | 89.8 | 89.6 | −0.1 |
| LiveBench-Math | nano | 126×4 | 73.8 | 74.7 | +0.9 | – | – | 82.7 | 83.2 | +0.5 |
| | DeepSeek | | 80.5 | 81.8 | +1.3 | – | – | 88.6 | 90.0 | +1.3 |
| | Luna | | 72.9 | 76.8 | **+3.9** | – | – | 85.2 | 85.4 | +0.2 |
| PUPA | nano | 221×4 | 73.6 | 91.9 | **+18.3** | 92.9 | **+19.3** | – | – | – |
| | DeepSeek | | 86.8 | 97.2 | **+10.4** | 96.2 | **+9.4** | – | – | – |
| | Luna | | 85.7 | 96.4 | **+10.7** | 97.3 | **+11.6** | – | – | – |
| ALFWorld | nano | 67×16 | 49.1 | 60.6 | **+11.5** | 63.4 | **+14.4** | – | – | – |
| | DeepSeek | | 65.0 | 76.5 | **+11.5** | 68.4 | **+3.4** | – | – | – |
| | Luna | | 68.2 | 79.1 | **+10.9** | 76.2 | **+8.0** | – | – | – |
| HoVer | nano | 300×4 | 50.3 | 58.2 | **+7.8** | 53.7 | **+3.3** | – | – | – |
| | DeepSeek | | 42.9 | 56.2 | **+13.2** | 55.0 | **+12.1** | – | – | – |
| | Luna | | 51.2 | 60.7 | **+9.4** | 55.1 | +3.8 | – | – | – |
| IFBench | nano | 294×4 | 70.3 | 74.1 | **+3.8** | – | – | – | – | – |
| | DeepSeek | | 77.2 | 80.4 | **+3.1** | – | – | – | – | – |
| | Luna | | 66.1 | 68.8 | +2.7 | – | – | – | – | – |

**Table 11** Scores behind the gain audit, in pp, averaged over evaluation repetitions. Released: the evaluation released with the method; baseline + one rule: scored under the released evaluation; repaired: after the repair in Table 9. Bold: 95% confidence interval excludes zero. –: no repair, no one-rule result, or not run (Appendix E). On LiveMathematicianBench, repaired scores use the 60 of the 106 test questions without the leaked option.

improves significantly. HotpotQA is the exception: GEPA improves retrieval for all three models under either evaluation (Appendix E.4).

SpreadsheetBench also shows why adapting to a shortcut is learned only when the gain from it is large. Writing values instead of formulas is worth about 30 pp under the released grader, and all eight runs of SkillOpt-Lite with GPT-5.4 nano discover it. Under the repaired grader the same rule is still worth about 6 pp, because some formulas use functions LibreOffice cannot evaluate, but on a 39-task validation set this is two or three tasks, comparable to run-to-run noise. Only 2 of the 12 artifacts re-evolved under the repaired grader keep the behavior, although half of them state the rule in their text.

| Benchmark | Method | nano | DeepSeek | Luna |
|---|---|---|---|---|
| OfficeQA | GEPA | – | −0.6 | +1.4 |
| | SkillOpt-Lite | – | +0.9 | +0.3 |
| SpreadsheetBench | GEPA | – | **+1.9** | +0.6 |
| | SkillOpt-Lite | – | **+5.8** | +0.0 |
| HotpotQA | GEPA | **+8.6** | **+8.9** | **+3.5** |
| | SkillOpt-Lite | −0.8 | **+2.2** | +0.0 |
| LiveMathematicianBench | GEPA | +3.1 | −4.3 | −3.5 |
| | SkillOpt-Lite | +1.2 | −3.9 | −3.5 |
| ALFWorld | HarnessLens | +3.2 | −0.4 | −0.8 |
| | Meta-Harness | −1.6 | **+30.3** | **+16.4** |
| | AHE | −0.1 | **+31.7** | **+23.6** |
| SpreadsheetBench | HarnessLens | +0.5 | −0.3 | +1.6 |
| | Meta-Harness | +0.6 | **+6.3** | +0.1 |
| | AHE | −3.6 | **+7.6** | **+8.8** |
| OfficeQA | HarnessLens | −0.3 | −0.3 | −1.1 |
| | Meta-Harness | +3.7 | **+2.8** | +0.3 |
| | AHE | +2.8 | **+3.7** | +1.2 |

**Table 12** Gain (pp) of artifacts evolved under the repaired evaluation over the same-model baseline under that evaluation. Bold: 95% confidence interval excludes zero. LiveMathematicianBench is scored on questions without the leaked option. No shortcut was found in the ALFWorld evaluation, so its harness rows use the released evaluation.

*Harness methods.* We also re-evolve HarnessLens, Meta-Harness, and AHE under the repaired evaluations on ALFWorld, SpreadsheetBench, and OfficeQA with GPT-5.4 nano, DeepSeek-V4-Flash, and GPT-5.6 Luna, with 8 evolution runs and 8 evaluations each (Table 12). On SpreadsheetBench and OfficeQA the gains are at most 9 pp, and HarnessLens returns its initial harness in 47 of 71 runs. Two larger gains need care. On ALFWorld, Meta-Harness and AHE improve DeepSeek-V4-Flash and GPT-5.6 Luna by 16 to 32 pp, because the initial agent keeps only the last two steps of its history; one run that changes nothing but this setting gains 29 pp. On SpreadsheetBench, AHE gains 11 to 13 pp for these two models on the 151 of 281 tasks whose instructions ask for a formula, because the grader compares values and writing the computed value passes, but only 3 to 4 pp on the other tasks. Both gains come from a single change that helps most tasks at once, the kind of gain that evolution can identify.

## E.4 Headroom After Repair

Repairing an evaluation also shows how much room a benchmark leaves for evolution. Terminal-Bench, which we did not audit, is included under its released evaluation.

*Benchmarks that no longer measure their target.* On two benchmarks, the score can be earned without the capability the benchmark is meant to test.

- **SearchQA: closed-book knowledge instead of retrieval.** The benchmark is meant to test answering from retrieved search snippets, but the models answer as well without them. Under semantic grading, removing all snippets changes the score from 84.2 to 81.9 for GPT-5.4 nano, from 89.4 to 90.2 for DeepSeek-V4-Flash, and from 89.8 to 90.7 for GPT-5.6 Luna. Snippets help only on the 126 of 1,400 questions whose snippets come from question-bank sites and contain both the clue and the answer. The remaining errors say little about the models: of 100 residual errors, which we attributed with an LLM and spot-checked by hand, 36 stem from wrong or ambiguous gold answers, such as Jeopardy clues that lost their category, or from answers that are in fact acceptable. The benchmark cannot even separate DeepSeek-V4-Flash from GPT-5.6 Luna ($-0.4$ pp, 95% CI $[-1.5, 0.8]$), and the released skill adds at most 1.6 pp.
- **LiveMathematicianBench: one option is always correct.** The benchmark turns theorems from recent arXiv papers into five-option questions that ask for the strongest statement that can be proved. Some

questions include the option "One of the remaining options is correct, but a stronger result can be proven." Across the 177 questions of the released splits, this option appears in 85 and is the correct answer in all 85; it appears in 46 of the 106 test questions. For example, one test question is built from the theorem that every flow-admissible, cyclically 5-edge-connected, cubic signed graph has a nowhere-zero 6-flow. The theorem itself is not among the options: one option weakens its conclusion to a $\mathbb{Z}_6$-flow, two strengthen it to cyclic 4-edge-connectivity or a 5-flow, one drops flow-admissibility, and the correct answer is the stronger-result option. A model that always picks this option answers all 46 such test questions correctly without reading the mathematics, and the released skill's entire gain comes from these questions. On the 60 test questions without the option, the baselines are 50–61, yet no artifact improves on them; with 60 questions the confidence interval is about ±5 pp.

*Little headroom.* On several benchmarks the baseline, or the baseline plus one sentence, is already close to full marks once answers are graded correctly, so evolution has little left to gain.

- **DocVQA.** With semantic grading the baseline reaches 88–94, and the released skill adds nothing.
- **AIME 2025.** DeepSeek-V4-Flash solves 96.5% of attempts without evolution. For all three models the repaired gains stay within 2.5 pp, comparable to the confidence interval of a 30-problem test set.
- **PUPA.** The evolved artifacts reach 96–97 for DeepSeek-V4-Flash and GPT-5.6 Luna, and one sentence added to the baseline reaches the same level.
- **ALFWorld, for strong models only.** Gemini-3.7-flash, an additional model outside our main grid, starts at 96.3, and the best of our four text-artifact methods reaches 97.0. In the SkillOpt-Lite paper [31], both SkillOpt-Lite and its strongest competitor score 100.0 with GPT-5.4-mini, GPT-5.4, and GPT-5.5, so the benchmark can no longer rank methods for these models.

*Headroom without an identifiable gain.* On other benchmarks much room is left, yet evolution gains little or nothing that the evaluation can identify.

- **Terminal-Bench.** It leaves the most room, 28 pp for DeepSeek-V4-Flash and 46 for GPT-5.6 Luna, yet the best method gains only 1.0 and 3.6 pp, below the standard deviation of a single run (Section 4.1).
- **OfficeQA.** Under its official evaluation it leaves 35–39 pp, but neither the released artifacts nor re-evolved text-artifact methods gain anything, and re-evolved harness methods gain at most 3.7 pp (Table 12).
- **LiveBench-Math.** It leaves 11–17 pp after repair, yet the repaired gains are only 0.2 to 1.3 pp.
- **IFBench.** Its gains of 2.7–3.8 pp are of the same order as its evaluation noise (Appendix E.2).

*Headroom with a real gain.* Three benchmarks keep a gain that the repair does not remove, and in each case its source can be named. On HotpotQA, GEPA keeps 5.5–7.6 pp under semantic grading, and re-evolution traces the gain to retrieval: recall of supporting articles rises by about 10 pp. On SpreadsheetBench, 5.7–6.7 pp survive LibreOffice recalculation, but the re-evolved artifacts that keep a gain do so by adapting to the recalculation engine, writing values and avoiding unsupported functions. On HoVer, gains of 8–13 pp come from retrieval strategies. These are the small, specific improvements whose identification Section 5 examines.

## E.5 Scope of the Audit

The audit tests whether the largest reported gains need more than a single rule; it does not show that an evolved artifact contains nothing else of value. Because we wrote each rule after inspecting the artifacts, the test shows that one rule is enough to reproduce a gain, not that the gain requires this rule, and HoVer shows that the test fails when a gain is spread over several strategies. The artifacts released with GEPA were evolved for GPT-4.1-mini, so their gains on our models are smaller than those in the original paper; we therefore compare, on the same model, the gain under the released evaluation with the one-rule result and with the repaired gain. Finally, a repair can be incomplete. For re-evolution on LiveMathematicianBench, we balance

the leaked option: in as many questions as it is correct in, we replace a wrong option with its text, so that it is correct only half the time where it appears. Half is still more often than the models answer these questions correctly (about 34% for DeepSeek-V4-Flash), so choosing the option remains profitable, and balancing only halves its benefit. We therefore score capability only on the questions that do not contain the option.

# F Candidate Pool and Selection

## F.1 Candidate Pools, Timeouts, and Evaluation Protocols

Sections 5.1 and 5.2 evaluate two fixed pools of $n$=100 skills designed for the two stages of candidate-level evaluation: the **Proposal Pool**, which tests out-of-sample skill transfer from 10 training tasks to 10 disjoint test tasks, and the **Selection Pool**, which benchmarks selectors and trajectory judges across 11 difficulty-stratified tasks.

*Proposal Pool (Section 5.1).* To evaluate proposal quality without selection in the loop, an offline SkillOpt-Lite-based proposer using DeepSeek-V4-Flash generates $n$=100 standalone skills from 78 execution trajectories across 10 Terminal-Bench v2.1 training tasks (kv-store-grpc, pypi-server, sqlite-db-truncate, path-tracing, filter-js-from-html, torch-tensor-parallelism, dna-insert, mcmc-sampling-stan, bn-fit-modify, tune-mjcf) following the protocol below. Each skill is then evaluated 16 times on each of 10 disjoint Terminal-Bench v2.1 test tasks (configure-git-webserver, break-filter-js-from-html, query-optimize, dna-assembly, rstan-to-pystan, vulnerable-secret, mailman, db-wal-recovery, torch-pipeline-parallelism, path-tracing-reverse), totaling $100 \times 10 \times 16 = 16{,}000$ solver rollouts with binary whole-task rewards under native task verifiers.

*Original timeout vs. extended execution budgets.* All 16,000 rollouts are executed with a 2-hour (7,200 s) agent budget. Because Terminal-Bench tasks normally enforce shorter native timeouts (900 s, 1,200 s, or 1,800 s per task), we also score five timeout cutoffs on the exact same set of finished trajectories by marking a rollout as failed whenever agent elapsed time exceeds the cutoff: the task's *original* timeout, 2×, 3×, 4×, and the full *2h* budget. Figure 9 shows Quality@$K$ and its matched equal-quality null across all five cutoffs under our primary empirical-Bayes estimator ($\kappa_{\max}$=100; Appendix F.2): the pooled $K$=1 success rate rises from 39.69% under original timeouts through 45.15% (2×), 48.14% (3×), and 49.13% (4×) to 49.23% at 2 h, while the $K$=100 excess above the matched null remains cleanly positive across every cutoff at +3.33, +3.22, +3.28, +3.30, and +3.41 pp (with non-overlapping 95% bands from $K \geq 5$ under original–4× timeouts and from $K \geq 4$ at 2 h). Figure 10 displays all 10 individual test tasks under both original and 2 h timeouts, showing that tasks whose multi-step compilation or rendering exceeds native limits (such as `path-tracing-reverse`, rising from 4.8% at original timeout to 41.4% at 2 h, and `rstan-to-pystan`, rising from 23.2% to 62.4%) explain most of the level shift between cutoffs, while near-homogeneous tasks track their matched equal-quality nulls closely within posterior bands.

*Selection Pool (Section 5.2).* To benchmark selection operators and LLM trajectory judges across diverse candidates and task difficulty levels, Section 5.2 evaluates the Selection Pool: $n$=100 skills generated from execution trajectories across six Terminal-Bench v2.1 tasks and three DeepSWE tasks, rolled out on 11 Terminal-Bench tasks spanning five baseline difficulty levels (six source tasks and five held out from skill generation; listed in Appendix C). Every artifact–task pair receives 8 rollouts, extended to 15 for the ten highest-scoring artifacts per task, and the unevolved baseline receives 12 rollouts per task. Its per-task selection curves are reported in Appendix F.4.

### F.1.1 Proposer Research Protocol

To generate $n$=100 distinct offline candidate skills from the training trajectories without an outer selection loop, the proposer launches parallel research sub-agents that inspect failure mechanisms on the training corpus and verify candidate fixes with small local probes. Below is the research and skill-authoring prompt used to synthesize the $n$=100 standalone skills in the Proposal Pool.

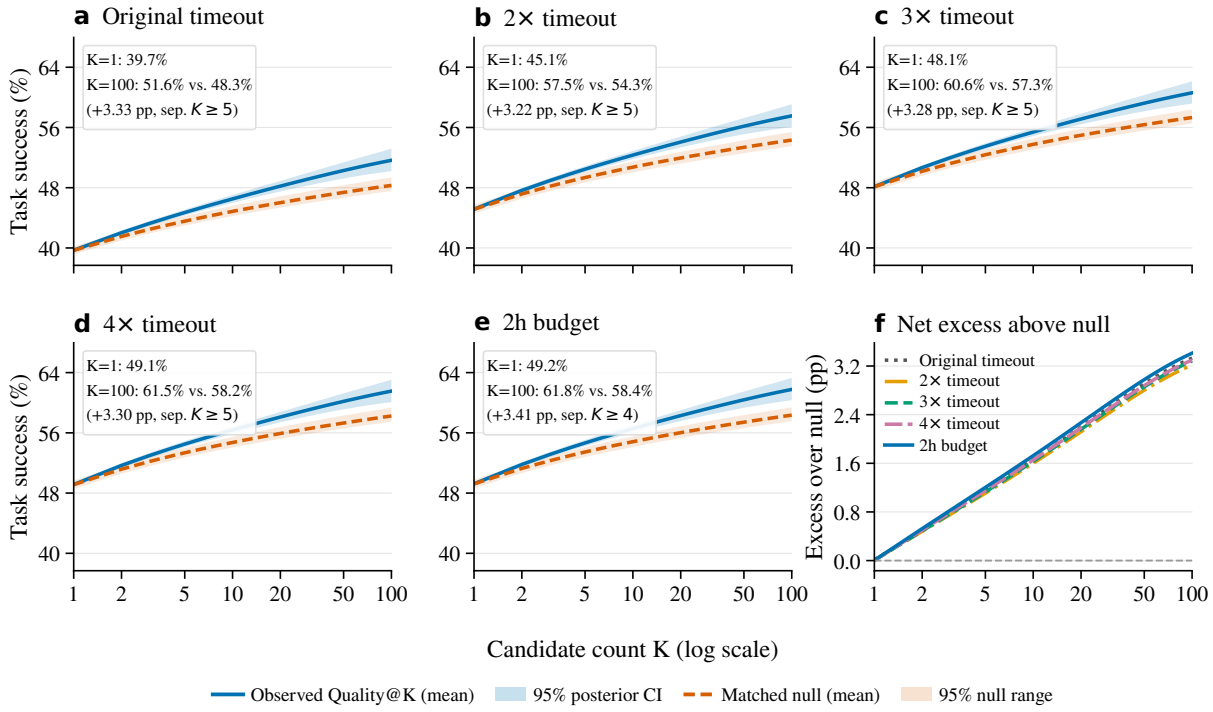



**Figure 9** Effect of agent timeout cutoff on Quality@$K$ vs. the matched equal-quality null on the Proposal Pool (10 disjoint Terminal-Bench v2.1 test tasks, $n$=100 skills, $\kappa_{\max}$=100 empirical-Bayes estimator). Panels (a)–(e) evaluate the same 16,000 trajectories under progressively relaxed cutoffs: original per-task timeout (900–1800 s), 2×, 3×, 4×, and the full 2-hour budget. Panel (f) overlays the net excess above the matched null across all five cutoffs.

**Research & Skill Authoring Guidance**

Your goal is to investigate the training corpus in `workspace/.skillopt/corpus/`, uncover the concrete mechanisms behind task successes and failures, verify them, and translate them into one hundred complete, standalone actionable skills at `workspace/artifacts/001/SKILL.md` through `workspace/artifacts/100/SKILL.md` (each with its own `report.md`), backed by an auditable evidence report in `workspace/report.md`.

There is no fixed reading quota, edit limit, or rigid output template. Read, search, write scripts, and run small executable probes in `workspace/` as much as the evidence warrants.

**1. Where to look: find the "moment of wrong commitment"**

Because task outcomes in the training pool are unbalanced, do not simply group failed runs by task name. Study **both** tasks with failures (comparing failed runs against passing runs on the same task to see the full workflow past early setup/install crashes) **and** pass-only or fail-only tasks (extracting how passing runs diagnose and recover from mid-run obstacles, or using local probes in `workspace/` when no passing run exists). Focus on the critical decision points where an agent locks into a flawed path:

- **State and deliverable lifecycle:** Check whether agents mutate raw inputs in-place before making a safe backup, spend the entire budget in scratch experiments without maintaining a valid working deliverable at the target path, or accidentally undo working state and stop required services during final cleanup.
- **Unverified references & circular self-tests:** Agents often believe a flawed self-test—or worse, **modify a correct implementation to match an inaccurate reference**. Check whether the agent trusted a library helper or reference script as an oracle without first verifying that its implicit behavior matches the task contract, or wrote a checker that simply reused the generator's own assumptions and offsets.
- **Premature scaling without small probes:** Check where agents launch expensive full-scale runs before using a tiny run to verify interfaces, output structure, and how runtime, memory, or output sizes scale.
- **Stopping at surface gains:** In optimization or reconstruction tasks, check where agents stop after an initial improvement over a naive baseline or a rough approximation, versus how successful runs use concrete diagnostics (profiles, execution plans, structural parsers, or quantitative residuals) to locate and remove the remaining bottleneck.

**2. Evidence discipline: verify mechanisms, do not just tell stories**

A failed run and a plausible narrative do not establish a cause, and a passing run's own explanation or code comment is only a hypothesis.

- **Inspect middle logs (beware output truncation):** Long command outputs in your environment are truncated in the middle. Never assume a broad dump showed the full trace. Before attributing a failure to downstream visibility or configuration, inspect the middle of the transcript to verify whether upstream steps (such as compilation, linking, or installation) actually succeeded.

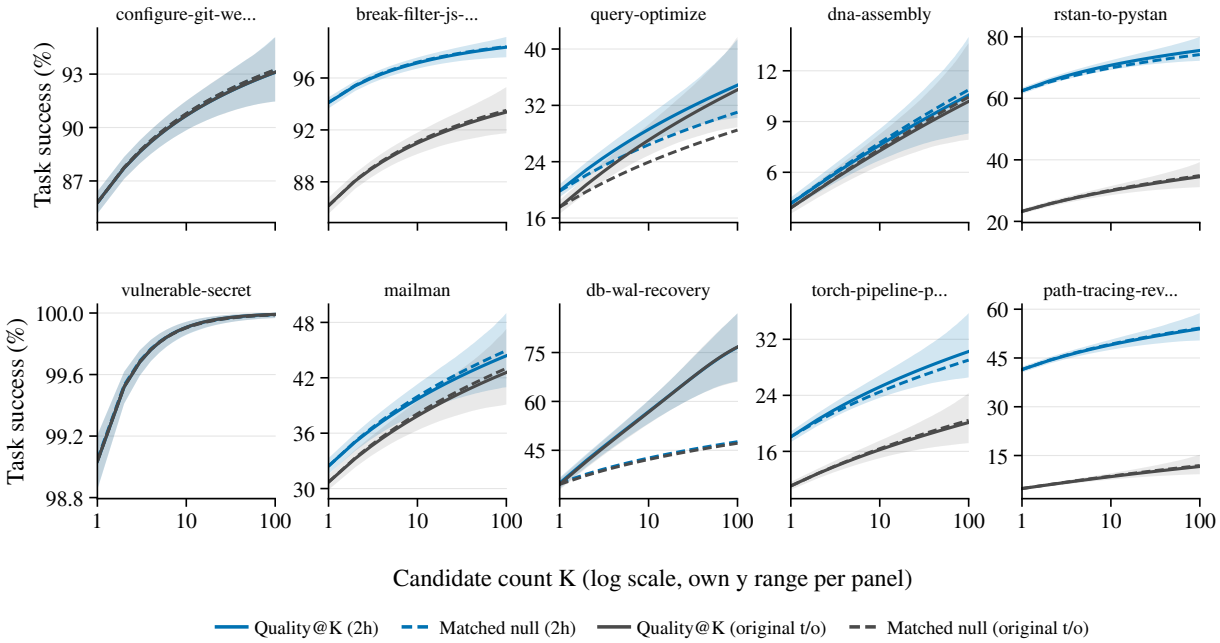



**Figure 10** Per-task Quality@$K$ (solid) and matched equal-quality null (dashed) across all 10 disjoint Terminal-Bench v2.1 test tasks of the Proposal Pool under the 2-hour budget (blue) and original per-task timeout (charcoal), using $\kappa_{\max}$=100 empirical Bayes; each panel uses its own vertical scale.

- **Account for exact assertions:** Repeated runs of a task are not independent tasks, and failing the same test function does not mean failing for the same reason. Check the actual verifier output per run to distinguish early format or missing-file failures from downstream numerical or behavioral failures.
- **Verify technical claims with small probes:** Agents in passing runs sometimes rely on lucky approximations or state inaccurate formulas, byte layouts, or API rules. Before encoding a specific mathematical formula, binary/database record format, parser behavior, or library rule into the skill, verify it with a small executable probe in `workspace/` or against local documentation.

**3. Authoring `skill.md`: concrete actions, de-anchored from training literals**

High-level principles are research lenses for **you (the proposer)**—they are **not** finished instructions for the solver. At the same time, a skill that is over-fitted to one training problem's surface details will fail to trigger on neighboring tasks.

- **Strip training-instance literals, keep technical depth:** Never leave training-specific filenames, JSON field names, sample IDs, or narrow one-task formulas as the trigger or example in `skill.md`. Broaden the trigger to the general technical scenario (e.g., "corrupt database/file recovery" rather than only "truncated file"; "structured sequence/protocol assembly" rather than one specific splicing formula), while keeping the concrete diagnostic commands, byte/format rules, and verification recipes sharp.
- **Turn principles into decision-point checks:** For each workflow you write, give the solver a complete, executable playbook:
  1. **When it triggers** and **what to protect or establish first** (e.g., full directory backup before first tool contact; tiny-budget schema and memory-growth check; saving an initial compilable baseline deliverable).
  2. **How to calibrate the reference check:** Teach the solver to verify any third-party reference API on a tiny 2–3 element hand-checked input before trusting it as an oracle, and to validate structured/composite outputs with an independent parser and end-to-end simulation (without reusing builder index offsets or dropping required protocol prefixes to guess grader behavior).
  3. **How to diagnose the next bottleneck:** Provide concrete diagnostic loops (inspecting execution plans/profiles, reducing candidate sets before heavy computation, driving reconstruction via quantitative residuals, and checking multi-component service chains from a fresh client).

**4. Outputs and honest boundaries**

All training evidence is mounted read-only at `workspace/.skillopt/corpus/` (see `corpus/README.md` for optional helpers). Use `workspace/` for your own scripts, notes, and executable probes. Do not retrieve external benchmark solutions or test-set material.

- Save exactly one hundred complete standalone skills at the artifact paths above, each free of training task names, sample IDs, and task-local literal constants.
- Save a concise, auditable report in `workspace/report.md` tracing the chain from **what you read** → **what you verified via probes** → **what you wrote into the skill** → **what counterexamples or open gaps remain**, citing exact message IDs, line ranges, and verifier assertions.

## F.2 Posterior Estimation of Quality@$K$

Each artifact's per-task quality is a latent success probability observed through a handful of binary evaluations (Appendix F.1). Raw success counts are too noisy to rank individual artifacts reliably, so we regularize with an empirical-Bayes hierarchical model and define two complementary estimators.

*Hierarchical model.* Let $n$ denote the pool size, $s_i$ and $n_i$ the number of successes and total evaluations for artifact $i$ on a given task, and $\theta_i \in [0,1]$ its latent success probability. All $n$ artifacts share a task-specific beta prior $\theta_i \sim \text{Beta}(a,b)$, where the mean $\mu = a/(a{+}b)$ and concentration $\kappa = a{+}b$ are estimated by cross-artifact method-of-moments from the pooled outcomes. In our primary estimator, $\kappa$ is bounded in $[0.25, \kappa_{\max}]$ with $\kappa_{\max}{=}100$ (already more than 6× the per-skill rollout budget $m{=}16$): whenever cross-artifact sample variance does not exceed pure binomial sampling noise, $\hat{\kappa}$ saturates at $\kappa_{\max}$ rather than collapsing deterministically to a zero-width point mass. For any finite $\kappa$, the posterior for each artifact is $\theta_i \mid s_i, n_i \sim \text{Beta}(a{+}s_i,\ b{+}n_i{-}s_i)$, with posterior mean $\bar{\theta}_i = (a{+}s_i)/(a{+}b{+}n_i)$. We evaluate this same $\kappa_{\max}{=}100$ moment estimator identically on both observed pools and simulated equal-quality null pools, and test sensitivity across $\kappa_{\max} \in \{30, 100, 300, \infty\}$ below.

*Posterior Oracle.* Quality@$K$ asks: if we draw $K$ artifacts uniformly from the pool, how good is the best one? Let $I_K$ denote a uniformly random size-$K$ subset of the $n$ artifacts (baseline excluded). The *posterior Oracle* integrates this quantity over the full posterior:

$$\widehat{\text{Quality@}K}_{\text{Oracle}} = 100\,\mathbb{E}_{\boldsymbol{\theta}|\mathcal{D}}\left[\mathbb{E}_{I_K}\left[\max_{i\in I_K}\theta_i\right]\right]. \tag{7}$$

We approximate the outer expectation with 8,000 posterior draws (256 draws per simulation inside each null dataset). Within each draw, the inner expectation is exact: sorting the $n$ sampled qualities in ascending order, the artifact at rank $r$ contributes with weight $\binom{r-1}{K-1}/\binom{n}{K}$ for $r \geq K$. Curves show draw means; shading denotes pointwise 95% credible intervals conditional on the fitted hyperparameters.

*Sensitivity to $\kappa_{\max}$.* Figure 11 examines sensitivity of the 2-hour Quality@$K$ curve on the Proposal Pool across $\kappa_{\max} \in \{30, 100, 300, \infty\}$. An overly tight cap such as $\kappa_{\max} = 30$ forces substantial prior dispersion even when an equal-quality null exhibits no overdispersion, inflating the 2 h matched null to 64.28% at $K{=}100$ so the +2.20 pp excess lies within overlapping bands. At $\kappa_{\max} = 100$, the 2 h curve reaches 61.77% vs. 58.36% null (+3.41 pp, separating cleanly for $K \geq 4$), and relaxing $\kappa_{\max}$ further to 300 or $\infty$ widens the excess smoothly to +4.75 pp ($K \geq 3$) and +6.42 pp ($K \geq 2$). Importantly, finite regularization at $\kappa_{\max} = 100$ avoids a subtle point-mass degeneracy that arises under $\kappa_{\max} = \infty$: on near-homogeneous tasks whose sample variance sits just below the binomial floor (e.g., `dna-assembly` or `mailman`), $\kappa_{\max} = \infty$ collapses the observed prior to $\delta_\mu$ (a flat curve) whereas roughly half of pure-binomial null simulations happen to draw positive finite-sample overdispersion and finite $\hat{\kappa}$, nudging the simulated null mean slightly above the collapsed observed curve in per-task views. Bounding $\kappa_{\max} = 100$ keeps null-like tasks centered inside their matched null bands while every $\kappa_{\max} \geq 100$ confirms the same positive headroom conclusion.

*Equal-quality null.* Both estimators rise with $K$ even when all artifacts are equally good, because selecting the luckiest among more candidates yields a higher expected score. To calibrate this statistical selection effect, we run 1,000 null simulations per task: all $n$ candidates are generated i.i.d. at the task-level pooled pass rate, and the exact same estimator (prior fitting and posterior Quality@$K$) is refit on each simulated pool. Each estimator has its own null curve (dashed in Figures 4a, 9, and 11). Null intervals describe central 95% simulation variability across equal-quality datasets rather than posterior uncertainty.

*Scope.* Quality@$K$ characterizes the finite pool of $n{=}100$ candidates, not the theoretical task optimum, and cross-task averages allow each task to select a different skill. For the Selection Pool (Section 5.2), quality ground truth uses a held-out rollout reference independent of each selector's observed rollouts (Appendix F.3).

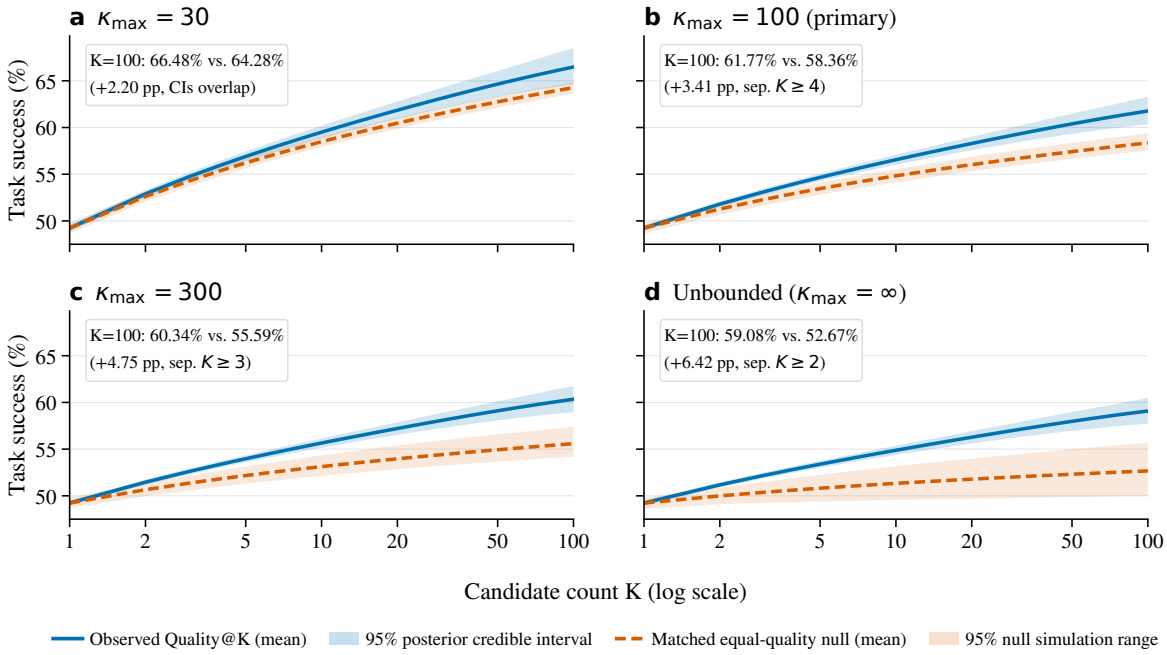



**Figure 11** Sensitivity of macro-averaged Quality@$K$ vs. matched equal-quality null under the 2-hour budget to the upper concentration cap $\kappa_{max} \in \{30, 100, 300, \infty\}$ of the empirical-Bayes Beta prior. A tight cap ($\kappa_{max}$=30) inflates dispersion in homogeneous null pools, whereas $\kappa_{max}$=100 (primary) through unbounded ($\kappa_{max}$=$\infty$) behave consistently and cleanly separate observed from null.

## F.3 Selection Operators and Evaluation Metrics

*Independent-reference protocol.* Each selector sees only the first rollout of each artifact (the native SkillOpt-Lite and GEPA rows, the first two). Artifact quality is evaluated against a held-out reference consisting of the remaining rollouts that the selector has not observed; the baseline reference likewise excludes its first two rollouts. This separation ensures that the quality ground truth used to assess selectors is independent of the information available to them. Artifact quality posteriors are fitted to this held-out evidence as described in Appendix F.2. All selectors share the same posterior draws, so selected quality is never larger than the oracle for the same draw, and all methods coincide at $K$=1.

*Native parent-selection ablation.* The SkillOpt-Lite and GEPA rows (marked † in Table 1) apply each method's own accept/reject logic to the first two rollout outcomes. SkillOpt-Lite replaces the current artifact only on strict improvement of the two-rollout mean; GEPA applies its frontier-pruning routine and samples a parent in proportion to surviving frequency. These rows use the same held-out quality reference as the other selectors.

*Trajectory judges.* DeepSeek-V4-Flash serves as the judge model on the 1,111 artifact–task pairs (including baselines) in the Terminal-Bench subset. Each call receives the task description, initial skill context, and the solver trajectory from the first rollout, without outcome labels. Pointwise scoring produces one score per pair (1,111 calls). Pairwise scoring evaluates 1,100 fixed pairs in both presentation orders (2,200 calls). The LLM judge (ACE) condition adapts the ACE Reflector's no-ground-truth template (error identification, root cause, corrective approach, and helpful/harmful/neutral tags), adding an explicit success-probability score. No ACE training loop or verifier feedback is used.

*Reported metrics.* *Pairwise accuracy*: posterior probability that the selector's chosen artifact has higher latent quality, with ties broken uniformly. *False acceptance*: among accepted baseline replacements, the posterior fraction with quality below baseline. *Improvement and regression rates*: posterior probability that

the selected artifact's quality exceeds or falls below baseline. Spearman correlations are within-task; tasks with constant selector scores are excluded.

*Can Higher-Resolution Selection Improve Evolution?* Our findings raise a natural question: if we reduce evaluation noise during evolution by averaging many validation rollouts at every selection step, does evolution accumulate larger gains? To test this directly, we run SkillOpt-Lite for up to eight rounds across 16 independent evolution runs, averaging 16 validation rollouts per candidate instead of one (and evaluating final artifacts with 32 independent test rollouts per task). Averaging 16 validation rollouts raises the mean final test score across runs from 61.4% to 64.4%, and raises the best evolved artifact's test score from 64.8% to 72.7%.

Inspecting the trajectories shows why: under single-rollout validation, an ordinary candidate occasionally draws an unusually high validation score by luck (e.g., 80%) and becomes the incumbent; subsequent candidates cannot beat this inflated threshold, stalling evolution within four rounds. Averaging 16 rollouts suppresses these false peaks so that accepted updates reflect real improvements, allowing the best evolved artifact to keep improving as more runs are launched. At the same time, 16× rollout validation is computationally expensive and makes greedy acceptance more conservative once a strong incumbent is found, motivating the cheaper, finer-grained behavioral diagnostics explored in Section 6.

### F.4 Per-Task Selection Curves

Figure 12 decomposes the aggregate selection curve of Figure 4b across each of the 11 Terminal-Bench v2.1 tasks in the Selection Pool (six source tasks T2, T5, T7, T9–T11 and five held-out tasks T1, T3–T4, T6, T8). Each panel plots selected-artifact success probability as a function of the candidate budget $K$ under the held-out rollout reference (Appendix F.3).

*Equivalent budget.* For each task and selector, the equivalent budget $K_{\text{eq}}$ is the smallest integer $K$ at which the oracle's posterior mean quality reaches or exceeds the selector's quality at $K$=100. It measures how many candidates an oracle selector would need to match a given selector after the latter has examined all $n$=100 candidates.

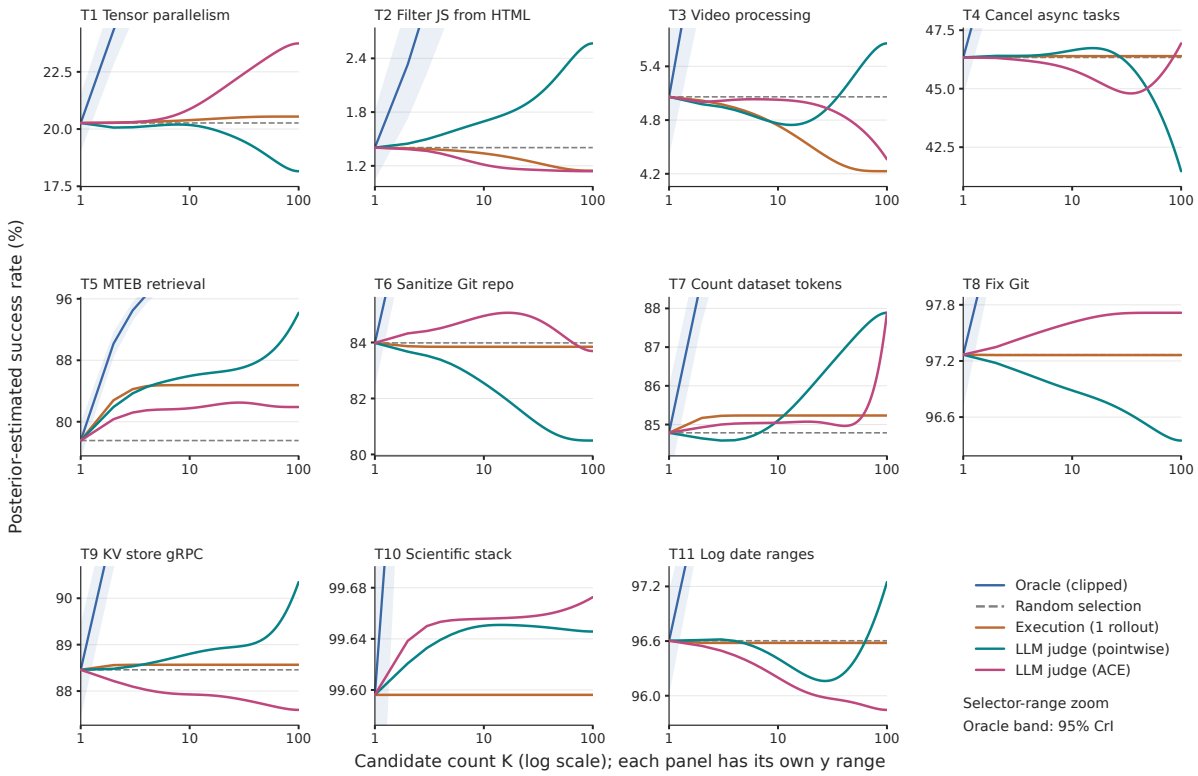



**Figure 12** Per-task selection curves underlying Figure 4b across the 11 tasks of the Selection Pool. Each panel shows selected-artifact success probability under the held-out quality reference (Appendix F.3); the oracle and its credible band are clipped to the selector range.

## G Trajectory Re-evaluation

We recover eight independent evolution runs per method and model for SkillOpt-Lite and GEPA with DeepSeek-V4-Flash and GPT-5.6 Luna. The 119 distinct within-run artifact versions are independently re-evaluated on the same 18 Terminal-Bench validation tasks, targeting eight rollouts per artifact–task pair (versus one during evolution).

*Variance estimation.* We estimate $\tau_t^2$ as the sample variance of the eight task-averaged re-evaluation means, subtracting the mean estimated variance of their evaluation noise. For task success estimate $\hat{p}_{rjt}$ from $n_{rjt}$ available repeats, the correction for run $r$ is $18^{-2}\sum_j \hat{p}_{rjt}(1-\hat{p}_{rjt})/(n_{rjt}-1)$ on the probability scale. Unlike Appendix B.2, which reports $\hat{\tau}$ itself and therefore sets negative estimates to zero, we retain negative corrected estimates here, because we compare $\hat{\tau}_t^2$ across rounds and clipping would bias the small early-round estimates upward.

## H Low-Variance Process Diagnostics

*Setup.* We reuse the trajectories of the variance experiments (Section 4) on DeepSeek-V4-Flash. The terminal reward is the benchmark's pass/fail verdict. The process metric is the fraction of check commands along the trajectory whose failure diagnostics are not masked, where a check command verifies intermediate state, such as running a test after an edit, and its diagnostics are masked if the agent cannot tell from the output that something failed. Both quantities are computed per trajectory, so their variances are directly comparable.

*Estimation.* For each configuration, we estimate the evaluation noise of each quantity with the within-artifact estimator of Appendix B.2, applied to the repeated evaluations of the same artifact.

For the baseline and all seven methods, the evaluation noise of the process metric is 13 to 57 times smaller than that of the terminal reward.

*Behavioral shifts visible only to the process metric.* The post-edit verification rate is the fraction of file modifications immediately followed by a test execution or sanity check. On DeepSeek-V4-Flash, Trace2Skill raises this rate by +8.73 pp over baseline ($z = 2.17$ against the evaluation noise of the difference). On GPT-5.6 Luna, AHE lowers it by 17.98 pp (variance of the difference 0.0018, $z = -4.19$). Neither shift is detectable in the terminal success rate of the same runs.

*Reading.* A process metric is a diagnostic, not a selector: a candidate that passes more checks need not solve more tasks, and a selector driven by such checks would face its own resolution question. Its value is that it makes behavioral change visible at resolutions where the terminal verdict is blind.

| Method | $\sigma^2_{\text{reward}}$ | $\sigma^2_{\text{behavior}}$ | Ratio |
|---|---|---|---|
| Baseline | 0.073 | 0.0013 | 57× |
| GEPA [2] | 0.074 | 0.0026 | 28× |
| HarnessLens [36] | 0.056 | 0.0044 | 13× |
| Meta-Harness [15] | 0.084 | 0.0015 | 57× |
| ReasoningBank [27] | 0.075 | 0.0021 | 35× |
| AHE [18] | 0.095 | 0.0022 | 44× |
| SkillOpt-Lite [31] | 0.071 | 0.0013 | 56× |
| Trace2Skill [24] | 0.077 | 0.0021 | 37× |

**Table 13** Variance of the evaluation noise: terminal pass/fail vs. a process metric (fraction of check commands without masked failure diagnostics), DeepSeek-V4-Flash.

# I Case Study

## I.1 Shareable Artifacts in the Candidate Pool

Section 5.1 shows that candidate skills vary substantially even when a single skill is shared across tasks. To illustrate what effective shared and task-specific skills look like in practice, this appendix prints three concrete skills from the Selection Pool (Appendix F.1): seed-027 and seed-099, tied for the best average gain across all 11 Terminal-Bench v2.1 tasks (+11.4 pp over baseline), and seed-058, which achieves a perfect score on mteb-retrieve in every rollout yet adds little when shared globally.

### I.1.1 Artifact seed-027 (best average gain, +11.4 pp)

**Software Engineering Skill**

- Read the task and relevant source before changing code.
- Preserve existing behavior while implementing the requested public API.
- Make a coherent implementation across affected modules.
- Verify normal behavior and edge cases with focused tests.
- Treat every precise validation, parsing, error, export, and format rule in the prompt as an explicit assertion.
- Run relevant existing tests and address regressions.

**Agent pitfall:** Substituting an idiomatic library name, conventional class name, or a generalized "reasonable" signature for the exact exported symbol, function signature, return shape, or error-message string that the task explicitly names.
**Trigger:** A prompt that gives a verbatim public identifier (class/function/entrypoint) or a literal error/status text that a hidden caller or assertion will match, while the agent invents a library-flavored alternative such as a generated servicer class name.
**Concrete correction:** Before coding, extract every verbatim identifier, signature, empty/zero-value return rule, and error fragment from the task; export and reference those exact strings on the public surface and error paths, and add a focused check asserting each named symbol and message exists exactly as spelled.
**Tiny verification:** For a spec that exports `Mean(xs []int) (float64, error)` and must return `(0, nil)` on an empty slice, expose exactly `Mean` (not `computeMean`) and return the documented `(0, nil)` instead of `(-1, err)`, then confirm with one empty-input call.

**Agent pitfall:** Improvising tokenization or embedding calls by hand (custom separators, toggling special-token flags, or a single generic task/prompt name applied to both query and corpus) instead of delegating to the named model/package and its documented encoding conventions.
**Trigger:** A data task that names a specific tokenizer, embedding model, or package and requires an exact numeric count or ranking derived from that model's own encoding behavior.
**Concrete correction:** Load the model only through the named package/entrypoint and let it both tokenize and encode every input; when the wrapper distinguishes query vs. passage or exposes special-token/clip parameters, set those explicitly per side rather than passing one default, and recompute the count or ranking a second time to confirm it is stable.
**Tiny verification:** Two short strings joined by `"\n"` and `" "` should tokenize to the same total under the named tokenizer's documented defaults; assert that changing the separator or the query/passage prompt flag alters the result, then keep the documented default so the final number matches the scorer's.

### I.1.2 Artifact seed-099 (tied best, +11.4 pp)

**Software Engineering Skill**

- Read the task and relevant source before changing code.
- Preserve existing behavior while implementing the requested public API.
- Make a coherent implementation across affected modules.
- Verify normal behavior and edge cases with focused tests.
- Treat every precise validation, parsing, error, export, and format rule in the prompt as an explicit assertion.
- Run relevant existing tests and address regressions.

**Agent pitfall:** Assuming an exact, machine-checked numeric result produced by a third-party text pipeline (a tokenizer count, an embedding similarity rank, a retrieved document) is invariant to how the input text is assembled, so you join source fields with an arbitrary separator or encode both the query and the reference set with the same default prompt setting.
**Trigger:** A script must output one exact number derived from feeding text to an encoder, and it builds that text by concatenating several source fields, or it calls a single encode pass for both a probe query and a corpus of documents.
**Concrete correction:** Replicate the source data-shape exactly: join fields with the same boundary character used in the corpus (a newline, not a space, can change the token count by one or more), and when the encoder wrapper is instruction-aware, call it once with the query prompt type for the probe and once with the passage prompt type for the documents so each side receives its real instruction prefix.
**Tiny verification:** `len(encode("a\nb"))` is not required to equal `len(encode("a b"))`, and the same text embedded under a query prompt differs from a passage prompt, so picking the wrong role can move a different document into 5th place.

**Agent pitfall:** Renaming or re-wrapping the exact public symbol (class, function, or variable) that the prompt explicitly names, substituting a more idiomatically descriptive name, on the assumption that the checker will call whatever the agent chose to define.
**Trigger:** The task text names an importable entrypoint in fixed wording, for example "in a class called X", "expose `BeginRepl(args, version)`", or a named API variable, and the agent instead ships a differently-named implementation.
**Concrete correction:** Declare the prompt-named symbol verbatim (exact name and signature) as a thin exported shim and delegate the real work to internal helpers; treat any renaming as a broken contract because the graded harness imports and calls the literal name, not a paraphrase.
**Tiny verification:** If a spec requires `class Server`, the harness does `from module import Server`; a class shipped under a name like `Servicer` is untouched by that import and fails even though it runs correctly by hand.

### I.1.3 Artifact seed-058 (best on mteb-retrieve)

**Software Engineering Skill**

- Read the task and relevant source before changing code.
- Preserve existing behavior while implementing the requested public API.
- Make a coherent implementation across affected modules.
- Verify normal behavior and edge cases with focused tests.
- Treat every precise validation, parsing, error, export, and format rule in the prompt as an explicit assertion.
- Run relevant existing tests and address regressions.

**Agent pitfall:** When feeding text to a subword tokenizer or a retrieval encoder from an external AI library, the agent guessed the assembled text shape (joining related fields with a space or nothing) and the special-token flag, and guessed the per-role prompt, instead of reading the documented contract.
**Trigger:** A counting or retrieval result that is sensitive to how the exact text is assembled and to which prompt is applied: combining two text fields with a space versus a newline, a guessed special-token setting, or encoding the query and the documents with the same prompt while task-aware encoders distinguish them, can shift a token count or change which document ranks $N$th.
**Concrete correction:** Check the source/package documentation for the exact text format and prompt roles before encoding. Join independent text fields with the documented separator (BPE tokenizers merge a leading space into the following subword, so a newline avoids cross-field merging) and set the tokenizer's special-token flag deliberately. With a retrieval encoder, encode the query with the query prompt and the documents with the passage prompt (e.g. a `prompt_type`/role argument), never one shared default for both.
**Tiny verification:** For a BPE tokenizer, tokenizing `"a" + " " + "b"` returns a different subword count than `"a" + "\n" + "b"`; embedding a query with the query prompt versus the passage prompt reorders which document has the $N$th-highest cosine similarity.

**Agent pitfall:** The agent implemented the requested behavior behind an idiomatically-named class, function, or registry key instead of the exact public identifier and signature the task states, so the code builds and runs under its own tests but the checker never reaches it.

**Trigger:** A hidden checker imports or instantiates a documented identifier (a service class name, a package function, a module registry key) or calls a documented public entrypoint with a stated signature; any differently-named or re-signatured variant is invisible to it even though it behaves correctly in isolation.
**Concrete correction:** Before writing code, enumerate the exact public identifiers the prompt names—class name, function name, full parameter list and return type, registry key—and name the artifacts precisely; keep existing public entrypoints and signatures unchanged and add new internals around them rather than renaming the reachable surface.
**Tiny verification:** `from module import ExactName` resolves only when the class/function is literally `ExactName`; a `class NearbyName` that subclasses the expected type is not found by an importer that looks up `ExactName`.