Self-Improvement Memory Enables Continual Recursive Self-Improvement

Xin Yu[1,4,*], **Fei Xia**[1], **Jiaqi Pan**[1], **Youbao Tang**[1], **Chengzhi Liu**[1], **Adams Yu**[2], **Sicong Jiang**[2], **Yanfei Chen**[3], **Lingzhou Xue**[4] **and Ting Yu**[1]
[1]Google Cloud AI
[2]Google DeepMind
[3]Google Cloud AI Research
[4]The Pennsylvania State University
{xinyulucas, feixia, jiaqip, tangyoubao, chengzhiliu, adamsyuwei, jiangsicong, yanfeichen, yuti}@google.com, lzxue@psu.edu

arXiv:2611.18856v1 [cs.LG] 13 Nov 2026

Harness optimization is often studied with a fixed policy model and benchmark. When either changes, the optimizer must decide which parts of an evolved harness to reuse or revise, but the harness alone may not explain why earlier edits were made or how they performed. In this work, we study Continual Recursive Self-Improvement (CRSI), a setting for continuing existing harness-optimization methods across compatible benchmarks or policy models. We propose Self-Improvement Memory (SIM) for the harness optimizer, rather than the task-solving policy. SIM groups edits by recurring harness problem and records their rationale, predicted and observed outcomes, and evaluation context. New attempts extend this history as optimization continues. In our experiments across four domains, SIM achieves higher final scores than harness-only and root-cause-only continuation in cross-dataset and pairwise cross-model comparisons. Across ten independent source runs in coding, paired continuations with SIM improve final scores by an average of 2.36 points over harness-only continuation ($p = 0.003$).

Keywords: Recursive self-improvement, agent harnesses, continual learning

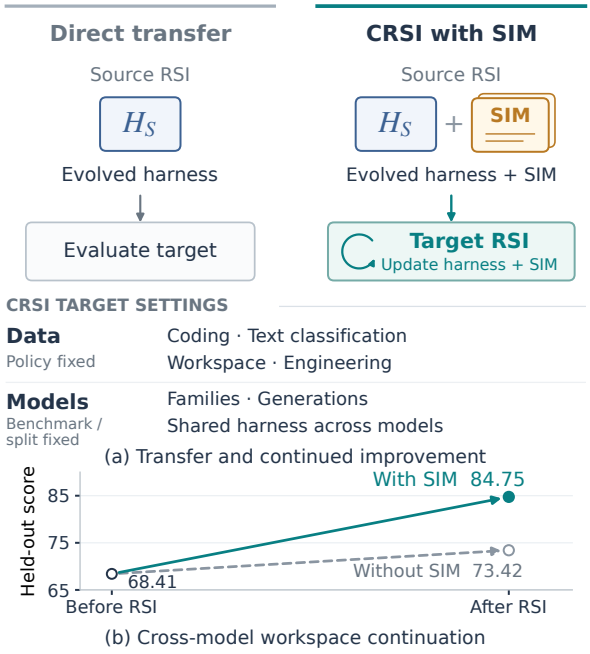

Figure 1 | **Continual improvement beyond direct transfer.** (a) CRSI carries the evolved harness and SIM into target RSI across data or models. SIM guides reuse and revision. (b) Both conditions continue RSI from the same Opus-evolved harness on Gemini's workspace tasks; arrows join measured endpoints with matched rounds and candidates.

1. Introduction

An LLM agent depends on both its policy model and the harness that manages context, tools, and execution. Improving the harness can raise performance without updating model weights, but finding useful changes by hand requires repeated engineering and evaluation (Lin et al., 2026; Xia et al., 2026). Recent work automates this process through recursive self-improvement (RSI), using execution feedback to repeatedly propose, evaluate, and select harness edits (Lee et al., 2026; Lin et al., 2026; Xia et al., 2026). Although evolved harnesses can improve per-

*This work was done while Xin Yu was an intern at Google.

Corresponding author(s): xinyulucas@google.com
© 2026 Google. All rights reserved

formance on other benchmarks or models (Lin et al., 2026; Wu et al., 2026; Xia et al., 2026; Xu et al., 2026b), optimization is usually conducted within a fixed setting. As models are updated and agents encounter new tasks, harness optimization often needs to continue beyond that setting. This raises two fundamental challenges.

Adapting to a new model. Model adaptation arises in two settings: *multi-model serving*, where one shared harness supports different policy models, and *model evolution*, where new versions replace earlier ones within a model family. Both motivate reusing an evolved harness and its improvement experience instead of independently optimizing a harness for every model or release. Yet modifications suited to one policy can interact differently with another, even when the benchmark stays fixed (Yang et al., 2026; Yu et al., 2026d). In our workspace experiment, an Opus-evolved harness starts at 68.41 on Gemini, below the seed harness's 73.89 (Figure 4). The optimizer must determine which inherited mechanisms still help and which need revision.

Adapting to new data. With the policy fixed, a different task distribution can expose failures that source optimization did not address. A completion gate that checks whether commands ran successfully, for example, may need to check whether a feasible solution is also optimal. In our text-classification experiment, the source-evolved FewShotAll harness starts at 72.92 on the target suite, below the seed's 75.72 (Figure 3). A useful source fix may therefore need revision before it helps the target.

In either transition, the evolved harness alone leaves important questions unanswered: why was an edit proposed, what was it expected to fix, and what happened when it was tested? Source optimization produces this evidence, but it need not accompany the harness. Without this history, target optimization may repeat earlier attempts or change inherited components without knowing their purpose. We therefore ask whether the experience gained during source optimization can help the harness adapt to new models or data.

In this work, we study **Continual Recursive Self-Improvement (CRSI)**, in which harness optimization continues after a benchmark or policy-model change while preserving a compatible executable interface. We propose **Self-Improvement Memory (SIM)** to carry improvement experience alongside the evolved harness (Figure 1). SIM groups selected edits by the harness problems they address and gives the proposer and selector access to their rationale, predicted and observed outcomes, and evaluation context. It creates entries for new problems and appends further attempts and outcomes to existing ones, helping the optimizer decide which modifications to retain, revise, replace, or remove.

Our contributions are as follows:

- **Continual harness optimization.** We define cross-dataset continuation with the policy fixed and cross-model continuation with the benchmark and split fixed. Existing harness-optimization methods can use this setting while retaining their proposal, evaluation, and selection procedures.
- **Self-Improvement Memory.** We introduce a structured record of harness problems, modifications, and outcomes that conditions both proposal and selection, preserving source experience while incorporating new attempts during target optimization.
- **Controlled evidence across four domains.** With matched starting harnesses, round counts, and candidate budgets, Full SIM outperforms harness-only and root-cause-only continuation in six cross-dataset and four Opus-to-Gemini settings. Across ten independent source runs in coding, paired continuations with SIM improve 9 of 10 runs, with a mean advantage of 2.36 points over harness-only continuation ($p = 0.003$).

2. Related Work

Recursive self-improvement of agent harnesses. Yu et al. (2026b) survey harness optimization workflows, recursive reuse, and the conditions under which evolved modifications transfer. Meta-Harness and AHE search

over harness code using execution traces and scores (Lee et al., 2026; Lin et al., 2026). Other approaches let the agent modify its own system (Zhang et al., 2026a,b,d) or learn a policy for editing harnesses (Shao et al., 2026). SIA instead combines harness edits with updates to the task model's weights (Hebbar et al., 2026).

Transfer of evolved harnesses. Evolved harnesses can retain their gains on held-out tasks, other models, or downstream benchmarks (Lee et al., 2026; Lin et al., 2026; Shi et al., 2026). Recent methods encourage transfer through benchmark-disjoint evolution and modular edits (Wu et al., 2026), regularized search (Xia et al., 2026), or separation of general and benchmark-specific mechanisms (Yu et al., 2026a). Their effectiveness still depends on how the model and harness interact (Yang et al., 2026; Yu et al., 2026d). PluginRSI also reuses its plugin library in subsequent optimization.

Memory of the improvement process. Prior work retains candidate artifacts (Lee et al., 2026), edit histories (Xia et al., 2026), improvement graphs (Zhang et al., 2026c), skill-revision records (Li and Hu, 2026), and playbooks (Zhang et al., 2026f). HarnessBank groups candidate harnesses by the problems they address (Luo et al., 2026). Xu (2026) continues multi-task evolution with free-text notes and source replay; Harness^2 reuses an improvement library across benchmarks (Xu et al., 2026a). Meta-Skill and Meta^2 transfer learned guidance for constructing or editing task-specific harnesses (Qian et al., 2026; Yuan et al., 2026). Continual adaptation also raises retention and interference concerns (Chen et al., 2026; Zhao et al., 2026). SIM links harness problems to modifications, predicted and observed outcomes, and provenance. We test SIM by continuing optimization from the same evolved harness with no SIM, root causes only, or the full record after a benchmark or model change. Appendix A expands these comparisons, including continued optimization and memory for task-solving agents.

3. Method

3.1. Problem Formulation

Harness RSI. Following Meta-Harness (Lee et al., 2026), we optimize an executable harness H around a fixed policy model M. The harness determines the model's context, tools, and execution flow. We write a setting as $\sigma = (M, \mathcal{X})$, with disjoint evolve and held-out evaluation sets $\mathcal{X}^{\text{evo}}$ and $\mathcal{X}^{\text{eval}}$. The score $J(H; M, \mathcal{X})$ follows the benchmark's aggregation rule (Appendix B). Only evolve-set feedback is available for proposing edits, selecting candidates, updating SIM, or deciding when to stop.

A run consists of N rounds. At round $t = 0, \ldots, N-1$, the evaluator executes H_t on $\mathcal{X}^{\text{evo}}$ and returns scores, traces, and errors as F_t. The proposer P uses F_t and the run's history $\mathcal{D}_t$ (earlier feedback and code diffs) to generate k candidate harnesses $\mathcal{C}_t$. The selector $\mathcal{S}$ screens proposals before execution and chooses among the evaluated candidates, weighing performance, complexity, domain specificity, and regression risk:

$$\begin{aligned} \mathcal{C}_t &= P(H_t, F_t, \mathcal{D}_t), \\ H_{t+1} &= \mathcal{S}(H_t, F_t, \mathcal{C}_t, F_t^{\mathcal{C}}), \end{aligned} \tag{1}$$

Here $F_t^{\mathcal{C}}$ contains feedback from the candidates that passed screening and were executed. If no candidate is accepted, $H_{t+1} = H_t$. The proposer and selector are LLM-based harness engineers; M is the policy that executes the tasks.

Continual setting. After source RSI in $\sigma_S = (M_S, \mathcal{X}_S)$, optimization resumes from the final retained harness H_S in $\sigma_T = (M_T, \mathcal{X}_T)$. In *cross-dataset* CRSI, the benchmark changes while the policy stays fixed ($\mathcal{X}_S \neq \mathcal{X}_T$, $M_S = M_T$). In *cross-model* CRSI, the policy changes on the same benchmark and split ($M_S \neq M_T$, $\mathcal{X}_S = \mathcal{X}_T$). Both require a shared executable harness interface. We measure target improvement as the final minus starting score under $J(\cdot\,; M_T, \mathcal{X}_T^{\text{eval}})$.

3.2. Self-Improvement Memory

SIM groups edits by the recurring harness problems they address. For a root-cause hypothesis

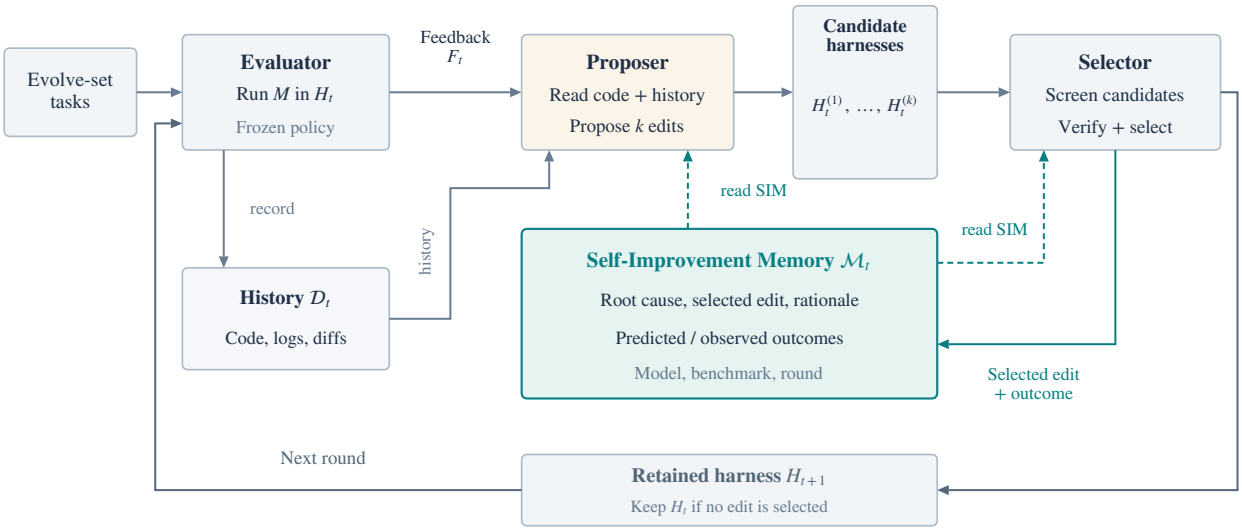

Figure 2 | **SIM in the agentic outer loop.** Starting from $(H_S, \mathcal{M}_S)$, the proposer and selector consult SIM (dashed arrows) to improve the harness on evolve-set tasks. After selection, SIM records the chosen edit and its measured outcome (teal arrow). Held-out evaluation stays outside the loop.

c_i, such as premature completion, an entry is

$$m_i = \big(c_i, [(a_{ij}, o_{ij}, p_{ij})]_j\big), \tag{2}$$

where j indexes attempts in order and $\mathcal{M}_t = \{m_i\}$ is the current memory. Each attempt stores the modification and its rationale a_{ij}, predicted and observed task outcomes o_{ij}, and provenance p_{ij} (domain, benchmark, policy, and round). Predictions of which tasks an edit will fix are recorded before evaluation and checked afterward.

At each round, the proposer receives the complete SIM in its system prompt, and the selector reads the same record when assessing candidates. Both read all entries without retrieval or truncation. The loop in Equation 1 becomes

$$\begin{aligned} \mathcal{C}_t &= P(H_t, F_t, \mathcal{D}_t; \mathcal{M}_t), \\ H_{t+1} &= \mathcal{S}(H_t, F_t, \mathcal{C}_t, F_t^{\mathcal{C}}; \mathcal{M}_t). \end{aligned} \tag{3}$$

Each proposal cites an existing root cause or names a new one. After selection, *only the chosen candidate* is recorded in e_t, including its measured outcomes even when the edit regresses or fails to fix a predicted task. The update $\mathcal{M}_{t+1} = U(\mathcal{M}_t, e_t)$ appends this record to the cited problem or creates a new entry, preserving earlier attempts. If no edit is accepted, $e_t = \emptyset$ and $U(\mathcal{M}_t, \emptyset) = \mathcal{M}_t$. Entries with identical root-cause text are merged without dropping their histories. Appendix D.1 gives the entry-matching details.

3.3. Continual RSI via SIM

Source RSI starts with an empty SIM and produces $(H_S, \mathcal{M}_S)$. Full SIM carries both into target optimization (Figure 2). The proposer and selector use each edit's rationale, outcome, and evaluation context to judge whether it still applies to the current failures. They can retain, revise, replace, or remove harness modifications while keeping the earlier memory records. Appendix D.2 gives the prompt instructions for using this history.

Transfer conditions. We compare four conditions: *Cold Start* $(H_0, \emptyset)$ from the seed harness; *Harness Transfer* $(H_S, \emptyset)$; *Problem-Only Memory* $(H_S, \mathcal{M}_S^c)$ with root causes only; and *Full SIM* $(H_S, \mathcal{M}_S)$. The three transfer conditions start from the same H_S, and all four use the same target policy, tasks, and round and candidate budgets. Each keeps its assigned memory condi-

tion throughout target RSI, while retaining the ordinary target-run history $\mathcal{D}_t$ (Appendix C.3). Source history enters through SIM; new target records follow the same memory condition.

4. Experiments

Our experiments address four questions:

- Does SIM improve continued RSI after a benchmark or policy-model change? (§4.2–4.3)
- Does the full SIM record help beyond retaining only root causes? (§4.2 and §4.3.1)
- How does SIM guide the reuse and revision of inherited harness edits? (§4.4; Appendix F.5)
- Is SIM's benefit consistent across independently evolved source harnesses? (§4.4)

4.1. Experimental Setup

Domains. We study four domains in which a source-evolved harness remains executable on the target (Appendix B, Table 4). *Coding* moves from Terminal-Bench 2.1 (TB2.1) to Terminal-Bench 4 (TB4) (Merrill et al., 2026; Terminal-Bench Team, 2026a,b), containerized shell tasks scored by their own tests. *Text classification* uses the online classification environment of Meta-Harness (Lee et al., 2026), from a three-dataset source suite to an eight-dataset target suite absent from source RSI. *Agentic workspace* moves from Harvey LAB (Harvey AI, 2026) to the equal-weight mean of JobBench (Li et al., 2026), GDPval (Patwardhan et al., 2026), and APEX-Agents (Vidgen et al., 2026), using scoring interfaces adapted from RRSI (Xia et al., 2026), with Gemini 3.8 Flash as the judge. *Engineering design* moves from EngDesign (Guo et al., 2025) to Frontier-Eng (Chi et al., 2026), scored by frozen verifiers. Tasks are split uniformly at random, without stratification, into evolve and held-out sets of approximately equal size; exact counts are in Appendix B. Only evolve-set scores and trajectories reach the proposer, selector, and SIM. The main comparisons use five rollouts per task, with every rollout included in the score.

Transitions, models, and budgets. Cross-dataset CRSI uses Gemini 3.8 Flash (Google DeepMind, 2026b) as the policy in both stages. For the cross-model transfer comparison, source RSI uses Claude Opus 5.5 (Anthropic, 2026), then Gemini 3.8 Flash continues on the same benchmark and split. Both stages use the target benchmarks listed above: TB4, the eight-dataset classification suite, the merged workspace benchmarks, and Frontier-Eng. Policy weights remain frozen. Gemini 3.8 Flash also serves as proposer and selector in the main comparisons. A source-RSI ablation replaces it with Claude Opus 5.5 in these roles, yielding larger gains from both coding seeds (Appendix E.4.2, Table 7). The seeds are the minimal AHE seed (Lin et al., 2026) and Terminus-2 (Merrill et al., 2026) for coding, NoMemory and Few-ShotAll (Lee et al., 2026) for text classification, and a ReAct agent (Yao et al., 2023) over an MCP tool gateway (Vidgen et al., 2026) for workspace and engineering. For the latter two domains, we adopt RRSI's environments, seed harnesses, and evaluation interfaces (Xia et al., 2026); Appendix C details our experimental setup. Every source and target run uses $N = 10$ rounds with $k = 5$ candidate proposals per round in the proposer–selector loop described in §3.1. Proposals rejected before execution receive no evaluation score; the five rollouts per task specify the separate evaluation allocation.

Comparison conditions. Cold Start begins from the seed harness without SIM. Harness Transfer, Problem-Only Memory, and Full SIM share the same source-evolved harness, with no SIM, root causes only, and the complete record, respectively. These memory conditions persist throughout target optimization, with matched target tasks, policy, round counts, and candidate budgets (§3.3).

Metrics and reporting. Scores are on a 0–100 scale to two decimals: pass@1 for coding, accuracy pooled over examples for text classification, benchmark-specific rubric, win-rate, and pass@1 scores for workspace, and pass rate or medal score for engineering (Appendix B). Before and

after scores evaluate the starting harness and the version retained after the final target round on the held-out set. Except for the dedicated paired study, each condition is one complete RSI run, and its five rollouts per task estimate execution performance rather than optimization variance. We report individual settings descriptively and use the ten-pair coding study to assess variation across independent source runs (§4.4).

4.2. Cross-Dataset CRSI

SIM improves final scores in all six dataset transitions. The memory ablation in Figure 3 compares continued optimization from the same evolved harness with no SIM, root causes only, or the full record. Full SIM finishes highest in all six settings, exceeding Harness Transfer by 2.43–4.35 points and Problem-Only Memory by 0.20–6.06 points. Problem-Only Memory has mixed effects: it ranges from 1.82 points below Harness Transfer to 2.40 above it. Keeping the edit history and evaluation evidence thus helps beyond retaining the diagnosis alone (Appendix E.4.1).

The margin over the strongest alternative ranges from 0.20 points on NoMemory to 4.35 on engineering. Section 4.4 tests whether the benefit recurs across independently evolved source harnesses. Appendix E.2 lists the exact starting and final scores for both transfer settings.

Source improvement does not ensure a better target start. Source RSI raises held-out source scores by 2.75–6.41 points in all six settings (Appendix E.1), yet the FewShotAll harness starts target RSI below the seed: 72.92 versus 75.72. Harness-only continuation reaches 74.12 and remains below Cold Start at 75.92, while Full SIM reaches 77.92.

On NoMemory, Cold Start gains more than Full SIM (3.40 versus 2.20 points) but finishes lower (72.00 versus 73.60). We therefore report both endpoints and use Harness Transfer, which shares Full SIM's initialization, as the direct control for the memory procedure.

Gemini model	Seed	Fresh RSI	CRSI
2.5 Flash	17.8	21.4	**22.4**
3.5 Flash	70.2	73.3	**76.8**
3.8 Flash	87.2	88.9	**90.5**
4 Argon	89.7	**95.6**	94.8

Table 1 | Success rates on held-out TB2.1 across Gemini generations. Seed uses the unevolved harness; Fresh RSI optimizes it independently. CRSI adds SIM in the first generation and inherits the previous generation's harness and SIM thereafter. Bold compares Fresh RSI and CRSI.

4.3. Cross-Model CRSI

We examine transfer between model families, continuation across model generations, and a harness shared by multiple models.

4.3.1. Harness Transfer across Model Families

When the Opus-evolved harness is first run with Gemini, it scores below Gemini's seed harness in every domain (Figure 4). The deficits range from 0.60 points in text classification to 5.48 in workspace. Target RSI does not always overcome this disadvantage: Harness Transfer still finishes below Cold Start in three domains. Full SIM finishes highest in all four and exceeds harness-only continuation by 3.60–11.33 points. It also finishes above Problem-Only Memory in all four domains (Appendix E.4.1). Workspace shows the largest difference: both begin at 68.41, but Full SIM reaches 84.75 while Harness Transfer reaches 73.42.

4.3.2. Continual Improvement across Model Generations

The generation chain carries the evolved harness and SIM through successive Gemini models on TB2.1 (Table 1; setup in Appendix C.5). Fresh RSI starts from the seed at each generation. In the first generation, both procedures use the same seed, so that comparison tests SIM during the initial run. *Continuation improves over fresh RSI in two of the three subsequent generations*: by 3.5 points on Gemini 3.5 Flash and 1.6 on Gemini 3.8 Flash, while finishing 0.8 points lower on

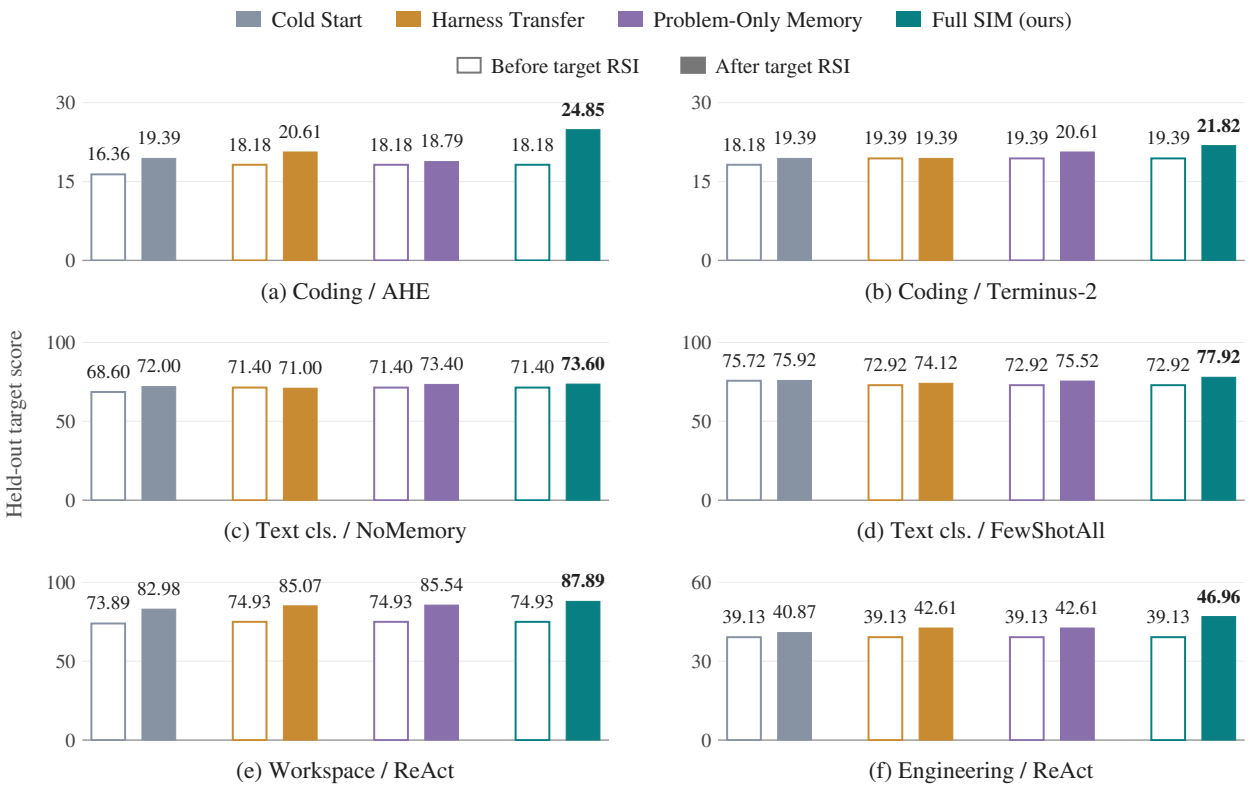

Figure 3 | Cross-dataset CRSI: held-out target scores before (open bars) and after (filled bars) target RSI. The three transfer conditions share the source-evolved harness; Cold Start uses the seed. Axis limits vary by domain. Table 5 lists exact scores.

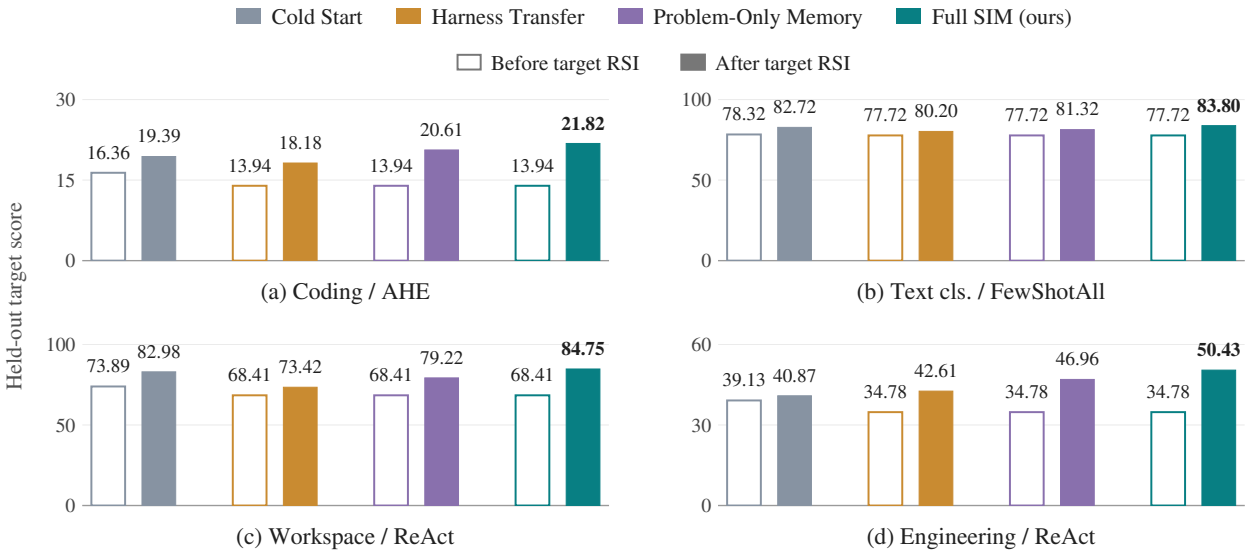

Figure 4 | Cross-model CRSI from Opus 5.5 to Gemini 3.8 Flash, with benchmark and split fixed. Open and filled bars show held-out scores before and after Gemini's target RSI. Transfer conditions share the Opus-evolved harness; Cold Start uses the seed (Table 6).

Harness construction	Evolve	Held-out	All
B1: Per-model harnesses	77.40	75.20	76.30
B2: Shared seed	72.36	73.28	72.82
S1: Sequential CRSI	76.72	74.20	75.46
S2: Unified synthesis	**78.92**	**75.24**	**77.08**
S3: Source policy used for RSI			
Gemini 3.8 Flash	74.28	71.24	72.76
Grok 4.20 Reasoning	68.24	66.84	67.54
Claude Opus 5.5	75.20	72.24	73.72

Table 2 | Text-classification scores averaged over Grok 4.20 Reasoning, Claude Opus 5.5, and Gemini 3.8 Flash. S3 rows name the policy used to evolve each harness; each harness is evaluated with all three models. All averages the equally sized evolve and held-out splits. Bold marks the highest mean in each column.

Run	Start	No SIM	SIM	Δ
1	18.18	20.61	22.42	+1.81
2	16.97	17.58	19.39	+1.81
3	17.58	20.00	20.61	+0.61
4	18.18	18.18	23.64	+5.46
5	17.58	18.79	20.61	+1.82
6	19.39	20.00	23.64	+3.64
7	18.79	20.00	21.82	+1.82
8	18.18	20.00	19.39	−0.61
9	19.39	18.18	23.03	+4.85
10	19.39	20.00	22.42	+2.42
Mean	18.36	19.33	**21.70**	**+2.36**

Table 3 | Paired coding continuations from ten independent Terminus-2 source runs. Start is the evolved harness's held-out TB4 score; No SIM and SIM are scores after ten target rounds. Δ is SIM minus No SIM (mean +2.36; paired $p = 0.003$).

Gemini 4 Argon.

4.3.3. *One Harness for Multiple Models*

We also ask whether one harness can serve Grok, Opus, and Gemini on the text-classification target suite (Table 2; setup in Appendix C.6). We compare sequential CRSI, unified synthesis using feedback from all three models, and transfer of harnesses evolved separately with each policy. *Unified synthesis reaches a similar mean score to separately evolved harnesses*: 75.24 versus 75.20 on held-out examples, above the shared seed's 73.28. Sequential CRSI reaches 74.20; the three S3 harnesses average 66.84–72.24 across the evaluation models. This descriptive comparison assesses a common harness; the matched starting-harness comparisons in §4.3.1 test the full SIM procedure.

4.4. Analysis: The Role of SIM in CRSI

The paired comparisons test whether SIM improves continuation from a common starting harness. We then follow individual memory entries to see how earlier diagnoses inform later edits.

SIM improves continuation from the same evolved harness. We repeat source RSI ten times from Terminus-2 on TB2.1 and initialize two ten-round TB4 runs from each resulting harness. Within each pair, the starting harness, policy, benchmark, and round and candidate budgets match. One run inherits that source run's SIM and appends target attempts; the other continues without SIM. The independent source run is the unit of analysis.

SIM finishes higher in 9 of 10 pairs (Table 3), with one pair favoring continuation without SIM by 0.61 points. From a shared mean starting score of 18.36, the mean final score is 21.70 with SIM and 19.33 without it. The mean paired advantage is 2.36 points (95% CI $[1.05, 3.68]$; two-sided paired $t(9) = 4.06$, $p = 0.003$). A two-sided signed-rank check gives $p = 0.006$; the paired standardized effect is $d_z = 1.28$ (statistical details in Appendix C.4).

SIM guides reuse and revision of inherited edits. Full SIM and Harness Transfer inherit the same completion gate from Terminus-2 source RSI (Figure 11b, Appendix F.3). With SIM, the proposer uses the gate's rationale and source outcomes to retain completion verification while adding target checks, including feasibility versus optimality. Without SIM, the run tightens the gate, twice repairs completion deadlocks,

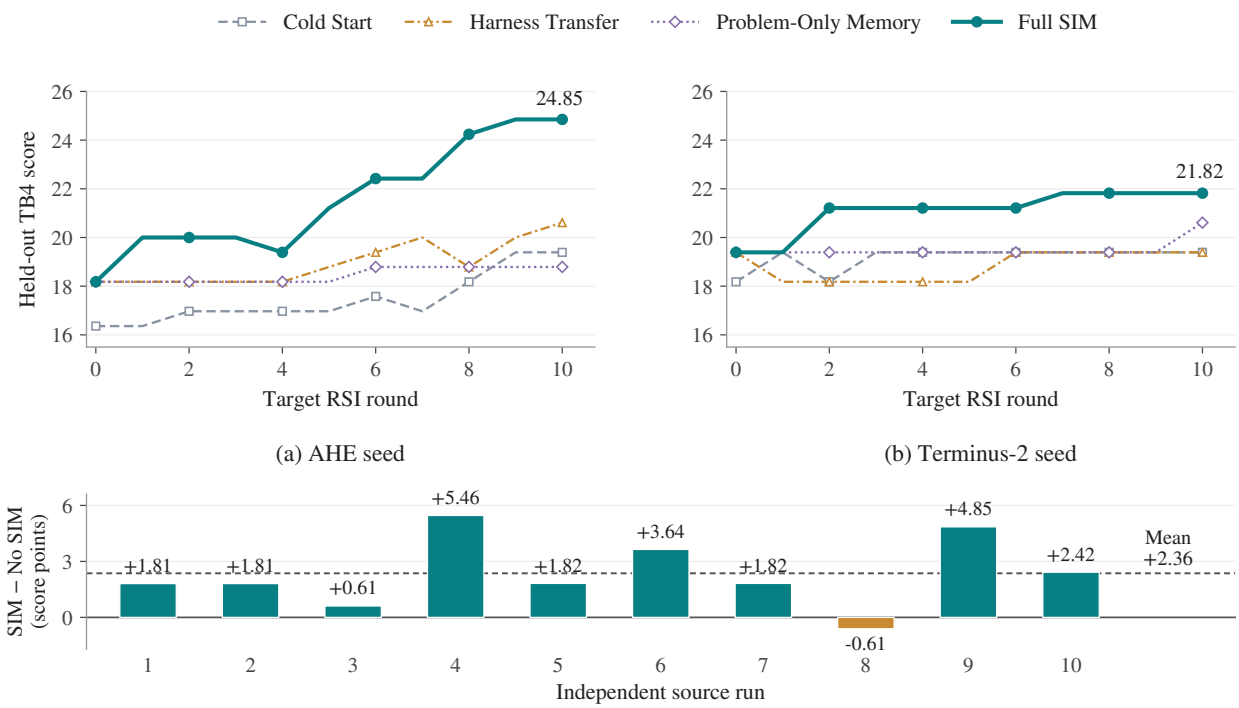

Figure 5 | Coding optimization (TB2.1→TB4). (a,b) Held-out TB4 scores of retained harnesses after each target round; round 0 is the starting harness. (c) Final-score differences (SIM minus No SIM) across ten paired continuations from independently evolved Terminus-2 harnesses. The dashed line marks the mean; Table 3 gives individual scores.

and misattributes the inherited gate to its own first round. The requirement carries over; its implementation needs revision. Appendices F.5 and F.6 summarize transfer patterns across domains and models and the lessons for SIM.

SIM preserves diagnoses while adding target evidence. P3 records a mismatch between deliverables and the verifier's interface contract (Figure 12; Appendix F.4). Target rounds 5 and 6 append revised checks and measured outcomes under this diagnosis, keeping source attempts and unresolved predictions available to later rounds.

SIM supports earlier target progress and source retention in coding. Full SIM is best or tied throughout rounds 1–10 for both coding seeds (Figure 5a,b). It exceeds the strongest control's final score (20.61) by round 5 from AHE and round 2 from Terminus-2. In the AHE source-retention trace, all transfer conditions start at 90.22 on held-out TB2.1 (Figure 8, Appendix E.3). Full SIM loses 3 successful rollouts out of 225, versus 4 for Problem-Only Memory and 6 for Harness Transfer, while gaining more on TB4 than Problem-Only Memory (6.67 versus 0.61 points).

5. Conclusion

CRSI continues harness optimization when the benchmark or policy model changes. SIM carries earlier diagnoses, edits, and outcomes into this process and records new attempts. From matched starting harnesses, Full SIM outperforms harness-only continuation in all six cross-dataset settings and four Opus-to-Gemini domains. The ten-pair coding study finds a mean advantage of 2.36 points across independent

source runs. The cases show how this history helps the proposer and selector reuse earlier reasoning while revising fixes for new requirements.

Limitations

Within our compute budget, independent RSI repetitions focus on coding; the other domains are evaluated with one run per condition. We study settings with compatible harness interfaces and keep the optimizer's planning and search procedures fixed. Our comparisons evaluate SIM as a complete procedure, combining inherited memory with updates during target optimization.

Future work can extend the evaluation to more repeated runs and longer transition sequences, and explore adaptation of planning, token allocation, and tool interfaces.

Ethics Statement

LLM-based optimization agents rewrite harness code in our experiments. The resulting harnesses run inside the benchmarks' execution environments, with policy weights fixed and scores computed by frozen verifiers or a fixed judging procedure. A self-modifying harness could satisfy an evaluator without solving the intended task. We keep held-out feedback away from the optimization agents, identify each benchmark's verifier or judge, and include an accepted edit that failed to repair its target tasks (Appendix F.2). Workspace scores may also reflect biases of the LLM judge. The experiments use existing publicly available benchmarks and datasets.

References

L. A. Agrawal, S. Tan, D. Soylu, N. Ziems, R. Khare, K. Opsahl-Ong, A. Singhvi, H. Shandilya, M. J. Ryan, M. Jiang, C. Potts, K. Sen, A. Dimakis, I. Stoica, D. Klein, M. Zaharia, and O. Khattab. GEPA: Reflective Prompt Evolution Can Outperform Reinforcement Learning. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/0e9e708b6f48e14fd0ac29e167413f76-Abstract-Conference.html.

Anthropic. Introducing Claude Opus 5.5, September 2026. URL https://www.anthropic.com/claude-opus-5-5.

J. Bai, L. Zhang, X. Yu, Y. Wu, Z. Wang, S. Li, W. Li, Z. Zhao, L. Xue, K. Maeng, X. Fan, and B. Peng. Verify Less, Evolve More: Training Idea-Level Critics for Verification-Efficient ML Evolving Agents. *arXiv preprint arXiv:2610.08993*, 2026. URL https://arxiv.org/abs/2610.08993.

F. Barbieri, J. Camacho-Collados, L. E. Anke, and L. Neves. TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, 2020. doi: 10.18653/v1/2020.findings-emnlp.148. URL https://aclanthology.org/2020.findings-emnlp.148/.

I. Casanueva, T. Temčinas, D. Gerz, M. Henderson, and I. Vulić. Efficient Intent Detection with Dual Sentence Encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45, 2020. doi: 10.18653/v1/2020.nlp4convai-1.5. URL https://aclanthology.org/2020.nlp4convai-1.5/.

X. Chen, Y. Huang, Y. Xu, Z. Zhao, E. McQuinn, J. Liu, and H. Yao. Are Self-Evolving Agents Continual Learners? *arXiv preprint arXiv:2610.25858*, 2026. URL https://arxiv.org/abs/2610.25858.

Y. Chi, D. Hong, D. Jiang, T. Luo, et al. Frontier-Eng: Benchmarking Self-Evolving Agents on Real-World Engineering Tasks with Generative Optimization. *arXiv preprint arXiv:2604.12290*, 2026. URL https://arxiv.org/abs/2604.12290.

A. Cohan, W. Ammar, M. van Zuylen, and F. Cady. Structural Scaffolds for Citation Intent Classification in Scientific Publications. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies,*

Volume 1 (Long and Short Papers), pages 3586–3596, 2019. doi: 10.18653/v1/N19-1361. URL https://aclanthology.org/N19-1361/.

D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. Ravi. GoEmotions: A Dataset of Fine-Grained Emotions. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054, 2020. doi: 10.18653/v1/2020.acl-main.372. URL https://aclanthology.org/2020.acl-main.372/.

Z. Fei, X. Shen, D. Zhu, F. Zhou, et al. LawBench: Benchmarking Legal Knowledge of Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7933–7962, 2024. doi: 10.18653/v1/2024.emnlp-main.452. URL https://aclanthology.org/2024.emnlp-main.452/.

Gemini Team. Gemini 2.5: Pushing the Frontier with Advanced Reasoning, Multimodality, Long Context, and Next Generation Agentic Capabilities. *arXiv preprint arXiv:2507.06261*, July 2025. URL https://arxiv.org/abs/2507.06261.

Google DeepMind. Gemini 3.5 Flash Model Card, May 2026a. URL https://deepmind.google/models/model-cards/gemini-3-5-flash/.

Google DeepMind. Gemini 3.8 Flash Model Card, September 2026b. URL https://deepmind.google/models/model-cards/gemini-3-8-flash/.

Google DeepMind. Gemini 4 Argon Model Card, October 2026c. URL https://deepmind.google/models/model-cards/gemini-4-argon/.

Gretel AI. Symptom to Diagnosis Dataset, 2023. URL https://huggingface.co/datasets/gretelai/symptom_to_diagnosis.

X. Guo, Y. Li, X. Kong, Y. Jiang, et al. Toward Engineering AGI: Benchmarking the Engineering Design Capabilities of LLMs. In *Advances in Neural Information Processing Systems*, 2025. URL https://arxiv.org/abs/2509.16204.

Harvey AI. Harvey LAB: The Legal Agent Benchmark, 2026. URL https://github.com/harveyai/harvey-labs/tree/v1.0.

P. Hebbar, Y. Manawat, S. Verboomen, A. Ivanova, et al. SIA: Self Improving AI with Harness & Weight Updates. *arXiv preprint arXiv:2605.27276*, 2026. URL https://arxiv.org/abs/2605.27276.

S. Karten, J. Zhang, T. Upaa, Jr, R. Feng, et al. Continual Harness: Online Adaptation for Self-Improving Foundation Agents. *arXiv preprint arXiv:2605.09998*, 2026. URL https://arxiv.org/abs/2605.09998.

Z. Ke, V. Patil, H. Shi, Y. Li, et al. Evo-HarnessBench: Can Your Agents Keep Pace with an Evolving Harness? *arXiv preprint arXiv:2609.04280*, 2026. URL https://arxiv.org/abs/2609.04280.

P. Keung, Y. Lu, G. Szarvas, and N. A. Smith. The Multilingual Amazon Reviews Corpus. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4563–4568, 2020. doi: 10.18653/v1/2020.emnlp-main.369. URL https://aclanthology.org/2020.emnlp-main.369/.

T. Khot, A. Sabharwal, and P. Clark. SciTail: A Textual Entailment Dataset from Science Question Answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. doi: 10.1609/aaai.v32i1.12022.

K. Kim, M. Kang, T. Kim, Y. Yang, et al. Memory Transfer Learning: How Memories are Transferred Across Domains in Coding Agents. *arXiv preprint arXiv:2604.14004*, 2026. URL https://arxiv.org/abs/2604.14004.

Y. Lee, R. Nair, Q. Zhang, K. Lee, et al. Meta-Harness: End-to-End Optimization of Model Harnesses. In *Conference on Language Modeling*, 2026. URL https://openreview.net/forum?id=tmbOUyFx3R.

Y. Li, Y. Feng, Z. Xu, Z. Ma, et al. JobBench: Aligning Agent Work With Human Will. *arXiv preprint arXiv:2605.26329*, 2026. URL https://arxiv.org/abs/2605.26329.

Z. Li and Y. Hu. SkillHone: A Harness for Continual Agent Skill Evolution Through Persistent Decision History. In *Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages

1819–1829. Association for Computational Linguistics, 2026. doi: 10.18653/2026.emnlp-industry.105. URL https://aclanthology.org/2026.emnlp-industry.105/.

J. Lin, S. Liu, C. Pan, L. Lin, et al. Agentic Harness Engineering: Observability-Driven Automatic Evolution of Coding-Agent Harnesses. *arXiv preprint arXiv:2604.25850*, 2026. URL https://arxiv.org/abs/2604.25850.

Z. Liu, Z. Shi, Y. Sang, B. He, et al. Adaptive Auto-Harness: Sustained Self-Improvement for Agentic System Deployment on Open-Ended Task Streams. *arXiv preprint arXiv:2606.01770*, 2026. URL https://arxiv.org/abs/2606.01770.

X. Luo, D. Xue, F. Wang, C. Hu, and Y. Deng. HarnessBank: Semantic GeneBank Search with Gated Verification for Agent-Harness Self-Evolution. *arXiv preprint arXiv:2607.13683*, 2026. URL https://arxiv.org/abs/2607.13683.

P. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala. Good Debt or Bad Debt: Detecting Semantic Orientations in Economic Texts. *Journal of the Association for Information Science and Technology*, 65(4):782–796, 2014. doi: 10.1002/asi.23062.

M. A. Merrill, A. G. Shaw, N. Carlini, B. Li, et al. Terminal-Bench: Benchmarking Agents on Hard, Realistic Tasks in Command Line Interfaces. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=a7Qa4CcHak.

S. Ouyang, J. Yan, I.-H. Hsu, Y. Chen, et al. ReasoningBank: Scaling Agent Self-Evolving with Reasoning Memory. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=jL7fwchScm.

T. Patwardhan, R. Dias, E. Proehl, G. Kim, et al. GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=hcuEdq6eKD.

C. Qian, K. Zhu, B. Li, Z. Wang, et al. Learning Meta-Skills for Agent Harness Design in Test-Time AI4AI. *arXiv preprint arXiv:2609.38143*, 2026. URL https://arxiv.org/abs/2609.38143.

N. Schneider, N. Stiefl, and G. A. Landrum. What's What: The (Nearly) Definitive Guide to Reaction Role Assignment. *Journal of Chemical Information and Modeling*, 56(12):2336–2346, 2016. doi: 10.1021/acs.jcim.6b00564. URL https://doi.org/10.1021/acs.jcim.6b00564.

S. Shao, K. Zhang, Q. Li, S. Wang, et al. Harness-R1: Learning to Edit Executable Runtime Harnesses from Agent Failure Trajectories. *arXiv preprint arXiv:2608.02276*, 2026. URL https://arxiv.org/abs/2608.02276.

Y. Shi, Y. Miao, Y. Chen, J. Zhang, et al. PluginRSI: Recursive Improvement of Agent Harnesses with Reusable Plugins. *arXiv preprint arXiv:2609.32423*, 2026. URL https://arxiv.org/abs/2609.32423.

N. Shinn, F. Cassano, A. Gopinath, K. Narasimhan, and S. Yao. Reflexion: Language Agents with Verbal Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 36, pages 8634–8652, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html.

Terminal-Bench Team. Terminal-Bench 2.1, May 2026a. URL https://www.tbench.ai/news/terminal-bench-2-1.

Terminal-Bench Team. Terminal-Bench 4.0, August 2026b. URL https://www.tbench.ai/news/terminal-bench-4-0.

B. Vidgen, A. Mann, A. Fennelly, J. W. Stanly, et al. APEX-Agents. *arXiv preprint arXiv:2601.14242*, 2026. URL https://arxiv.org/abs/2601.14242.

G. Wang, Y. Xie, Y. Jiang, A. Mandlekar, et al. Voyager: An Open-Ended Embodied Agent with Large Language Models. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=ehfRiF0R3a.

Y. Wang, H. Zhu, Z. Hu, Y. Yuan, et al. Rethinking the Evaluation of Harness Evolution for Agents. *arXiv preprint arXiv:2607.12227*, 2026. URL https://arxiv.org/abs/2607.12227.

T. Wei, Z. Shi, M. Lin, B. He, et al. Evo-Harness: Context-to-Harness Skill Compilation for Self-Evolving Agents. In *Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing*, pages 22343–22358, 2026. URL https://aclanthology.org/2026.emnlp-main.969/.

S. Wu, J. Ren, Y. Li, H.-S. Li, et al. ModularRSI: Modular and Generalizable Recursive Harness Self-Improvement. *arXiv preprint arXiv:2609.14857*, 2026. URL https://arxiv.org/abs/2609.14857.

P. Xia, R. Han, Z. Wang, Y. Chen, et al. RRSI: Regularized Recursive Self-Improvement of Agent Harnesses. *arXiv preprint arXiv:2609.24972*, 2026. URL https://arxiv.org/abs/2609.24972.

Y. Xiao, Z. Lu, J. Wang, Y. Zhang, et al. Recurrent Harness for Agentic Continual Learning. *arXiv preprint arXiv:2610.12668*, 2026. URL https://arxiv.org/abs/2610.12668.

Q. Xu. Self-Evolving Harness on Multiple Tasks with the Agent as Its Own Optimizer. *arXiv preprint arXiv:2609.38372*, 2026. URL https://arxiv.org/abs/2609.38372.

R. Xu, Y. Chen, Z. CuiZhu, B. D. Mishra, et al. Harness^2: Recursive Agent Harnessing for an Open World. *arXiv preprint arXiv:2610.14224*, 2026a. URL https://arxiv.org/abs/2610.14224.

T. Xu, H. Wen, and M. Li. Adapting the Interface, Not the Model: Runtime Harness Adaptation for Deterministic LLM Agents. *arXiv preprint arXiv:2605.22166*, 2026b. URL https://arxiv.org/abs/2605.22166.

C. Yang, X. Zhao, T. Wu, and C. Kästner. Better Harnesses, Smaller Models: Building 90% Cheaper Agents via Automated Harness Adaptation. *arXiv preprint arXiv:2607.08938*, 2026. URL https://arxiv.org/abs/2607.08938.

S. Yao, J. Zhao, D. Yu, N. Du, et al. ReAct: Synergizing Reasoning and Acting in Language Models. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2210.03629.

W. Yu, L. Miculicich, B. D. Mishra, J. Chen, et al. Harmony Harness: Multi-Objective Self-Evolution of Agent Harnesses. *arXiv preprint arXiv:2610.15150*, 2026a. URL https://arxiv.org/abs/2610.15150.

X. Yu, C. Liu, W. Long, H. Zhang, and L. Xue. From Harness Optimization to Recursive Self-Improvement: A Survey of Language Model Agents. *Preprints.org*, 2026b. doi: 10.20944/preprints202610.2293.v2. URL https://www.preprints.org/manuscript/202610.2293. Version 2, preprint.

X. Yu, L. Zhang, J. Bai, Y. Wu, Z. Wang, S. Li, W. Li, L. Xue, X. Fan, and B. Peng. ML-ToolBench: Learning Tool-Augmented Agents for Machine Learning Development. *arXiv preprint arXiv:2609.36679*, 2026c. URL https://arxiv.org/abs/2609.36679.

Z. Yu, B. Bi, S. K. Pentyala, S. Mehrotra, et al. Co-Evolving Harnesses and Models: On-Policy Correction Helps Weaker Models Catch Up Where Imitation Fails. *arXiv preprint arXiv:2609.09134*, 2026d. URL https://arxiv.org/abs/2609.09134.

L. Yuan, Y. Cai, W. Liu, S. Yang, F. Han, Y. Lee, U. M. Sehwag, and F. Huang. Meta^2: Precomputing Reusable Knowledge for Online Harness Evolution in Long-Horizon Agents. *arXiv preprint arXiv:2611.05139*, 2026. URL https://arxiv.org/abs/2611.05139.

H. Zhang, S. Zhang, K. Li, C. Zhang, et al. Self-Harness: Harnesses That Improve Themselves. *arXiv preprint arXiv:2606.09498*, 2026a. URL https://arxiv.org/abs/2606.09498.

J. Zhang, S. Hu, C. Lu, R. Lange, et al. Darwin Gödel Machine: Open-Ended Evolution of Self-Improving Agents. In *International Conference on Learning Representations*, 2026b. URL https://openreview.net/forum?id=pUpzQZTvGY.

J. Zhang, S. Liu, H. Wang, T.-E. Lin, et al. Causal Improvement Graph for Agentic Harness Optimization. *arXiv preprint arXiv:2610.05039*, 2026c. URL https://arxiv.org/abs/2610.05039.

J. Zhang, B. Zhao, W. Yang, J. Foerster, et al. Hyperagents. *arXiv preprint arXiv:2603.19461*,

2026d. URL https://arxiv.org/abs/2603.19461.

Q. Zhang, C. Hu, S. Upasani, B. Ma, et al. Agentic Context Engineering: Evolving Contexts for Self-Improving Language Models. In *International Conference on Learning Representations*, 2026e. URL https://openreview.net/forum?id=eC4ygDs02R.

T. Zhang, H. Wang, K. Xu, and D. N. Metaxas. Turbo Harness: Instance-Adaptive Harness Optimization. *arXiv preprint arXiv:2609.40330*, 2026f. URL https://arxiv.org/abs/2609.40330.

X. Zhang, J. Zhao, and Y. LeCun. Character-level Convolutional Networks for Text Classification. In *Advances in Neural Information Processing Systems*, volume 28, pages 649–657, 2015. URL https://arxiv.org/abs/1509.01626.

A. Zhao, D. Huang, Q. Xu, M. Lin, Y.-J. Liu, and G. Huang. ExpeL: LLM Agents Are Experiential Learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. URL https://arxiv.org/abs/2308.10144.

S. Zhao, Y. Zhang, and J. Li. Can Coding Agents Sustain Continual Evolution without Forgetting? *arXiv preprint arXiv:2611.03219*, 2026. URL https://arxiv.org/abs/2611.03219.

C. Zheng, C. Xue, B. Liang, J. Yang, et al. SEAGym: An Evaluation Environment for Self-Evolving LLM Agents. *arXiv preprint arXiv:2606.17546*, 2026. URL https://arxiv.org/abs/2606.17546.

Guide to the Appendices

Appendices B–D specify evaluation, experimental setup, and SIM implementation. Appendix E gives additional numerical results and ablations; Appendix F explains the observed edits through case studies.

Appendix	Contents
A	**Extended Related Work**: harness optimization, transfer, and memory.
B	**Benchmarks and Evaluation**: datasets, splits, feedback access, and scoring.
C	**Experimental Setup**: models, seeds, budgets, comparison conditions, and repeated runs.
D	**SIM Implementation**: how agents read, update, and store SIM.
E	**Additional Results**: source and target scores, source retention, and ablations.
F	**Case Studies**: source RSI, continual adaptation, memory records, and transfer patterns.

A. Extended Related Work

This appendix expands §2 around three questions: how harnesses are improved, how optimization continues after a change of setting, and what role memory plays in that process.

A.1. How Are Agent Harnesses Improved?

Harness RSI optimizes prompts, tools, control flow, memory, and context management around a policy model. GEPA uses trajectory reflection and recombination to optimize prompts (Agrawal et al., 2026). Meta-Harness exposes prior implementations, execution traces, and scores through a shared filesystem (Lee et al., 2026). AHE makes harness components and modification outcomes observable, checking edit-level predictions against subsequent task outcomes (Lin et al., 2026). In ML development, MLToolBench learns diagnostic tool use through supervised and reinforcement learning (Yu et al., 2026c).

Other methods change how improvement itself is performed. Self-Harness combines weakness mining, harness proposals, and regression-based validation (Zhang et al., 2026a). DGM and Hyperagents explore self-modification and open-ended evolution of agent programs (Zhang et al., 2026b,d); Hyperagents also reports transfer of meta-level improvements across domains. Harness-R1 learns a harness-editing policy from failure trajectories using reinforcement learning (Shao et al., 2026), while SIA updates both the harness and the task model's weights (Hebbar et al., 2026). Bai et al. (2026) train idea-level critics to screen ML modifications before costly empirical verification. MLToolBench learns task-agent behavior, while the critics learn candidate assessment; SIM preserves edit histories for the proposer and selector. CRSI defines continuation after a benchmark or policy change for compatible proposal–evaluation–selection loops (§3.1).

A.2. How Does Optimization Continue after Transfer?

Transfer of evolved harnesses. AHE evaluates transfer from Terminal-Bench to SWE-bench Verified without re-evolving the harness (Lin et al., 2026); Meta-Harness and PluginRSI report transfer across solver models (Lee et al., 2026; Shi et al., 2026). Several methods aim to improve this transfer. ModularRSI uses benchmark-disjoint evolution data, contrasts successful and failed trajectories, and directs edits toward recurring deficiencies in functional modules (Wu et al., 2026). RRSI constrains bundled modifications and redirects stalled search toward underexplored components (Xia et al., 2026). Harmony Harness separates general mechanisms from benchmark-specific adaptations (Yu et al., 2026a). The interaction between a model and its harness can nevertheless require model-specific adaptation (Yang et al., 2026; Yu et al., 2026d).

Continued optimization. Prior work also reuses improvement experience during further optimization. PluginRSI reuses an evolved plugin library when optimizing from the initial harness (Shi et al., 2026). Xu (2026) continues multi-task harness evolution on a new benchmark with free-text evolution notes and source-

task replay. Harness^2 retains adopted harnesses and refinement lessons in a library and tests its reuse across benchmarks (Xu et al., 2026a). Related settings include task streams and continuous episodes (Karten et al., 2026; Wei et al., 2026; Xiao et al., 2026). Adaptive Auto-Harness separates a persistent evolution workspace from deployable harnesses, with optional routing among stored branches (Liu et al., 2026). Evo-HarnessBench studies retention and adaptation as externally supplied tools, skills, and specialist agents expand (Ke et al., 2026).

Evaluating continuation. Held-out evaluation and baselines with matched feedback and inference budgets help distinguish improvements from additional optimization resources (Wang et al., 2026). SEAGym tracks validation, transfer, replay, and cost throughout evolution (Zheng et al., 2026). Studies of shifting task streams and coding benchmarks also document interference and forgetting (Chen et al., 2026; Zhao et al., 2026). We report the evolved harness's score before and after target RSI, compare memory conditions from the same starting harness with matched round and candidate budgets, and separately track source retention in one coding transition (Appendix E.3).

A.3. What Does Memory Store, and Who Uses It?

Memory for harness optimization. Keeping improvement history has several precedents. Meta-Harness retains candidate artifacts (Lee et al., 2026); RRSI records hypotheses, diffs, score and cost changes, and acceptance decisions, including rejected edits (Xia et al., 2026). The Causal Improvement Graph links evidence, hypotheses, interventions, and outcomes (Zhang et al., 2026c). SkillHone preserves diagnoses, revisions, and evaluation evidence (Li and Hu, 2026), and Turbo Harness summarizes completed optimization into a playbook for instance-specific edits (Zhang et al., 2026f). Harness-Bank groups candidates by modified component and diagnosed failure, retaining high-performing harnesses for recombination (Luo et al., 2026). Organizing improvement around problems therefore also has precedent. Meta-Skill freezes learned support-design principles to construct a harness for each test task (Qian et al., 2026). Meta^2 freezes reusable evolution guidance for edits during task execution, resetting the harness between task runs (Yuan et al., 2026).

Memory for task solving. Reflexion stores verbal reflections to guide later attempts (Shinn et al., 2023), and Voyager maintains a growing library of executable skills (Wang et al., 2024). ExpeL extracts reusable insights from task experience (Zhao et al., 2024); ReasoningBank and ACE retain strategies and lessons (Ouyang et al., 2026; Zhang et al., 2026e). ACE incrementally extends, refines, and deduplicates its playbook from execution trajectories. Memory Transfer Learning compares task-experience representations across coding domains and models, finding abstract insights more transferable than concrete traces in its studied settings (Kim et al., 2026). These memories support the agent executing tasks.

SIM in continual harness optimization. SIM supplies the proposer and selector with a record organized by harness problem, linking selected modifications to their rationale, predicted and observed outcomes, and evaluation context (§3.2). The record accompanies the evolved harness across a benchmark or model transition and gains new attempts during target RSI. Our comparisons with no SIM and root causes only test this complete procedure, combining inherited history with target-side updates (Appendix E.4.1).

B. Benchmarks and Evaluation

This section specifies the protocol for the main cross-dataset and Opus-to-Gemini comparisons and the paired coding study. We list the benchmarks and split sizes, explain which feedback can guide optimization, and then define the scoring rules for each domain. The model-generation and shared-harness experi-

ments have additional settings described in Appendices C.5 and C.6, respectively.

B.1. Which Benchmarks Do We Use?

Benchmarks. We study four domains in which a source-evolved harness remains executable in the target environment. For *coding*, Terminal-Bench 2.1 (TB2.1) is the source and Terminal-Bench 4 (TB4) the target (Merrill et al., 2026; Terminal-Bench Team, 2026a,b). For *text classification*, we use the online classification environment of Meta-Harness (Lee et al., 2026). The source suite combines LawBench (Fei et al., 2024), the Gretel adaptation of Symptom2Disease (Gretel AI, 2023), and USPTO-50k (Schneider et al., 2016). The target suite contains eight datasets absent from source RSI: AG News (Zhang et al., 2015), GoEmotions (Demszky et al., 2020), Banking77 (Casanueva et al., 2020), Financial PhraseBank (Malo et al., 2014), SciCite (Cohan et al., 2019), TweetEval-Hate (Barbieri et al., 2020), Amazon Reviews (Amazon5) (Keung et al., 2020), and SciTail (Khot et al., 2018). For *agentic workspace tasks*, Harvey LAB (Harvey AI, 2026) is the source. JobBench (Li et al., 2026), GDPval (Patwardhan et al., 2026), and APEX-Agents (Vidgen et al., 2026) jointly form the target environment. For *engineering design*, EngDesign (Guo et al., 2025) is the source and Frontier-Eng (Chi et al., 2026) the target. Table 4 gives the task inventory and rollout allocation. Appendix B.2 defines the split and feedback-access rules; Appendices B.3–B.6 detail each domain's scoring protocol.

Source and target roles. In *cross-dataset CRSI*, the policy model is fixed while the benchmark changes from the source to the target listed above. In *cross-model CRSI*, the benchmark and its task split are fixed while the policy changes: source RSI is run with the source policy on that benchmark, and target RSI continues with the target policy. The Opus-to-Gemini comparisons use the corresponding T inventory in Table 4 for both stages.

B.2. How Are the Splits Used?

The tasks (examples for text classification) in Table 4 are split uniformly at random into approximately equal evolve and held-out halves, without stratification by category, difficulty, or source dataset. When a benchmark holds an odd number of tasks, the evolve half takes the even count. The two sets are disjoint, and no task or rollout is removed from the scoring denominator after the split.

Only evolve-set scores and execution trajectories are available for proposals, selection, SIM updates, and stopping decisions. The task-solving policy executes the harness on held-out tasks to measure starting, final, and per-round diagnostic scores (§4.4). These held-out scores and trajectories are not passed to the harness-optimization agents.

Within these primary comparisons and the paired coding study, the scoring allocation is five rollouts per task on both splits, matched across conditions; the overlapping Frontier-Eng component contributes zero under the rule in Appendix B.6. The reported harness versions, independent-run counts, and statistical comparisons are specified in Appendix C.4.

B.3. Coding: Terminal-Bench 2.1 and 4

Environment and metric. Terminal-Bench tasks run in containers with a task description, shell access, and task-specific tests. We report pass@1, estimated from the binary outcome of each task rollout. For n tasks with five rollouts each, let $y_{ir} \in \{0, 1\}$ indicate whether rollout r of task i passes the task verifier. The percentage score is

$$\widehat{\text{pass@1}} = \frac{100}{n} \sum_{i=1}^{n} \frac{1}{5} \sum_{r=1}^{5} y_{ir}. \tag{4}$$

Failed rollouts. A rollout that fails for infrastructure reasons stays in the denominator and scores zero.

Domain	Benchmark	Total	Evolve tasks / rollouts	Held-out tasks / rollouts	Metric
Coding	S: TB2.1	89	44 / 220	45 / 225	pass@1
	T: TB4	66	33 / 165	33 / 165	pass@1
Text cls.	S: 3 datasets	412	206 / 1,030	206 / 1,030	accuracy
	T: 8 datasets	1,000	500 / 2,500	500 / 2,500	accuracy
Workspace	S: Harvey LAB	160	80 / 400	80 / 400	rubric
	T: JobBench	65	32 / 160	33 / 165	rubric
	T: GDPval	220	110 / 550	110 / 550	win rate
	T: APEX-Agents	480	240 / 1,200	240 / 1,200	pass@1
	T: merged	765	382 / 1,910	383 / 1,915	mean of 3
Engineering	S: EngDesign	67	34 / 170	33 / 165	pass rate
	T: Frontier-Eng	47	24 / 120	23 / 115	medal

Table 4 | Benchmark inventory used in our experiments, with source (S) and target (T) roles for cross-dataset CRSI. Split sizes are tasks / rollouts (examples / rollouts for text classification). Each inventory is split uniformly at random into approximately equal halves, with five nominal rollouts per task. Frontier-Eng's overlapping EngDesign component receives no credit but remains in the denominator. The three workspace targets are evolved jointly and reported as their equal-weight mean; the merged row sums their task and rollout counts.

B.4. Text Classification

The online classification environment follows Meta-Harness: the harness controls example storage, retrieval, and prediction context. The example counts in Table 4 are totals across each suite, not allocations for individual datasets.

Example memory during evaluation. The task agent reads only labeled examples supplied in advance from training/evolve data. Held-out labels are used for scoring and are not fed back into its example memory. For a fixed harness, each of the five evaluation repetitions starts from the same initial example-memory state. This memory supplies prediction context to the task agent and is separate from SIM, which supports harness optimization.

Accuracy across examples. We compute accuracy over all examples in a split, pooling the three source datasets or the eight target datasets. For split $q \in \{\text{evo}, \text{eval}\}$, let n_q be the total number of examples and $y_{ir}(H;q) \in \{0,1\}$ indicate whether harness H predicts the correct label for example i in rollout r. The percentage accuracy is

$$A_{\text{cls}}(H;q) = \frac{100}{5n_q} \sum_{i=1}^{n_q} \sum_{r=1}^{5} y_{ir}(H;q). \tag{5}$$

Each example has equal weight, so a dataset contributes in proportion to its example count. The same aggregation is used on the evolve and evaluation splits.

B.5. Agentic Workspace Tasks

Task scoring and judging. We adapt the benchmark-specific scoring interfaces described in RRSI (Xia et al., 2026, Appendix B.3–B.6). Our runs use Gemini 3.8 Flash as the judge for all workspace benchmarks; this differs from RRSI's mixed judge configurations and GDPval voting panel. Harvey LAB reports the fraction of rubric criteria satisfied, and JobBench uses its weighted rubric score. GDPval compares each generated deliverable with the benchmark's expert deliverable and reports the win rate against that reference. APEX-Agents reports pass@1, with success determined by satisfaction of the task's rubric on a single rollout; the repeated rollouts estimate this success probability. The judg-

ing interface and model are fixed across comparison conditions within each benchmark. We express each benchmark score on a 0–100 scale before taking the equal-weight mean.

Equal-weight benchmark aggregation. For split $q \in \{\text{evo}, \text{eval}\}$, let $S_b(H; q)$ be the benchmark-specific percentage score of harness H on benchmark b. The workspace score is

$$S_{\text{workspace}}(H; q) = \frac{1}{3} \sum_{b \in \mathcal{B}} S_b(H; q), \qquad (6)$$

where $\mathcal{B}$ contains JobBench, GDPval, and APEX-Agents. Each benchmark contributes one third, independent of its number of tasks or rubric criteria. The evolve-set performance objective used to select target harnesses and the reported evaluation aggregate use this same weighting. Within each benchmark, scoring retains that benchmark's own evaluation convention; averaging occurs after computing its score. The workspace row of Table 5 reports this aggregate.

B.6. Engineering Design

Frontier-Eng has no MCP gateway of its own and uses the runner and gateway of EngDesign. EngDesign uses a binary verifier pass; Frontier-Eng uses a medal score against frozen thresholds, with gold, silver, and bronze mapped to 1.00, 0.67, and 0.33. Neither engineering metric uses an LLM judge.

Source overlap. Frontier-Eng includes an EngDesign-derived component containing seven sub-problems (Chi et al., 2026). Following RRSI, we assign zero credit to this overlapping component while retaining it in the original denominator, including after the 24/23 split.

C. Experimental Setup

This section specifies the models, seed harnesses, comparison conditions, and optimization budgets. The common setup covers cross-dataset and Opus-to-Gemini transfer (§4.2–4.3.1); the last two subsections describe continuation across model generations and shared-harness construction (§4.3.2–4.3.3). Benchmark scoring is defined in Appendix B, and SIM implementation in Appendix D.

C.1. Which Models and Budgets Are Used?

Task-solving policy. Gemini 3.8 Flash is the policy in both stages of cross-dataset CRSI. The cross-family comparison uses Claude Opus 5.5 as the source policy and Gemini 3.8 Flash as the target policy. Policy weights remain fixed.

Harness-optimization agents. The proposer and selector use Gemini 3.8 Flash, including when Opus executes the tasks. The ablation in Appendix E.4.2 replaces the proposer and selector with Opus during source RSI.

Optimization budget. Each RSI run uses $N = 10$ rounds with $k = 5$ candidate proposals per round. Candidate proposals and task rollouts are separate budgets: the main comparisons use five rollouts per task on each split (Appendix B). Proposals screened out before execution receive no rollout scores.

Cost per round. A coding round takes about 4 h, including 1.7 h for evaluation and 34 min for candidate sampling. We match rounds and candidate counts across methods; token use and runtime can differ.

C.2. Which Harnesses Do We Start From?

We use publicly available seed harnesses from prior work. The two seeds in coding and text classification test different starting harnesses; repeated optimization runs are described in Appendix C.4.

Coding. The AHE seed starts with a basic shell tool and allows edits to prompts, tools, middleware, skills, sub-agents, and memory (Lin et al., 2026). Terminus-2 is distributed with Harbor and is also used by RRSI (Merrill et al., 2026; Xia et al., 2026).

Text classification. NoMemory stores no training examples; FewShotAll stores the observed labeled examples and includes them as few-shot context (Lee et al., 2026).

Workspace and engineering. We use RRSI's ReAct toolbelt agent for file manipulation, PDF reading, and code execution (Vidgen et al., 2026; Xia et al., 2026; Yao et al., 2023). Its 6 files (1,743 lines) are shared by the source and all three targets, allowing the Harvey LAB harness to run on the targets without modification. Engineering uses the same agent with a bare system prompt and task wrapper, evaluated by fixed task-specific verifiers (Appendix B.6). This seed has 7 files (1,857 lines), with prompts in a separate module. For these two domains, we adopt RRSI's environments, seed harnesses, and evaluation interfaces. Our models and budgets are specified above; our splits, rollout counts, and workspace judge are specified in Appendix B.

C.3. How Do the Four Methods Differ?

Figure 6 summarizes the four conditions from §3.3:

- **Cold Start**: the seed harness H_0, without SIM.
- **Harness Transfer**: the source-evolved harness H_S, without SIM.
- **Problem-Only Memory**: the same H_S with root causes only, omitting modifications, outcomes, and provenance.
- **Full SIM**: the same H_S with the complete source record $\mathcal{M}_S$, updated during target RSI.

Each condition persists throughout target RSI. All methods retain the ordinary target-run history $\mathcal{D}_t$. All four use the same target policy, evolve tasks, round and candidate budgets, and evaluation allocation, and receive feedback through the same RSI loop. Harness Transfer uses the evolved-harness starting point studied in prior work (Lee et al., 2026; Lin et al., 2026; Xu, 2026).

Condition	Harness	SIM during target RSI
Cold Start	H_0	None
Harness Transfer	H_S	None
Problem-Only	H_S	Root causes only
Full SIM	H_S	Full record; updated on target

Figure 6 | Transfer conditions. The last three rows share the source-evolved harness. The listed memory condition persists throughout target RSI; the full record contains root causes, edits, outcomes, and provenance. Target round and candidate budgets match.

C.4. How Are Runs Reported and Repeated?

Reported harnesses. Each round retains one harness for the next round; H_S is the final retained source harness. The *before* score evaluates the harness used to initialize target RSI, the *after* score evaluates the version retained at the end of the final target round, and the target improvement is their difference. Curves show each retained version's measured score, without taking a running maximum over earlier scores.

Independent runs. Repeated task rollouts and independent optimization runs are distinct sources of variation. Each main cross-dataset and Opus-to-Gemini condition uses one complete RSI run. The paired coding study repeats source RSI ten times from Terminus-2; each evolved harness initializes two target continuations, with and without SIM (§4.4, Table 3). Task rollouts estimate execution performance. The ten independent source runs support across-run intervals for the paired study; scores from single-run conditions are reported descriptively. For the paired study, the paired t interval and two-sided test use the ten final-score differences; the signed-rank check uses average ranks for ties and enumerates the 2^{10} sign assignments.

C.5. How Is Transfer across Generations Tested?

We continue optimization through Gemini 2.5 Flash (Gemini Team, 2025), 3.5 Flash (Google DeepMind, 2026a), 3.8 Flash (Google DeepMind, 2026b), and 4 Argon (Google DeepMind, 2026c), keeping TB2.1 as the benchmark. Fresh RSI starts from the seed harness at each generation.

CRSI starts from the same seed in the first generation, then passes the evolved harness and SIM to each subsequent model. The first row tests SIM during an initial RSI run; only the later rows test transfer from an earlier generation.

Table 1 reports scores to one decimal, with the results discussed in §4.3.2.

External scores use different evaluation settings. For context, the published score is 16.9 for Gemini 2.5 Flash with Terminus-2 on TB2.0 (Merrill et al., 2026), and 76.2 and 89.4 on TB2.1 for Gemini 3.5 and 3.8 Flash (Google DeepMind, 2026a,b). These are external references, not matched controls. We include no external score for Gemini 4 Argon.

C.6. How Is a Shared Harness Tested?

We test a harness shared by Grok 4.20 Reasoning, Claude Opus 5.5, and Gemini 3.8 Flash on the text-classification target suite (Table 2). The evolve and held-out sets each contain 500 examples, with five rollouts per example (2,500 rollouts per split).

Harness construction. We compare two baselines and three ways to build a shared harness:

- **B1: Per-model harnesses.** Each policy model runs its own RSI and is evaluated with its own evolved harness.
- **B2: Shared seed.** All three models use the same unevolved harness.
- **S1: Sequential CRSI.** The order is Grok 4.20 Reasoning, Gemini 3.8 Flash, then Claude Opus 5.5, from weaker to stronger policies in this setting. Each model passes the evolved harness and SIM to the next.
- **S2: Unified synthesis.** Each round pools trajectory feedback from all three models to propose updates to one shared harness.
- **S3: Single-agent transfer.** We evolve a harness separately with each of the three policy models, then evaluate each resulting harness with all three models.

This experiment tests the performance of a common harness; Table 6 separately tests the effect of transferring SIM from a matched starting harness.

Score aggregation. All averages over models use equal weights. B1 averages the three model–harness pairs, each model using its own harness. B2, S1, and S2 average the three models using a shared harness. Each S3 row names the source policy used to evolve a harness and reports its mean score across the same three evaluation models. The All column is $(\text{Evolve} + \text{Held-out})/2$, the average over the two equally sized splits.

Each procedure is run once, and optimization cost is not compared. The results in §4.3.3 describe the observed means; small differences do not establish statistical equivalence or a ranking for each model.

D. SIM Implementation

We supplement §3 with entry-matching details, prompt guidance, and the stored record format.

D.1. How Are Entries Matched?

For the update in §3.2, the proposer itself decides whether a diagnosis matches an existing root cause; no separate similarity model is used. Its choice is saved in the round's change manifest before evaluation. Entries are merged deterministically when their root-cause text is identical, preserving all attempts.

D.2. How Do the Proposer and Selector Use SIM?

The prompt guidance for §3.3 is summarized below:

1. Treat entries as prior experience, with flexibility to choose a different edit.
2. Check the domain, benchmark, and model behind each entry.
3. Start with an existing diagnosis when it explains the current failure.
4. Reuse the rationale and reassess the edit after a benchmark or model change.

5. Revise, replace, or override a harness modification when new evidence warrants it.
6. Revisit a tested direction only when new evidence justifies it.
7. Create a new problem entry when existing diagnoses do not explain the failure.
8. Give task-specific evidence more weight in similar settings and less when both benchmark and model change.

D.3. What Does a SIM Record Contain?

Each run stores SIM in one JSON file, with the following schematic structure:

```
{ "next_id": 10,
  "carried_from": { "arm": <source run>, "domain": "
     coding_enriched",
                    "into": <target run> },
  "aliases": { "P6": "P5" },              // P6 was
     merged into P5
  "problems": [ ...entries P1 P2 P3 P4 P5 P7 P8 P9... ] }
```

Each problem entry contains `{id, name, opened_round, attempts[]}`. `name` gives the root-cause hypothesis, and `opened_round` records its first diagnosis. Each attempt contains: `{round, sample, accepted, backfilled, what, files, predicted, fixed, missed, passed, total, domain}`. `what` and `files` describe the edit, and `domain` identifies its setting.

`predicted` lists the tasks expected to be fixed before evaluation; `fixed` and `missed` record the outcomes afterward. `passed/total` gives the successful and total trials in that attempt's evolve-set evaluation. `accepted` means the selector chose the edit in that round, even if some predicted fixes failed. The edit may later be changed or replaced while its record stays in SIM, as in the P8 case in Appendix F.2. Predictions for the same task in different rounds are separate observations.

`carried_from` identifies the source and target runs; `aliases` keeps the identifiers of merged entries (for example, P6→P5). Read each entry's ID and `opened_round` within its recorded run context. Appendix F gives records as stored, including P3 in Figure 12.

E. Additional Results

We report source-RSI scores (§E.1), exact target scores (§E.2), source retention (§E.3), and memory and harness-agent ablations (§E.4). The model-generation, shared-harness, and ten-pair coding results are in the main text (Tables 1, 2, and 3); edit histories and cases are in Appendix F.

Evaluation and reporting follow Appendices B and C.4; differences are absolute score points.

E.1. Does Source RSI Improve the Seed Harness?

Source RSI improves all six held-out source scores. Source RSI improves the held-out source score in all six settings, by 2.75–6.41 points (Figure 7). In five settings, the held-out gain exceeds the evolve-set gain; the exception is FewShotAll, with held-out and evolve-set gains of 6.41 and 6.60 points. Section 4.2 examines whether these source gains yield a better target starting point.

E.2. How Do the Methods Compare on the Target?

Tables 5 and 6 give the numerical results plotted in Figures 3 and 4, respectively.

E.3. How Much Source Performance Is Retained?

Both memory conditions retain more source performance in this run. In Figure 8, Full SIM and Problem-Only Memory end at 88.89 and 88.44, respectively, from the common starting score of 90.22. Their lowest source score is 88.44, a drop of 1.78 points. Harness Transfer falls to 86.67 in round 2 and ends at 87.56; Cold Start begins at 85.78 and ends at 85.33. Section 4.4 relates this trace to target progress: the two memory conditions retain similar source performance, while Full SIM improves more on TB4.

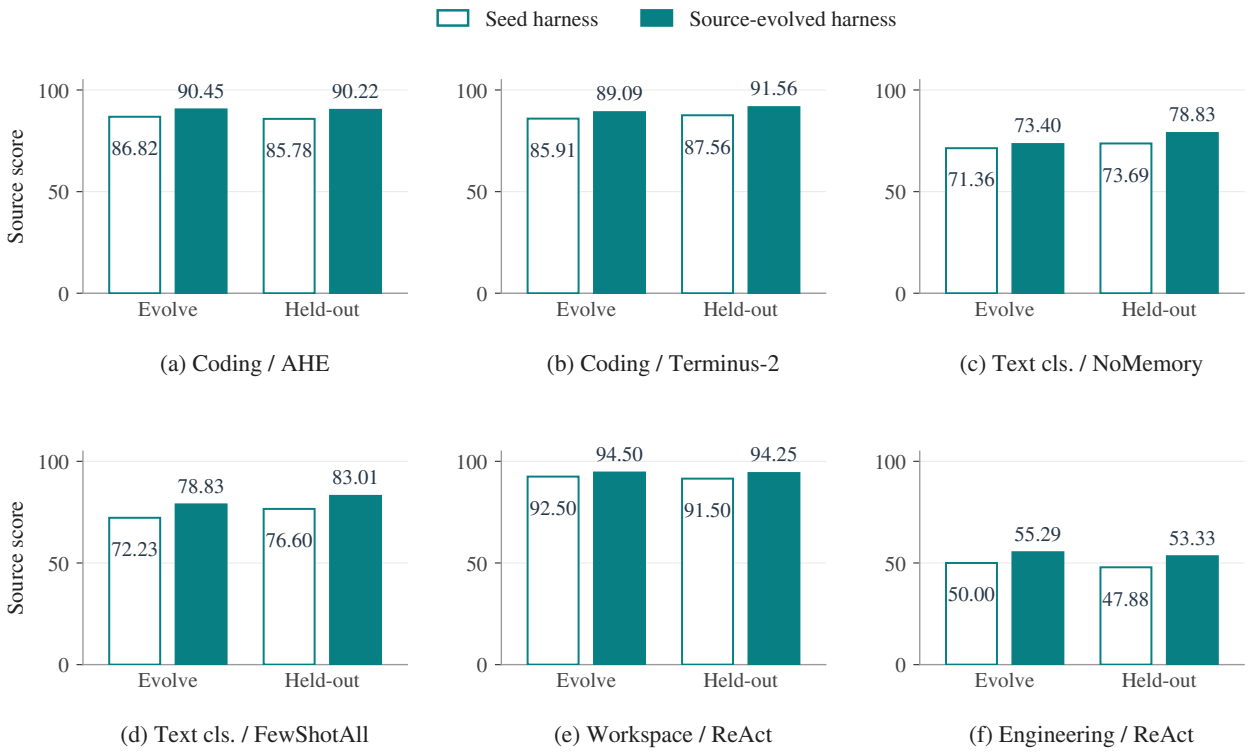

Figure 7 | Source RSI with Gemini 3.8 Flash: scores of the seed harness H_0 and the source-evolved harness H_S on the source evolve set and the held-out source evaluation set. (a) Coding on TB2.1 from the AHE seed; (b) coding on TB2.1 from Terminus-2; (c) text classification on the source suite from NoMemory; (d) the same from FewShotAll; (e) workspace tasks on Harvey LAB; (f) engineering design on EngDesign. H_S is the starting harness of Harness Transfer, Problem-Only Memory, and Full SIM in Figure 3. Open and teal bars show seed and source-evolved scores for each split, with both values labeled. All panels share a zero baseline and the same score scale.

	Cold Start		Harness Transfer		Problem-Only Memory		**Full SIM (ours)**	
Domain / seed	Before	After	Before	After	Before	After	Before	After
Coding / AHE	16.36	19.39	18.18	20.61	18.18	18.79	18.18	**24.85**
Coding / Terminus-2	18.18	19.39	19.39	19.39	19.39	20.61	19.39	**21.82**
Text cls. / NoMemory	68.60	72.00	71.40	71.00	71.40	73.40	71.40	**73.60**
Text cls. / FewShotAll	75.72	75.92	72.92	74.12	72.92	75.52	72.92	**77.92**
Workspace / ReAct	73.89	82.98	74.93	85.07	74.93	85.54	74.93	**87.89**
Engineering / ReAct	39.13	40.87	39.13	42.61	39.13	42.61	39.13	**46.96**

Table 5 | Cross-dataset CRSI: held-out target scores before and after target RSI for the six settings of Figure 3. The first column identifies the domain and seed harness. Bold marks the highest final score in each setting.

E.4. What Do the Ablations Show?

E.4.1. Is Recording Only Root Causes Enough?

From the same evolved harness, Full SIM finishes above both no-memory and root-cause-only continuation in all ten settings (Tables 5 and 6). This comparison tests each memory condition throughout target RSI, covering both inherited information and target-stage updates

	Cold Start		Harness Transfer		Problem-Only Memory		**Full SIM (ours)**	
Domain / seed	Before	After	Before	After	Before	After	Before	After
Coding / AHE	16.36	19.39	13.94	18.18	13.94	20.61	13.94	**21.82**
Text cls. / FewShotAll	78.32	82.72	77.72	80.20	77.72	81.32	77.72	**83.80**
Workspace / ReAct	73.89	82.98	68.41	73.42	68.41	79.22	68.41	**84.75**
Engineering / ReAct	39.13	40.87	34.78	42.61	34.78	46.96	34.78	**50.43**

Table 6 | Cross-model CRSI (Claude Opus 5.5 to Gemini 3.8 Flash): exact held-out scores for Figure 4. Before and After refer to Gemini's target RSI. Both stages use the same benchmark and split; the three transfer conditions share the Opus-evolved starting harness. The first column identifies the domain and seed harness. The Terminus-2 Cold Start counterpart for coding scores 18.18 before and 19.39 after target RSI. Bold marks the highest final score in each row.

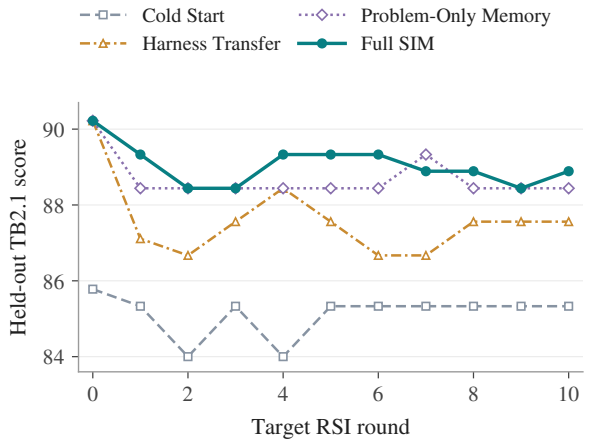

Figure 8 | Source scores during target RSI (coding, TB2.1→TB4, Gemini 3.8 Flash): held-out TB2.1 score (45 tasks, 225 rollouts; one rollout is 0.44 points) of the harness retained after each target round from the AHE seed, with one run per condition. The three transfer conditions start from the source-evolved harness; Cold Start begins from the seed.

(Appendix C.3). The main results discuss the margins and the mixed effects of retaining root causes alone (§4.2 and §4.3.1).

E.4.2. Does the Choice of Harness Agent Matter?

Opus scores higher, with a smaller held-out margin on Terminus-2. Table 7 replaces Gemini 3.8 Flash with Claude Opus 5.5 as the proposer and selector during TB2.1 source RSI. The held-out score is 1.06 points higher with Opus for AHE and 0.05 points higher for Terminus-2. This experiment varies the optimization agents; the policy-model transitions are studied separately in §4.3.

F. Case Studies

We distinguish improvements within a fixed setting from continued optimization after transfer. The source-RSI cases show changes to command transport and parsing (§F.1). The CRSI cases follow completion control from TB2.1 to TB4 (§F.2), compare continuations with and without SIM (§F.3), and trace new evidence appended to P3 (§F.4). These are complementary views of the coding transfer study, not independent repetitions. We then summarize transfer patterns in the four domains (§F.5) and draw together the lessons for SIM (§F.6). The detailed coding cases use Terminus-2 and Gemini 3.8 Flash.

Evidence in the cases. These cases document edit histories rather than estimate each edit's contribution to the final score. Code fragments are simplified; displayed SIM records retain their stored text, including truncations. Appendix D.3 defines the fields.

F.1. How Does RSI Improve the Source Harness?

These examples stay within source RSI: feedback identifies a harness failure and leads to

			Gemini 3.8 Flash		Claude Opus 5.5	
Set	**Seed**	**Before**	After	Gain	After	Gain
Evolve	AHE	86.82	90.45	+3.63	**91.23**	+4.41
	Terminus-2	85.91	89.09	+3.18	**92.21**	+6.30
Held-out	AHE	85.78	90.22	+4.44	**91.28**	+5.50
	Terminus-2	87.56	91.56	+4.00	**91.61**	+4.05

Table 7 | Source RSI on TB2.1 with Gemini 3.8 Flash or Claude Opus 5.5 as the harness agent, the optimization agent that proposes and selects harness edits: the score of the seed harness, and the score after source RSI with each harness agent. Gain is the change from Before; bold marks the higher After score within each row.

an edit, without a benchmark or policy change. The transport case includes a measured task outcome; the parser example documents the repair mechanism.

F.1.1. `adaptive-rejection-sampler`*: unreliable command transport*

Task. Implement adaptive rejection sampling in R as `/app/ars.R`, with input validation and tests.

Root cause (*c*). The agent wrote the whole program in a single shell command of about 22.8 KB, exceeding the command-transport limit reported in this run (about 16 KB). The command never ran, so the file was never created and the verifier failed with `No ars.R file found`.

Modification (*a*). The transport behind the existing `run_shell_command` tool now splits literal input into chunks of at most 2,048 characters, sends special keys such as `Enter` as separate key events, and raises an error when a send fails instead of failing silently (simplified):

```
MAX_KEY_CHUNK_SIZE = 2048
def _tmux_send_keys(self, keys):
    cmds = []
    for key in keys:
        if key in SPECIAL_KEYS:  # key event
            cmds.append(f"{prefix} {quote(key)}")
            continue
        for chunk in self._split_key(key):
            cmds.append(f"{prefix} {quote(chunk)}")
    return cmds
```

Outcome (*o*). Successful rollouts increased from 1 to 5 out of 5 on this task.

F.1.2. `gpt2-codegolf`*: fragile tool-call parsing*

Task. Write a compact GPT-2 inference program in C (`/app/gpt2.c`) that loads the weights, decodes greedily, and compiles with `gcc -O3 -lm` under a code-size limit.

Root cause (*c*). Terminus-2 expects a JSON response whose `keystrokes` fields hold shell input. When the model wrote C code containing unescaped quotation marks into that field, for example

```
"keystrokes": "...FILE *f = fopen("model.bin", "rb")
    ;..."
```

`json.loads` failed and the command was rejected. The parser could already add missing closing braces and strip surrounding text, but not repair quotes inside a field.

Modification (*a*). A repair step, run only after normal parsing fails, locates the real closing quote of each string field as the one followed by the next expected key (for `keystrokes`, `"duration"` or the closing brace) and escapes every quote before it.

Outcome (*o*). The change adds a recovery path for malformed calls; no post-edit task success rate is reported for this parser repair. Its dependence on the model's output format is discussed in §F.5.1. The P3 record in §F.4 concerns a separate prompt intervention.

F.2. How Does CRSI Adapt an Evolved Harness?

This case follows completion control across a benchmark change, with the policy fixed. Source RSI builds an execution check on TB2.1; target RSI inherits it and adds checks for solution quality on TB4. P5 records the source problem of premature completion; P8 adds a target diagnosis about solution quality. Both concern the inherited completion gate, but they are separate entries, with no explicit parent/refinement link.

F.2.1. Source: Execution-Verified Completion

Task and evaluation interface. *"Configure a git server so that from my own machine I can* `git clone user@server:/git/server`*, commit* `hello.html`*, push, and then* `curl http://server:8080/hello.html` *shows 'hello world'."* The agent must build four parts that compose: a bare repository at `/git/server`, a `post-receive` hook that deploys to the web root, an SSH-reachable account named `user`, and an HTTP service on port 8080. The verifier runs outside the container and clones as `user` over SSH, so a setup that works locally is not evidence that the externally visible interface is correct.

Observed failure and diagnosis. In the described failing rollout, the agent configured the repository, hook, and web server, read *"I will set up login myself"* as meaning the account would be provisioned externally, and never created `user`. It tested the flow with the pre-existing `ubuntu` account, saw every local command succeed, and declared `task_complete: true`. SSH then refused the nonexistent user, the push failed, and the page returned 404. The seed harness could not catch this, because its completion check only asks the model to repeat its claim (simplified from `terminus_2.py`):

```
if is_task_complete:
    if self._pending_completion:   # second claim
        return episode + 1         # terminate
    self._pending_completion = True  # first: ask again
    obs = self._confirmation_message(out)
```

and the confirmation message asks about intent, not evidence: *"Are you sure you want to mark the task as complete? ... If so, include* `"task_complete": true` *in your JSON response again."* **Root cause:** ending the episode is an agent privilege that requires no execution evidence.

Harness intervention. The intervention introduces completion-control middleware in the harness. Source RSI built it over three accepted rounds:

Round	Edit
r6	New `execution_gating_interlock.py` (137 lines), wired into `terminus_2.py`: a cross-turn state machine that refuses completion when the previous turn had a parse error or no command actually executed, capped at three rejections.
r8	Also intercepts simulated multi-turn text, truncated output with discarded command blocks, and completion after at most one real command (adds `terminus_json_plain_parser.py`).
r10	Detects terminal crashes, assertion failures, HTTP 404s, and test errors at the moment completion is requested (adds both prompt templates).

The interlock tracks execution across turns and uses four rejection paths (simplified):

```
def check_completion_attempt(self, cmds, out, pending=
     False):
    if self._consecutive_rejections >= self._max_rej:
        return False, ""  # never lock out a done task
    if self.detect_terminal_error(out):      # Case 0
        return self._reject(...)
    if pending:
        return False, ""
    ran = any(c.keystrokes.strip() for c in cmds)
    if self._last_turn_had_error and not ran:   # A
        return self._reject(...)
    if self._last_turn_discarded and not ran:   # B
        return self._reject(...)
    if self._total_executed <= 1 and not ran:   # C
        return self._reject(...)
    return False, ""
```

The interlock enforces the completion checks; prompt reminders provide additional guidance. By the end of source RSI the file had grown to 275 lines.

Reported task outcome. Successful rollouts increased from 1 to 5 out of 5 on

`configure-git-webserver`. The same P5 intervention also raised `circuit-fibsqrt` from 1 to 5 out of 5: that task tests Fibonacci and square-root circuits with `/app/sim.c`, including large-integer inputs, and the agent had stopped after failed commands. These are two task outcomes for the same source intervention.

The three source attempts remain together in P5 (Figure 9). Two target proposals cited P5, but neither was selected, so no target attempt was appended to that entry.

F.2.2. Target: Feasibility and Solution Quality

Task and evaluation interface. Route nine photonic waveguides described in `/app/layout_spec.json` and write the waypoint paths to `/app/routing_result_1.json`. Each path must run from its exact start to its exact target coordinate using 90-degree bends of radius 50 μm and valid s-bends, stay inside the board, avoid obstacle clearance regions and self-intersections, and keep a minimum separation between nets. The score is the negative weighted route cost summed over all nets. The task ships its own feasibility checker, `check_routing.py`, which answers whether the constraints are satisfied, while the reward is a cost function. A solution can pass every check and still score poorly.

Observed failure and diagnosis. The evolved harness contains the source interlock, intact and wired. On this task the agent produces a feasible routing, writes the output, runs `check_routing.py`, sees it pass, encounters no terminal error, and requests completion. None of the four rejection conditions is triggered: commands ran, no parse error occurred, nothing was discarded, and no visible terminal failure was detected, so the gate allows completion. The target failure adds a question about sufficient evidence for completion:

	The source gate asks	**This task also needs**
Question	Has enough executed for the workflow not to have obviously failed?	Is the feasible solution good enough to stop improving?
Evidence	Execution activity, errors, crashes	Objective value against an achievable bound

Harness intervention and installed behavior. Target RSI retained the state machine and bounded rejection, expanding the component from 275 to 437 lines under several target entries. The change attributed to this deficiency, accepted in target round 2, adds a detector for the signature of a constraint-only check that also prints a score, and turns a match into a notice at the confirmation step (simplified):

```
PATTERNS = [re.compile(
    r"(?:ALL CHECKS PASSED|ASSERTIONS PASSED"
    r"|CONSTRAINTS? SATISFIED).*?"
    r"(?:score|cost|loss|metric)\s*[:=]\s*[-+]?\d+",
    re.I | re.S), ...]

def detect_optimization_feasibility_notice(cls, out):
    if any(p.search(out[-4000:]) for p in PATTERNS):
        return ("[ExecutionGatingInterlock NOTICE]: "
                "passing constraint checks (e.g. "
                "check_routing.py) only proves
     FEASIBILITY, "
                "NOT OPTIMALITY ... compare your current
      "
                "score against theoretical bounds ...")
```

The pre-completion checklist gained a matching clause asking the agent to verify that it "has not stopped at the first feasible solution" and has "quantified [its] suboptimality gap against the theoretical bound." In target round 4, the adopted candidate added a fifth rejection path, blocking completion once when feasibility output with a metric was detected. Later edits changed this behavior: the final harness retains the source body of `check_completion_attempt` and its four rejection paths. The error patterns feeding Case 0 were expanded to catch deliverables that report themselves infeasible (e.g. `"route_feasible": false`). The final behavior is therefore to block reported infeasibility and warn about possible suboptimality; SIM retains the earlier blocking attempt and its outcome.

```
{ "id": "P5",
  "name": "premature completion without execution verification",
  "opened_round": 6,
  "attempts": [
    { "round": 6, "sample": "s2", "accepted": true, "domain": "coding_enriched",
      "what": "Implemented ExecutionGatingInterlock in Terminus2 to intercept premature
               task completion attempts when previous commands failed to execute
               (e.g. on unescaped JSON parse errors) or when no commands wer...",
      "files": "execution_gating_interlock.py, terminus_2.py",
      "predicted": ["circuit-fibsqrt", "configure-git-webserver"],
      "fixed":     ["circuit-fibsqrt", "configure-git-webserver"],
      "missed": [], "passed": 49, "total": 60 },
    { "round": 8, "sample": "s1", "accepted": true, "domain": "coding_enriched",
      "what": "Intercept premature completion when LLM generates multi-turn simulated
               text or output-length truncations with unexecuted discarded command
               blocks or completes with trivial session execution (<=1 comma...",
      "files": "execution_gating_interlock.py, terminus_2.py,
                terminus_json_plain_parser.py  +1 more",
      "predicted": ["constraints-scheduling"],
      "fixed":     ["constraints-scheduling"],
      "missed": [], "passed": 45, "total": 60 },
    { "round": 10, "sample": "s1", "accepted": true, "domain": "coding_enriched",
      "what": "Enhance ExecutionGatingInterlock in Terminus2 to detect active terminal
               crashes, assertion failures, HTTP 404s, and test errors when task
               completion is requested, prevent immediate termination when co...",
      "files": "execution_gating_interlock.py, templates/terminus-json-plain.txt,
                templates/terminus-xml-plain.txt  +1 more",
      "predicted": ["configure-git-webserver", "protein-assembly", "train-fasttext"],
      "fixed":     ["configure-git-webserver"],
      "missed":    ["protein-assembly", "train-fasttext"],
      "passed": 47, "total": 60 } ] }
```


Figure 9 | SIM entry P5 as carried into the target run's record (`what` fields truncated as stored). All three attempts are source rounds, identified by `domain: "coding_enriched"`.

Recorded task outcome. The recorded predictions and outcomes distinguish the intended repair from its measured result:

Attempt	Predicted	Fixed
P8, round 2	`photonic-waveguide-routing`, `production-planning`	none
P8, round 4	`photonic-waveguide-routing`, `production-planning`	none

Both attempts were selected, but neither fixed its targeted tasks. Their predictions remain marked as unresolved in P8 (Figure 10). This case illustrates adaptation and memory updates; aggregate target performance is reported separately in §4.2.

The two entries distinguish the inherited execution problem from the new solution-quality problem:

	P5 (carried)	P8 (target-native)
Provenance	`coding_enriched`, source round 6	`tb4`, target round 2
Missing evidence	Execution did not occur or did not succeed	Execution succeeded but the solution is only feasible
Marked accepted	3 of 3, all on the source	2 of 2, both on the target
Realized predictions	4 of 6	0 of 4
Diagnosis opened in	Source RSI	Target RSI
Component	Interlock created	Interlock extended
Completion behavior	Four source rejection paths	Quality advice in final harness

```
{ "id": "P8",
  "name": "premature termination on feasible but suboptimal solutions
           in optimization tasks",
  "opened_round": 2,
  "attempts": [
    { "round": 2, "sample": "s3", "accepted": true, "domain": "tb4",
      "what": "Add objective optimality vs constraint feasibility guidelines to prompt
               templates and enhance ExecutionGatingInterlock with optimality audit
               checklist and feasibility score notice to prevent premature...",
      "files": "execution_gating_interlock.py, templates/terminus-json-plain.txt,
                templates/terminus-xml-plain.txt  +1 more",
      "predicted": ["photonic-waveguide-routing", "production-planning"],
      "fixed": [],
      "missed":    ["photonic-waveguide-routing", "production-planning"],
      "passed": 5, "total": 33 },
    { "round": 4, "sample": "s4", "accepted": true, "domain": "tb4",
      "what": "Active optimization gating and terminal error precision in
               ExecutionGatingInterlock: enforce completion interception when
               constraint feasibility outputs with score/metric are detected
               without verified...",
      "files": "execution_gating_interlock.py",
      "predicted": ["photonic-waveguide-routing", "production-planning"],
      "fixed": [],
      "missed":    ["photonic-waveguide-routing", "production-planning"],
      "passed": 10, "total": 36 } ] }
```


Figure 10 | SIM entry P8, opened by target RSI in round 2 (`what` fields truncated as stored). Its provenance, `domain: "tb4"`, places it inside the target run. The round-4 modification was adopted and later changed; its record remains in SIM.

Memory preserves the history as the harness changes. P5 preserves why the source interlock was introduced; P8 records the target attempts to assess solution quality, including the blocking attempt that was later changed. Together they retain the rationale and outcomes beyond the edits present in the final harness.

F.3. How Does SIM Guide Continuation?

The SIM run adds target checks to the inherited gate. Figure 11b follows the completion gate inherited from TB2.1. It compares the continuation described in §F.2 with a run from the same harness without SIM. In round 2, the proposer with SIM reads the gate's source entry, P5, which diagnoses premature completion without execution verification. The proposer keeps the gate and asks the agent to compare its score with theoretical bounds before declaring completion: satisfying constraints establishes feasibility, but a better solution may still exist. Later edits add checks of specification parameters and lifecycle state. Without SIM, the run edits the same gate four times, twice to repair completion deadlocks. Its round-3 rationale attributes the gate and the deadlocks to round 1, although the gate was inherited from source RSI.

Appendix F.2 gives the task-level evidence and memory records behind these revisions.

F.4. How Does SIM Extend an Existing Record?

Figure 12 shows the record behind the round-5 and round-6 gate edits in Figure 11b. Source RSI diagnoses a mismatch between internally correct deliverables and the external schema or API contract checked by the verifier. The round-4 source edit adds system-prompt guidance and fixes one of four predicted tasks, leaving three unresolved. Target RSI appends two attempts under the same diagnosis: a schema-deliverables and test-audit extension in round 5, and a defensive-contract detector in round 6. They fix zero of two and one of two predicted tasks, respectively.

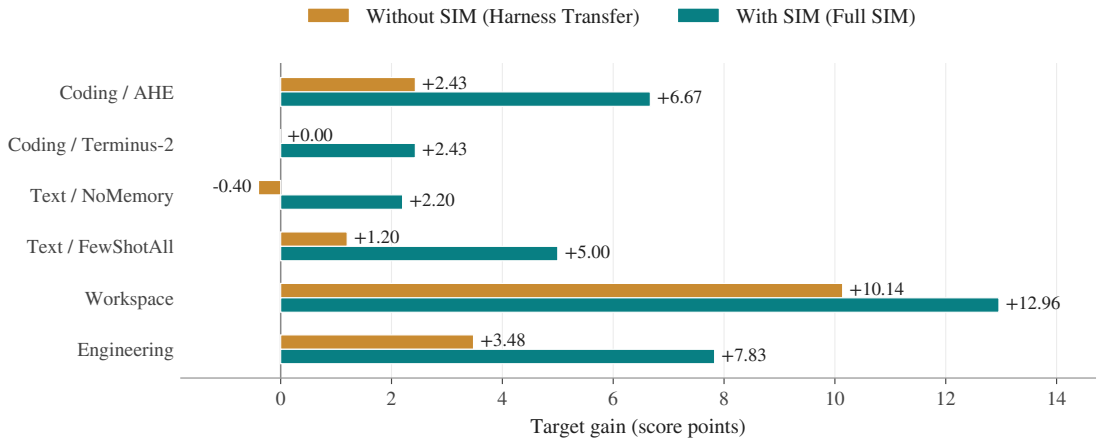

(a) Target gain from the same inherited harness

Round	**With SIM**: same gate + source record	**Without SIM**: same gate, no record
R1	Gate + specification-parameter audit	Gate: stronger error patterns
R2	Gate + feasible $\neq$ optimal check	Prompt: configuration completeness
R3	tmux transport fix	Gate: fix completion deadlocks
R4	Gate + feasibility-score block (P8)	—
R5	Gate + coupled-field and test-suite audit	Gate: cross-turn test-failure latching
R6	Gate + defensive function contracts	—
R7	Gate + infeasibility flags	Gate: fix latched-error deadlock
R8	—	—
R9	Gate + lifecycle-state check	—

(b) Retained edits during target RSI

Figure 11 | Target gains and edit histories with and without SIM. (a) The six cross-dataset target gains from Figure 3, measured from the common source-evolved starting harness. (b) Selected edits in target rounds 1–9 from Terminus-2 (TB2.1→TB4). Both runs inherit the same harness and completion gate. "Gate +" adds a check; — marks no selected edit; round 10 is not listed. The round-4 P8 block was later changed (§F.2); rounds 5 and 6 extend P3 (Figure 12).

P3 therefore retains one diagnosis while accumulating evidence from different interventions and settings.

F.5. What Transfers across Datasets and Models?

Table 8 summarizes which mechanisms are reused, adapted, or extended, and which depend on the target setting. The domain summaries distinguish recorded changes from conditions that limit reuse; a retained mechanism can remain inactive. They broaden the coding cases to the other domains without treating every mechanism as a separate SIM ablation.

F.5.1. Coding

Source RSI revises execution, parsing, and completion. The evolved tool layer retains stdout and stderr after timeouts, checks background processes and their exit codes, recovers misplaced artifacts, and sends long commands through a paste buffer (§F.1). Terminus-2's parser attempts to repair malformed JSON before rejecting it and detects safety refusals. Completion control requires evidence before accepting a completion claim: AHE's middleware forces continuation when the agent stops without a tool call, and Terminus-2's cross-turn interlock blocks completion after unexecuted com-

Domain / shift	Reused directly	Reused with adaptation	Inherited and extended	Setting-dependent
Coding Dataset	Shell, terminal, and process layer	Paste-buffer transport (key normalization, foreground detection)	Completion contract, with new target predicates	Source completion predicates; refusal handling (idle)
Model	Terminal and process layer; interlock accounting	—	Lifecycle control, with recalibrated interception heuristics	Parser repairs tuned to one model's output
Text cls. Dataset	BM25 index and retriever; query sanitization	Top-k, score filter, and class quotas (recalibrated constants)	Multi-template routing skeleton	Data-specific post-processing (SMILES normalization)
Model	Indexing; query sanitization	Retrieval budgets (re-tuned to model capacity)	Routing skeleton, with rewritten rubrics	Output schema and decoding contract
Workspace Dataset	Soft gate before `final_answer`	Evidence probe (file system → invocation history)	Truncation limits and context budget	Evaluator-specific registration assumption; substring heuristics
Model	System-prompt injection pipeline	Budget-triggered sparse intervention	Rubric taxonomy with counterexamples	Calibrations to one model's failure rate or API schema
Engineering Dataset	Central tool-call interception point	Deliverable check → continuous checkpointing	Context buffers and truncation budgets	Prompt guidelines without SIM (reverted)
Model	Deliverable observation layer; re-execution check	Intervention schedules (time and step thresholds)	Composable soft-gate predicates	Intervention intensity; environment-specific bypasses

Table 8 | Mechanism reuse across datasets and models. The last column includes inactive or removed components as well as assumptions and heuristics that require reassessment in a new setting. A dash means that no mechanism is listed. §F.5.1–§F.5.4 give the supporting cases.

SIM entry P3 Policy: Gemini 3.8 Flash

Root cause. A deliverable can be internally correct yet violate the verifier's external contract: schema, coupled attributes, API signature, or split hygiene.

Stage	Selected modification	Recorded outcome
Source TB2.1, round 4 Carried forward	*System prompt.* Execution and Verification Guidelines: deliverable integrity, schema conformance, defensive API compatibility, evaluation hygiene, and invariant checks instead of silent suppression.	**1 of 4 fixed.** Fixed `gpt2-codegolf`; `protein-assembly`, `torch-pipeline-parallelism`, and `train-fasttext` remained unresolved.
Target TB4, round 5 Appended	*Verification layer.* Add a schema-deliverables clause to the completion interlock and prompt templates; extend the checklist to programmatic assertions and a test-suite audit.	**0 of 2 fixed.** Neither `atrx-vep-crispr` nor `intrastat-meldung` was fixed.
Target TB4, round 6 Appended	*Verification layer.* Add a defensive-contract detector to the completion checklist.	**1 of 2 fixed.** Fixed `embedding-drift-monitor`; `biped-contact-dynamics` remained unresolved.

Figure 12 | SIM entry P3 across source and target RSI. The source attempt realizes one of four task-fix predictions; target rounds 5 and 6 realize zero of two and one of two. All attempts retain their outcomes and provenance under the same root-cause hypothesis, including unresolved predictions.

mands or errors (Appendix F.2).

Core execution handling is retained on TB4. The command-execution routines carry over with the shared container and terminal interfaces, while session checks adapt as described below. The recorded cold-start run does not rediscover the tool-layer changes within eight rounds; after ten rounds it has a larger parser without outperforming the inherited harness.

Paste-buffer checks adapt to TB4. TB4 exposes failures absent from TB2.1: special key sequences must be normalized for interrupts to take effect, and foreground processes must be identified so that a running command is not classified as a stalled session. The paste-buffer structure remains, with checks adjusted to these target failures.

Completion checks add target criteria. Both inherited harnesses keep completion control while revising the evidence required on TB4. Appendix F.3 follows the selected edits; Appendix F.2 gives the task evidence and the P5/P8 records. Source predicates that no longer fire are pruned rather than re-tuned. Both branches also add a reminder triggered by the remaining step budget in their first round, a target-specific mechanism that the source did not need.

Refusal handling remains inactive. Refusal handling is carried forward but has no triggering tasks on TB4. Its presence alone therefore does not establish a transfer benefit or a loss.

Parser repairs depend on the policy model. The Opus- and Gemini-evolved harnesses have near-identical terminal and process structures, which are retained after transfer. Parser repairs instead address model-specific output errors: unescaped control characters, illegal escapes, and truncated responses. Each repair needs checking against the target model's outputs. Completion control likewise retains counts of executed commands, discarded blocks, and consecutive rejections, while recalibrating checks

for simulated execution and premature completion. The target run also adds format constraints and step-by-step guidance absent from source RSI.

F.5.2. Text Classification

Source RSI builds structured retrieval. Memory writing moves from no retention (NoMemory) or unstructured storage (FewShotAll) to retrievable indices: multi-granularity BM25 over words, character n-grams, and chemical-formula n-grams, later extended with class-based candidate pools, hard-negative error indices, and class-prototype centrality. Prompt assembly bounds retrieval to the top eight examples with positive BM25 scores, strips benchmark boilerplate from queries before retrieval, and selects contrastive, class-balanced exemplars. Post-processing normalizes data-specific formats, such as SMILES strings for USPTO-50k, because open-vocabulary generation requires exact matches.

Retrieval survives with adjusted thresholds. The target suite changes the input types and label spaces. The BM25 index and retriever classes remain intact, while retrieval constants are recalibrated and query sanitization gains delimiters for the new datasets. The routing structure is retained, with rubrics and exemplars rewritten for the target. With SIM, target RSI adds smoothed class quotas and contrastive ordering; class balancing applies to the target's finite label sets rather than the source's open-vocabulary outputs.

SMILES normalization stays inactive. SMILES normalization used in source RSI remains inactive on the target suite, whose outputs no longer use that format. The cold-start run independently develops relevance-gated retrieval and bounded context.

Indexing transfers while output rules change. Three of four runs evolve BM25 indices from different starting harnesses, using features suited to the input, such as chemical n-grams for SMILES and character bigrams for Chinese legal text. Indexing and query sanitization depend on the input data and carry over across models; retrieval budgets such as top-k are re-tuned. Output rules depend more on policy behavior: the answer-first schema addresses Gemini's truncated reasoning. Target RSI also rewrites inherited rubrics more extensively on Banking77 and GoEmotions than on SciTail and TweetEval-Hate.

F.5.3. Agentic Workspace

Source RSI adds output handling and audits. Over-long tool outputs are preserved through head-and-tail truncation, with quotas raised three- and five-fold, respectively. A one-time audit gate before `final_answer` can append a tool message and return the policy to the loop. A deliverable auditor supplies the gate with file names and sizes from candidate output directories.

The termination gate survives a rubric change. Source RSI and a cold-start target run independently add a tool message before the same early-return branch, with different predicates and messages. The target gate fires only when no deliverable is registered, files were written during code execution, and the gate has not fired before; the second condition keeps purely conversational tasks from being intercepted. This insertion point depends on the ReAct loop and remains usable when the evaluation rubric changes.

Registration checks replace file-system probes. The source probe samples the file system, whereas target failures come from invocation history: files created on disk but never registered through `write_file`. Target checks therefore use tool-invocation history rather than file-system probes. The target escalates this check in three single-shot tiers: an advisory notice when file saving is detected, a mandatory instruction once 75% of the step budget passes without registration, and a conditional block at `final_answer`.

Registration checks depend on the evaluator. The registration check assumes that the evalua-

tor credits files registered through `write_file`; it would need reconsideration for an evaluator that reads output directories directly. Its substring matches over tool arguments (such as `/tmp/`, `.docx`, or `savefig`) also depend on argument formatting.

Prompt infrastructure transfers across models. The Opus source run adds editable system-prompt guidelines and a six-level severity taxonomy with counterexamples, leaving control flow and tool execution unchanged. The injection mechanism uses message structures, while the taxonomy describes the unchanged evaluator. Gemini's cold-start run takes a different direction: a notice at 80% of the step budget asks it to fill pending fields so that a deliverable exists before `final_answer`; seven later rounds add no improvement in this trace. The timing and wording of such interventions depend on the policy's behavior.

F.5.4. Engineering Design

Source RSI adds design and verification guidance. The source run adds prompt guidance on analytical calculations, self-contained deliverables, tool-failure recovery, and re-reading written files before completion. It also imports a syntax parser and a tool registry without calling them during source RSI.

Target RSI expands checks and context budgets. The tool-registry import is unused in the source harness but becomes the entry point that the Harness Transfer target run uses to enforce centralized output truncation at the registry call. Three target runs independently raise the shared context constants: output-token caps, head and tail lengths, and summarization triggers. The Full SIM target run adapts the source's one-time post-generation check into a checkpoint protocol that writes intermediate artifacts exceeding the baseline score to disk, protecting them from unannounced timeouts. The cold-start target run's retained edits are limited to prompt changes: reconstructing the objective from the evaluation script, embedding self-tuning optimization code in deliverables, and clamping outputs with a fallback to the baseline implementation.

SIM retains the source guidelines in this trace. The Harness Transfer target run inherits the source guidelines but reverts them to the seed prompt in its first round, whereas the Full SIM target run retains and extends them throughout. This pair retains similar code structures but differs in whether the inherited guidance is preserved.

Deliverable checks survive model changes. A deliverable tracker that records deliverable state independently of model output transfers without refactoring, and a check that blocks `final_answer` when a deliverable was modified but not re-executed transfers unchanged. Budget prompts emerge as wall-clock triggers in an Opus run (every 600 s, at most three times) and step-budget ratios in Gemini runs (0.8 and 0.75). Both use the gate before `final_answer`, with checks on different state variables. The recorded zero-score rates differ between Opus (0.09%) and Gemini (4.49%); intervention thresholds therefore need checking for each policy. Handlers for a missing sandbox binary remain idle where the environment is fully provisioned.

Interfaces explain much of the observed reuse. Across these cases, benchmark changes alter the evidence needed for completion, while model changes alter how the harness handles policy failures. With the benchmark fixed, a model change leaves the success criterion intact but can require different detection and guidance. Stable interfaces favor reuse; output repairs and intervention thresholds need checking against the target setting.

F.6. What Do the Cases Tell Us about SIM?

The cases suggest four roles for SIM in continuing harness optimization.

Earlier rationale informs what to retain. In engineering, the SIM run retains and extends

inherited guidance, while harness-only continuation reverts it (§F.5.4). Recording why an edit was made gives the optimizer a basis for assessing its continued relevance, beyond its presence in the inherited harness.

A transferable rationale can require a different implementation. The coding runs inherit the same completion gate, but the SIM run extends it to check target-specific requirements (§F.3). The need for completion evidence carries over; what counts as sufficient evidence changes. This illustrates how earlier reasoning can guide adaptation without treating the source implementation as fixed.

New outcomes add evidence to an existing diagnosis. P3 links a source diagnosis to further target interventions and their outcomes (§F.4). Keeping these attempts together lets later rounds inspect which remedies were tried, under which settings, and which predictions remain unresolved. The reusable experience includes the evidence behind a diagnosis, not only its name.

Experience remains available after an edit is replaced. P8 retains the selected blocking attempt and its unresolved predictions after later edits change the behavior (§F.2). The final harness alone would omit this part of the search history. SIM keeps the attempt, its diagnosis, and its unresolved outcome together as the harness changes.