GUARD: Stabilizing Recursive Self-Improvement with Regret Guidance

Minghe Zhang[1,*], **Xuetao Kang**[1,*], **Yifei Ming**[1], **Yanfei Chen**[1], **Han Yu**[1], **Yufan Zhuang**[1] and **Chen-Yu Lee**[1]
[1]Google Cloud, [*]Equal contribution

arXiv:2611.23476v1 [cs.AI] 17 Nov 2026

Autonomous AI agents pair a frozen base model with an editable *harness* of prompts, skills, tool rules, and memory. Recursive self-improvement (RSI) adapts this harness from experience without weight updates, yet self-editing without external graders often backfires: reflection and scalar scores conflate *what an edit newly broke* with *what remains unsolved*, causing local regressions, skill forgetting, and memory bloat. We introduce GUARD (*Grader-free Unsupervised Environment Design [UED] with Adaptive Rule Distillation*), a closed-loop framework for reference-free continual harness evolution. Our key insight is to decompose each task into pre-execution self-checks and guide harness updates via *directional regret*. Each round grades trajectories against these checks, pinning mastered rules while prioritizing newly broken behaviors, and selects among parallel candidate edits by critic-free group-relative fitness. A *Soft-HURT* gate then promotes an edit only when gains outweigh regressions, scoping specialized fixes to prevent over-triggering and guaranteeing non-regression on its gated buffer. At test time, one frozen harness runs in a single pass. Across benchmarks on memory, skill evolution, continual task streams, and domain reasoning, `GUARD` consistently improves performance without catastrophic forgetting or memory pollution while substantially cutting deployment token usage.

1. Introduction

Modern language model agents combine a frozen base model with a surrounding *harness* of prompts, reusable skills, tool interfaces, memory, and control logic [31, 42, 20, 45]. In practice, this harness decides whether the same model checks the right state before acting, calls tools with valid arguments, recovers from failed commands, and carries useful rules across sessions [52, 40]. Much recent progress in deployed agents comes from *harness engineering*, yet tuning harnesses by hand scales poorly: engineers can inspect only a small slice of failed runs, and fixed designs quickly go stale as tools, schemas, and user workflows change [5, 49]. *Recursive self-improvement* (RSI) aims to automate this loop by turning past experience into persistent edits that improve future runs [54, 38, 44, 26]. Because frontier models often sit behind black-box APIs with frozen weights and open-world tasks rarely provide ground-truth rewards for weight-space reinforcement learning [60, 62, 13, 46], treating the editable harness as a non-parametric policy around the frozen model offers a practical, inspectable way to improve agents after deployment [23, 36].

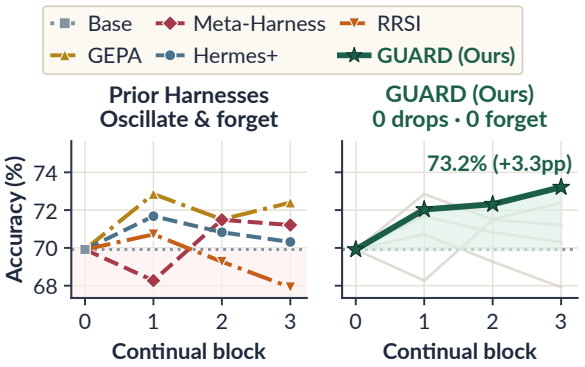


Figure 1 | **Continual harness evolution across shifting tasks.** *Left:* On Drift Gauntlet [5] (four continual task streams under single-model `Gemini 4`), prior persistent harnesses oscillate across adaptation blocks and drop below the base agent as new edits overwrite earlier skills. *Right:* `GUARD` improves monotonically at every checkpoint with zero block drops and zero post-block forgetting.

Early automated methods search over prompts or agent code offline using labeled training splits before deployment [18, 12, 55, 58, 20]. The resulting static harnesses struggle when test-time

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

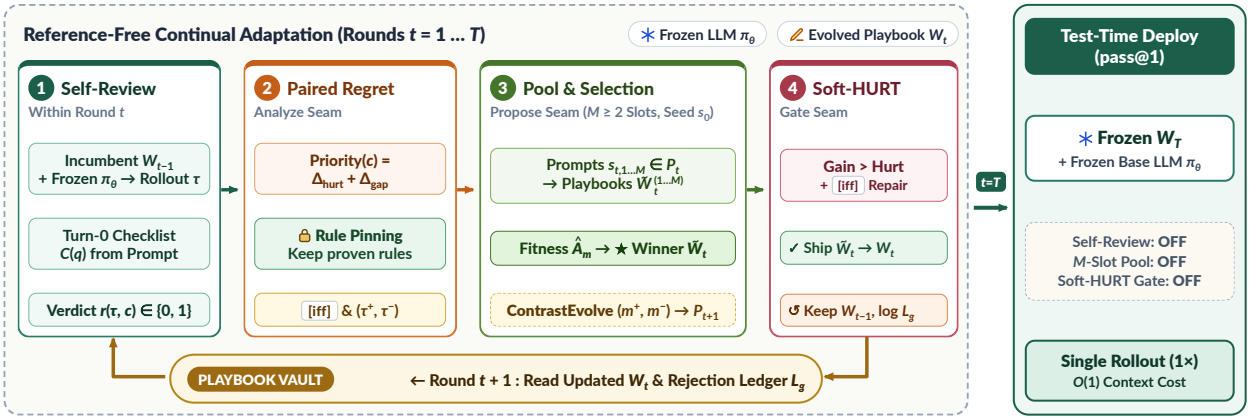


Figure 2 | **GUARD: regret-guided continual recursive self-improvement of agent harnesses** (*Grader-free UED with Adaptive Rule Distillation*). **Adaptation phase (§4):** across rounds, the frozen model executes tasks with the current harness while a reference-free self-reviewer checks simple task requirements on each trajectory; `Analyze` separates newly broken behaviors from remaining gaps and pins working rules; `Propose` compares passing and failing runs to write targeted rules that state when they apply; and `Gate` (*Soft-HURT*) keeps an update only when gains outweigh regressions while logging rejected edits so bad changes are not repeated. **Test-time evaluation:** candidate pooling, self-review, and gating are turned off, deploying the single frozen harness for standard single-run evaluation (`pass@1` on all primary benchmarks, matching the base agent's rollout budget).

tasks shift. Newer approaches regularize offline search for out-of-distribution generalization [45], probe per-task harnesses online while saving reusable lessons [49], or append reflections and skills across episodes [59, 30, 50, 1]. Yet continually updating a single persistent harness across rounds without reference answers or external graders often degrades performance over time [5]. Because every edit changes subsequent rollouts, free-form reflection and scalar task scores mix two distinct failure signals: behaviors the latest edit *newly broke* and capability gaps the agent *still has not solved*. Mixing these signals across rounds causes three failure modes: *local overfitting* (*over-triggering*), where a broad rule written for recent failures breaks tasks that already worked; *catastrophic forgetting*, where shifts in workflow, topic, or domain overwrite earlier skills [5]; and *memory pollution*, where unverified additions accumulate stale or conflicting rules that bloat the prompt. Over multiple rounds, self-editing can leave an agent worse than its starting point (Figure 1).

> ***Central Question.*** *How can an autonomous agent continually improve its persistent harness in a reference-free setting without interfering with or catastrophically forgetting previously learned skills?*

We address this challenge by casting harness adaptation around a frozen model as gradient-free policy optimization in text space, guided by two ideas from reinforcement learning: curriculum regret from unsupervised environment design (UED) [7, 14, 15, 35] and critic-free group-relative scoring from Group Relative Policy Optimization (GRPO) [37, 3]. We never update the model weights or train a value critic: the harness is the only part of the policy that changes (§2). We implement this approach in **GUARD** (*Grader-free UED with Adaptive Rule Distillation*; Figure 2, §4). Before executing a task, `GUARD` compiles the task prompt and tool schema into a fixed checklist of atomic binary checks (§4.1) and grades completed runs with a two-layer (sandbox execution plus semantic trace) self-reviewer, because checking one requirement on a finished trace is simpler than satisfying every constraint during generation. Across rounds, `Analyze` (§4.2) pins

criteria that stay satisfied and ranks newly broken criteria ahead of unsolved gaps; `Propose` (§4.3) samples a parallel group of candidate edits, scores them against each other on the same tasks, and contrasts the best and worst candidates to rewrite its proposal prompts; and `Gate` (§4.4) promotes the winner only when its gains exceed induced regressions (*Soft-HURT*), narrowing rules that cause isolated drops with `[iff]` preconditions and logging rejected edits. At test time, candidate pooling, self-review, and gating are turned off, and the agent runs a single pass (`pass@1`) with its frozen harness.

Three choices separate GUARD from prior harness learners. Reflexion and GEPA update from free-form reflection or scalar and Pareto scores, so they cannot tell a new regression from an old gap; GUARD ranks checklist criteria by paired directional regret. META-HARNESS searches agent code against raw trace archives, and RRSI rewrites one candidate per round behind a score-band gate that leaves per-task drops unbounded; GUARD scores a parallel group of candidates on the same tasks and promotes the winner only when its gains exceed induced regressions, narrowing rules with `[iff]` preconditions instead of deleting them. We evaluate GUARD across four benchmark settings (§5–§6, Appendix C): PAST-BENCH [50], EVOHARNESSBENCH [17], DRIFT GAUNTLET [5], and professional legal and financial reasoning suites [2, 16, 27]. Our contributions are threefold:

- **Trace-grounded checklist self-review and directional regret (`Analyze`, §2–§4.2).** We formulate reference-free harness adaptation over a frozen model using paired directional regret on pre-execution task checklists. We prove that dense checklist grading removes the $1/\epsilon$ passive-discovery penalty on rare task invariants (Theorem 2), that paired priority ranks every fresh regression above every non-regressed criterion in the single-rollout binary case and pins exactly the mastered criteria (Proposition 1), and that gating keeps self-review scores non-decreasing on each round's gated task buffer (Theorem 1).
- **GRPO-inspired candidate pooling and contrastive prompt evolution without weight updates (`Propose`, §4.3).** Inspired by group-relative reinforcement learning [37, 3], GUARD samples a parallel group of candidate harnesses from one domain-agnostic seed prompt, evaluates them on the same task batch, and ranks them by critic-free group-standardized fitness (a selection step, not a gradient step) while contrasting the best and worst group candidates to rewrite proposal prompts across rounds. In a controlled single-model ablation at equal budget, parallel two-candidate pooling beats sequential single-candidate search and cuts rejected gate rounds by 64%, while deploying a single harness in one pass at test time.
- **Soft-HURT promotion gating and continual generalization (`Gate`, §4.4, §6, Appendix C).** Pairing Soft-HURT promotion with in-round `[iff <condition>]` precondition repair yields 19/19 non-regressive execution-gated promotions across EVOHARNESSBENCH, DRIFT GAUNTLET, and the professional suites. GUARD is the only evaluated method with no negative backward or forward transfer on hosted EVOHARNESSBENCH (FWT is zero because the gate abstained at later stages) and cuts tokens per task by 54.5% versus the base agent, incurs no stream-mean drop across all 12 block transitions and zero post-block forgetting on all four DRIFT GAUNTLET streams, and ties the memory-pollution floor on PAST-BENCH.

2. Background and Problem Formulation

Agents, harnesses, and task streams. We model an autonomous agent $A = (\pi_\theta, W_{t,g})$ as a frozen language model policy π_θ paired with an inspectable *harness state* $W_{t,g}$ at adaptation round $t \in \{0, 1, \ldots, T\}$ ($W_{0,g}$ empty; subscript g omitted when clear) [20, 45]. Tasks arrive from groups $g \in \mathcal{G}$ indexing roles, task families, or, when a continual stream hides group labels and drift boundaries [5], the pre-execution (*Turn-0*, before the solver's first action) observable tool and schema signature of query $q \in Q_g$. We call $W_{t,g}$ a *non-parametric* policy component because it is editable text rather than model weights: condition-

ing on it steers rollouts $\tau \sim \pi_\theta(\cdot \mid q, W_{t,g})$, while θ stays fixed during both adaptation and testing. In reinforcement-learning terms, an action is a harness edit, its reward is the gated self-review fitness R_m (§4.3), and only W is optimized. While this formulation accommodates general text or code harnesses [26], we instantiate $W_{t,g}$ as structured Markdown files, termed a *harness playbook*, organizing workflow instructions, entity-grouped domain rules, reusable skills (`SKILL.md`), persistent memories (`MEMORY.md`), and optional pre-submit checks, with specialized exceptions scoped by explicit `[iff <condition>]` preconditions.

Reference-free continual evolution and non-regressive RSI. Let $V^\star(q,\tau) \in [0,100]$ denote the unobserved task evaluator, unavailable throughout adaptation. Observing only task prompts, tool interfaces, execution trajectories, and runtime states in interaction history $\mathcal{H}_{t,g}$, a reference-free harness *evolver* $\mathcal{E}$ recursively updates the incumbent across adaptation rounds:

$$W_{t,g} \;=\; \mathcal{E}\big(W_{t-1,g},\, \mathcal{H}_{t,g}\big), \qquad t = 1, \ldots, T \le T_{\max}. \quad (1)$$

At test time, evolution freezes and each unseen query $q \in Q_g$ executes under the single promoted harness $W_{T,g}$ to maximize $\mu_g(W_{T,g}) = \mathbb{E}_{q\in Q_g,\, \tau\sim\pi_\theta(\cdot|q,W_{T,g})}[V^\star(q,\tau)]$. Because $\mathcal{E}$ updates $W_{t,g}$ adaptively from finite online rollouts without $V^\star$, continual generalization under Eq. (1) requires improving μ_g without local overfitting on solved tasks or catastrophic forgetting of mastered behaviors as t and task distributions shift [45, 5]: scalar round scores mix edit-induced regressions with unresolved capability gaps and compound negative transfer across rounds.

What reference-free does and does not mean. Throughout this paper, *reference-free* (and *grader-free*) is an operational constraint on the adaptation loop $\mathcal{E}$: neither the solver, the self-reviewer, nor the distiller ever accesses held-out reference answers, test cases, hidden verifiers $V^\star$, oracle skills, or official grading rubrics. It does *not* mean learning from zero supervision. Specifically, adaptation draws on four observable sources of structure: (i) benchmark-authored training task prompt distributions, (ii) public tool declarations and environment schemas, (iii) the agent's own interaction trajectories and observable sandbox states (with no human expert demonstrations), and (iv) structured criterion-level judgments from the language-model self-reviewer (§4.1).

3. Related Work

Table B.1 (Appendix B.1) compares GUARD with representative self-improvement methods and environments across the components they optimize, their closed-loop mechanisms, and their adaptation and context scalability.

Verbal reflection, memories, and playbooks. Reflexion [38], ExpeL [59], ReasoningBank [30], and AWM [43] store verbal reflections or workflows; MemGPT [31], Voyager [42], and WikiSkill [40] maintain memories, skill libraries, and append-only wikis; Dynamic Cheatsheet [39] and ACE [56] incrementally update playbooks; and Context Language Models (CLM) [36] update model weights to compact context via Group Relative Policy Optimization (GRPO) [37], whereas Training-Free GRPO [3] summarizes group rollouts into an experience library with frozen weights. GUARD differs in what forms the group: it scores M candidate harnesses on the same task batch rather than rollouts of one fixed context, contrasts the best and worst candidates to rewrite proposal prompts, and accepts the winner only through a regression-aware gate. AutoGuide [9], AutoManual [4], and CLIN [24] distill conditional guidelines, rule manuals, or causal memories from experience; GUARD's `[iff]` rules are likewise conditional, but enter the harness only through regression-aware ranking and gating. Because these updates neither isolate edit-induced regressions nor gate per-task drops, they remain vulnerable to the plasticity–pollution and plasticity–stability trade-offs measured by PASTBench [50], EvoHarnessBench [17], and Drift Gauntlet [5]. Accepting an update only when prior behavior does not degrade echoes regression-free model updates [47] and gradient-constrained continual learning [21], which rely on labels or gradients that a reference-free harness loop does not have.

Unsupervised environment design and zero–data RL. UED [7, 14, 15] prioritizes high-regret POMDP levels via antagonist regret (PAIRED [7]), maximal Monte Carlo (MaxMC) regret [14, 15, 32], or learnability $p(1-p)$ [35], whereas we transfer regret from level curation to directional criterion ranking over pre-execution task checklists [61] without an antagonist. Zero-data RL self-play (Absolute Zero [60], Self-Challenging Agents [62], R-Zero [13], and Agent0 [46]) eliminates human labels via code verifiers or majority voting, yet requires weight gradients (precluding black-box APIs), falters on stateful tool mutations, and updates opaque model weights instead of inspectable conditional rules.

Prompt, workflow, and harness optimization. Prompt optimizers such as DSPy [18], TextGrad [53], TextReg [10], OPRO [51], WizardLM [48], Promptbreeder [8], EvoPrompt [11], and GEPA [1] search instructions via scalar objectives, textual gradients, or Pareto archives; TextReg [10] filters low-recurrence patches under labeled single-turn supervision, which can suppress valid edge-case exceptions. Workflow/code optimizers ADAS [12], AFlow [55], and Darwin Gödel Machine (DGM) [58] mutate code via archives [19, 28] or MCTS against validation ground truth, while AutoHarness [22], SelfHarness [57], Meta-Harness [20], RRSI [45], and Harness^2 [49] evolve harnesses around frozen weights (alongside MLGym [29], PostTrainBench [34], RSIBench-Data [25], and RSIGym [26]). Among open-world designs, Harness^2 probes per-task harnesses online from W_0 at extra test-time rollout cost, RRSI's score-band gate leaves per-task drops unbounded, and persistent harnesses (ReasoningBank, GEPA, Meta-Harness) suffer severe forgetting under hidden drift [5]. `GUARD` instead evolves a single persistent harness in a reference-free loop and deploys it at $O(1)$ context in T and single-pass (`pass@1`) test cost (Appendix B.5).

4. `GUARD`: Regret-Guided Harness Evolution

During adaptation, `GUARD` instantiates the reference-free evolver $\mathcal{E}$ in Eq. (1) as a four-stage closed loop (Figure 2; Algorithm 1; constants in Appendix B): pre-execution checklist derivation and two-layer self-review (§4.1), paired regret ranking (`Analyze`, §4.2), contrastive candidate pooling (`Propose`, §4.3), and Soft-HURT promotion (`Gate`, §4.4), with formal guarantees in §4.5. At test time, this adaptation loop is disabled and each unseen task runs a single rollout under the frozen harness $W_{T,g}$.

4.1. Reference-Free Self-Review via Task Checklists

Pre-execution task checklist $C(q)$. Extending IFEval's verifiable constraint decomposition [61] and checklist-based evaluation [33, 6, 41] to multi-turn tool workflows, the self-reviewer reads *only* the task prompt, visible tool declarations, and schema once at Turn 0 (before the solver's first action) to derive and cache a fixed checklist $C(q)$ of atomic binary predicates (*criteria* $c \in C(q)$; Appendix B) across two classes: (i) *execution-verifiable* predicates (clean sandbox execution, valid entity IDs/columns, idempotent mutations) and (ii) *semantic* predicates (honoring every prompt constraint, recipient, format, and safety rule without applying superseded rules).

Two-layer grading and verification asymmetry. Layer 1 checks execution predicates deterministically on an ephemeral sandbox copy; Layer 2 grades semantic predicates individually from the prompt, output, and tool/state trace, returning $r(\tau, c) \in \{0, 1\}$ and a failure rationale. Over graded predicates $C^{\mathrm{gr}}(\tau) \subseteq C(q)$ ($|C^{\mathrm{gr}}(\tau)| \geq 1$; unconfigured checks excluded; an unparseable verdict counts as failed or falls back to deterministic trace checks), trajectory τ receives criterion

score $S(\tau)$ and strict score $S^{\text{strict}}(\tau)$:

$$S(\tau) = 100 \cdot \frac{1}{|C^{\text{gr}}(\tau)|} \sum_{c \in C^{\text{gr}}(\tau)} r(\tau, c),$$
$$S^{\text{strict}}(\tau) = \begin{cases} 0 & \text{on a blocking sandbox error,} \\ 100 & \text{if all graded criteria pass,} \\ S(\tau)/2 & \text{otherwise.} \end{cases} \tag{2}$$

Pre-execution decomposition uses *generation–verification asymmetry*: long-horizon generation couples N constraints and suffers "seen-but-not-carried" omissions, while checking one atomic predicate after execution requires localized $O(1)$ attention grounded by Layer 1 sandbox states. Cross-task $\Delta_{\text{hurt},t}$ and Soft-HURT gating then prevent single-instance overfitting [23].

4.2. Paired Regret Ranking (`Analyze` Stage)

Definition 1 (Criterion obedience and paired regret)**.** *For trajectories $\mathcal{T}_c(W)$ grading criterion c under W ($\mathcal{T}_c(W_{t-1}) = \emptyset$ skipped), criterion obedience is*

$$\text{obey}(c; W) = \frac{1}{|\mathcal{T}_c(W)|} \sum_{\tau \in \mathcal{T}_c(W)} r(\tau, c) \in [0, 1]. \tag{3}$$

Entering round $t \geq 1$ with incumbent W_{t-1}, comparing the two most recent incumbents W_{t-2} and W_{t-1} on the same tasks gives:

$$\Delta_{\text{hurt},t}(c) = \max\big(0,\ \text{obey}(c; W_{t-2}) - \text{obey}(c; W_{t-1})\big),$$
$$\Delta_{\text{gap},t}(c) = 1 - \text{obey}(c; W_{t-1}),$$
$$\text{Priority}_t(c) = \Delta_{\text{hurt},t}(c) + \Delta_{\text{gap},t}(c) \in [0, 2], \tag{4}$$

where $\Delta_{\text{hurt},t}(c) = 0$ if $t = 1$ or c is not graded under both compared playbooks.

Both terms are non-negative and parameter-free: a regressed criterion receives its incumbent gap plus its drop, outranking equally obeyed non-regressed criteria (Proposition 1); with one binary verdict per criterion after a promotion, $\text{Priority}_t(c) \in \{2, 1, 0\}$ separates regressions ($1 \to 0$), unsolved gaps ($0 \to 0$), and mastered checks ($1 \to 1$ or $0 \to 1$). In regret terms, $\Delta_{\text{gap},t}$ is the incumbent's shortfall against an oracle harness that satisfies every criterion, and $\Delta_{\text{hurt},t}$ is the negative transfer introduced by the latest promotion.

Rejected rounds, diagnostic lists, and trajectory contrasts. After a rejection ($W_{t-1} = W_{t-2}$), execution scoring instead takes $\Delta_{\text{hurt},t}$ from W_{t-1} to the rejected winner $\widetilde{W}_{t-1}$ (or its repair $\widetilde{W}^{\text{rep}}_{t-1}$), so the next proposal sees which criteria the refused edit broke, while $\Delta_{\text{gap},t}$ and pinning stay anchored to the incumbent; critic scoring always compares consecutive incumbents (§5). The signs of the two terms split criteria in the regret bundle $\mathcal{B}_{t,g}$ into three diagnostic lists, each sorted by descending $\text{Priority}_t(c)$ (ties broken by $\Delta_{\text{hurt},t}(c)$): (1) *Regressed* ($\Delta_{\text{hurt},t} > 0$), asking the proposer to wrap offending rules in `[iff <trigger>]` preconditions instead of discarding specialized fixes [10]; (2) *Unsolved Gap* ($\Delta_{\text{hurt},t} = 0, \Delta_{\text{gap},t} > 0$), targeting persistent gaps; and (3) *Mastered* ($\text{Priority}_t = 0$, i.e., $\text{obey}(c; W_{t-1}) = 1 \wedge \Delta_{\text{hurt},t} = 0$), pinning obeyed rules against overwriting. These lists serve as diagnostic annotations, not hand-coded edit operators: the proposer receives them with the incumbent, trajectory contrasts, and the ledger window, and writes candidate playbooks end-to-end. Finally, on tasks with differing scored rollouts, $\mathcal{B}_{t,g}$ pairs top and bottom trajectories (τ^+, τ^-) (ordered by S^{strict}, ties by S) whenever

$$\Delta^{\text{strict}}(\tau^+, \tau^-) = S^{\text{strict}}(\tau^+) - S^{\text{strict}}(\tau^-) > 0$$
$$\vee\ \Delta^{\text{cont}}(\tau^+, \tau^-) = S(\tau^+) - S(\tau^-) \geq \delta_{\text{con}}, \tag{5}$$

($\delta_{\text{con}} = 10$). This condition captures partial checklist progress even when blocking errors or critical caps flatten S^{strict}.

4.3. Group-Relative Candidate Pooling and Contrastive Evolution (`Propose` Stage)

In reinforcement learning, Group Relative Policy Optimization (GRPO) [37] estimates policy advantage without a value network by sampling a group of rollouts from a shared state and standardizing their rewards within the group. Following the training-free group-relative perspective [3], the `Propose` stage applies this principle to non-parametric harness optimization without updating model weights: it samples $M \geq 2$ candidate harnesses in parallel from $(W_{t-1,g}, \mathcal{B}_{t,g})$, selects the winner by within-group standardized fitness ($\arg\max_m \hat{A}_m$), and contrasts the best and worst group candidates (ContrastEvolve) to update the strategy-prompt pool $\mathcal{P}_{t,g} = \{s_{t,1}, \ldots, s_{t,M}\}$ across rounds. Unlike GRPO, `GUARD` takes no gradient step and uses no clipping or KL penalty; $\hat{A}_m$ only ranks candidates and picks the contrast pair (m^+, m^-), and the Soft-HURT gate (§4.4) plays the role of a trust region.

What the group buys is credit isolation without a learned critic: all M candidates are scored on the same $(W_{t-1,g}, \mathcal{B}_{t,g}, Q_g)$, so batch-level task difficulty cancels out.

Uniform initialization and scoring. Round 1 initializes all M parallel candidates with the same domain-agnostic prompt s_0 (`SHARED_SEED_PROMPT`, Appendix B.4):

$$\mathcal{P}_{1,g} = \{s_{1,1}, \ldots, s_{1,M}\}, \qquad s_{1,1} = \cdots = s_{1,M} = s_0, \tag{6}$$

drawing diversity without operator priors from sampling $\widetilde{W}_{1,g}^{(m)} = \mathrm{Distill}(s_0, \mathcal{B}_{1,g}, W_{0,g})$. Over $N_{\mathrm{rep}} \geq 1$ replicates $\{\tau_{q,r}\}_{r=1}^{N_{\mathrm{rep}}}$ and noise floor $\delta_{\mathrm{noise}} \geq 0$, task scores and helped/hurt sets are

$$\begin{aligned} S_q(W) &= \frac{1}{N_{\mathrm{rep}}} \sum_{r=1}^{N_{\mathrm{rep}}} S(\tau_{q,r}), \\ S_q^{\mathrm{strict}}(W) &= \frac{1}{N_{\mathrm{rep}}} \sum_{r=1}^{N_{\mathrm{rep}}} S^{\mathrm{strict}}(\tau_{q,r}), \\ Q^{\pm}(\widetilde{W}, W) &= \left\{ q \in Q_g : \pm\big(S_q(\widetilde{W}) - S_q(W)\big) > \delta_{\mathrm{noise}} \right\}, \end{aligned} \tag{7}$$

with means $\bar{S}(W), \bar{S}^{\mathrm{strict}}(W)$ supplied by *execution scoring* (M group rollouts/round) or *critic scoring* (one joint pass over accumulated trajectories on stateful suites). With $\mathrm{Perf}_m = \bar{S}^{\mathrm{strict}}(\widetilde{W}_t^{(m)}) + \frac{1}{4}\bar{S}(\widetilde{W}_t^{(m)})$ and $Q_m^{\pm} = Q^{\pm}(\widetilde{W}_t^{(m)}, W_{t-1})$, scalar and standardized fitness are

$$\begin{aligned} R_m &= \mathrm{Perf}_m - \beta\,|Q_m^-| + \gamma\,|Q_m^+|, \qquad \beta > \gamma \geq 0, \\ \hat{A}_m &= \big(R_m - \mu_R\big)/\sigma_R, \\ \sigma_R &= \sqrt{\tfrac{1}{M}\textstyle\sum_{m=1}^{M}(R_m - \mu_R)^2} + 10^{-6}, \end{aligned} \tag{8}$$

($\mu_R = \frac{1}{M}\sum_m R_m$), where $\beta > \gamma$ penalizes hurts over helps and $m^\star = \arg\max_m \hat{A}_m = \arg\max_m R_m$ selects $\widetilde{W}_t = \widetilde{W}_t^{(m^\star)}$ (with $M = 2$, $\hat{A}_m \approx \pm 1$ whenever $R_1 \neq R_2$, so selection keeps the higher-fitness candidate).

Contrastive prompt evolution, deployment, and complexity. Selecting distinct best and worst candidates in the group,

$$m^+ = \arg\max_{m \in \{1,\ldots,M\}} R_m, \qquad m^- = \arg\min_{m \neq m^+} R_m \tag{9}$$

(ties broken by $\bar{S}$, $|Q_m^-|$, candidate index; mirroring MaxMC spread [15]), ContrastEvolve scrubs task IDs to `<TASK_ID>`, contrasts $(s_{t,m^\pm}, \widetilde{W}_{t,g}^{(m^\pm)}, R_{m^\pm}, |Q_{m^\pm}^{\pm}|)$, and outputs `learnings`, an exploitation prompt s_{t+1}^+, and a precondition-narrowing prompt s_{t+1}^-, updating

$$\mathcal{P}_{t+1,g} = \{s_{t+1}^+\}^{\lfloor M/2 \rfloor} \cup \{s_{t+1}^-\}^{M - \lfloor M/2 \rfloor}. \tag{10}$$

At test time, $\mathcal{P}_{t,g}$, the self-reviewer, and $\mathcal{L}_g$ are discarded: deployment runs a single (1×) rollout per task (`pass@1` on all primary benchmarks) under frozen $W_{T,g}$. During adaptation, parallel M-candidate pooling isolates credit on identical $(W_{t-1,g}, \mathcal{B}_{t,g}, Q_g)$, bootstraps Round-1 contrasts, and reduces single-sample gate rejections (§6.4, Appendix C.4). Execution scoring costs $O(TMQ)$ rollouts with $O(TQ)$ parallel rollout depth (or $O(TQ)$ online rollouts under critic scoring), while bounded playbooks B_W, ledger window $L_{\max}$, and capped bundles $B_{\mathcal{B}}$ keep proposer and test context $O(1)$ in T (Table B.1, Appendix B.5).

4.4. Soft-HURT Promotion Gate and Rejection Ledger (`Gate` Stage)

Soft-HURT promotion gate. Using per-task scores S_q and hurt set $Q^-(\widetilde{W}, W)$ from (7), the group-averaged regression penalty beyond δ_{noise} is

$$\mathrm{Hurt}(\widetilde{W}, W) = \frac{1}{|Q_g|} \sum_{q \in Q^-(\widetilde{W}, W)} \big(S_q(W) - S_q(\widetilde{W})\big), \tag{11}$$

and the Soft-HURT rule promotes a candidate over W_{t-1} only when net gain exceeds induced hurt:

$$\begin{aligned} &\mathrm{Promote}(\widetilde{W}, W_{t-1}) \iff \\ &\quad \bar{S}(\widetilde{W}) - \bar{S}(W_{t-1}) > \mathrm{Hurt}(\widetilde{W}, W_{t-1}). \end{aligned} \tag{12}$$

Parameterized by single-task noise floor $\delta_{\mathrm{noise}} \geq 0$, hurt-task budget $h_{\max} \geq 0$, and trade-off margin $\mu_{\min} \geq 0$, Soft-HURT admits $\widetilde{W}_t$ when $|Q^-(\widetilde{W}_t, W_{t-1})| \leq h_{\max}$, $|Q^+(\widetilde{W}_t, W_{t-1})| \geq \mathbb{I}[|Q^-(\widetilde{W}_t, W_{t-1})| > 0]$, and

$$\begin{aligned} &\bar{S}(\widetilde{W}_t) - \bar{S}(W_{t-1}) > \\ &\quad \max\big(\mathrm{Hurt}(\widetilde{W}_t, W_{t-1}),\ \mu_{\min}\,\mathbb{I}[|Q^-(\widetilde{W}_t, W_{t-1})| > 0]\big). \end{aligned} \tag{13}$$

When $\widetilde{W}_t$ fails solely due to isolated regressions ($1 \leq |Q^-(\widetilde{W}_t, W_{t-1})| \leq h_{\mathrm{rep}}$), `GUARD` triggers at most one in-round `[iff]` scope-guard repair $\widetilde{W}_t^{\mathrm{rep}}$ and promotes $\widetilde{W}_t^{\mathrm{rep}}$ if its re-scored trajectories satisfy the same Soft-HURT rule ($\mu_{\min} = 0$

Algorithm 1 GUARD: reference-free regret-guided harness RSI during adaptation (at test time, each unseen task executes a single pass under frozen $W_{T,g}$). Distill$(s, \mathcal{B}, W)$ compiles a candidate harness playbook under strategy prompt s from regret bundle $\mathcal{B}$ and incumbent W, and ContrastEvolve updates the best and worst strategy prompts via Eq. (10).

Require: groups $g \in \mathcal{G}$ with task sets Q_g; pool size $M \geq 2$; uniform seed prompt s_0 (`SHARED_SEED_PROMPT`); round budget $T_{\max}$
Output: final incumbent $W_{T,g}$ of every group g, where $T \leq T_{\max}$ is the last round run
Derive and cache pre-execution checklists $C(q)$ from public prompts and tool schemas (§4.1)
for each group $g \in \mathcal{G}$ **do**
  $W_{0,g} \leftarrow$ empty playbook; $\mathcal{L}_g \leftarrow$ empty ledger; $\mathcal{P}_{1,g} \leftarrow \{s_0, \ldots, s_0\}$ (M parallel candidates, Eq. (6))
  Solve and self-review Q_g under $W_{0,g}$ on fixed-batch suites ▷ measured baseline
end for
for $t = 1, \ldots, T_{\max}$ **do**
  for each group $g \in \mathcal{G}$ **do**
    On stateful critic-scored suites, execute round-t online episode under $W_{t-1,g}$ and self-review
    Compute $\Delta_{\text{hurt},t}(c)$, $\Delta_{\text{gap},t}(c)$, $\text{Priority}_t(c)$ from $W_{t-2,g}, W_{t-1,g}$ (§4.2) ▷ Eq. (4)
    Partition c into Regressed ($\Delta_{\text{hurt},t} > 0$) / Unsolved Gap ($\Delta_{\text{hurt},t} = 0 \wedge \Delta_{\text{gap},t} > 0$) / Mastered ($\text{Priority}_t = 0$, pinned)
    $\mathcal{B}_{t,g} \leftarrow$ incumbent $W_{t-1,g}$, routed partitions, (τ^+, τ^-) contrasts, active windowed $\mathcal{L}_g$
    for each candidate $m \in \{1, \ldots, M\}$ with strategy prompt $s_{t,m} \in \mathcal{P}_{t,g}$ **do**
      $\widetilde{W}_{t,g}^{(m)} \leftarrow \text{Distill}(s_{t,m}, \mathcal{B}_{t,g}, W_{t-1,g})$
      Score $\widetilde{W}_{t,g}^{(m)}$ on Q_g ▷ per-candidate execution rollouts or joint critic pass
    end for
    $m^\star \leftarrow \arg\max_m \hat{A}_m$ ($= \arg\max_m R_m$) ▷ Eq. (8)
    $\widetilde{W}_{t,g} \leftarrow \widetilde{W}_{t,g}^{(m^\star)}$ ▷ group-standardized fitness
    if Soft-HURT gate (§4.4) accepts $\widetilde{W}_{t,g}$ **then**
      $W_{t,g} \leftarrow \widetilde{W}_{t,g}$
    else if $1 \leq |Q_{m^\star}^-| \leq h_{\text{rep}}$ and repaired $\widetilde{W}_{t,g}^{\text{rep}}$ passes Soft-HURT **then** ▷ `[iff]` repair
      $W_{t,g} \leftarrow \widetilde{W}_{t,g}^{\text{rep}}$
    else
      $W_{t,g} \leftarrow W_{t-1,g}$; record deduplicated active rejection in $\mathcal{L}_g$
    end if
    Select $m^+ \leftarrow \arg\max_m R_m$ ▷ Eq. (9)
    Select $m^- \leftarrow \arg\min_{m \neq m^+} R_m$
    $(\texttt{learnings}, s_{t+1}^+, s_{t+1}^-) \leftarrow$ ContrastEvolve$\big((s_{t,m^\pm}, \widetilde{W}_{t,g}^{(m^\pm)}, R_{m^\pm}, |Q_{m^\pm}^\pm|)\big)$ ▷ task IDs scrubbed
    $\mathcal{P}_{t+1,g} \leftarrow \{s_{t+1}^+\}^{\lfloor M/2 \rfloor} \cup \{s_{t+1}^-\}^{M - \lfloor M/2 \rfloor}$ ▷ Eq. (10)
  end for
  $T \leftarrow t$
  if optional early-stop criterion fires (disabled in all reported runs) **then break**
  end if
end for
return $W_{T,g}$ for every group $g \in \mathcal{G}$

on repaired candidates; suite parameters in Appendix B). Any other candidate is rejected, so every promoted update obeys (12).

Rejection ledger and stopping. Rejected gate decisions log deduplicated active entries to $\mathcal{L}_g$ (an on-disk JSONL log; candidate-level runners store gate metrics and scrubbed rule excerpts with readback cap $L_{\max} \in \{12, 40\}$, while the professional-suite runner stores active criterion-removal excerpts; Appendix B), rendered into $\mathcal{B}_{t+1,g}$ at $O(1)$ prompt cost in T so refused edits are not repeated verbatim while `[iff]` refinements remain eligible. Rejected candidates are never re-executed ($\Delta_{\text{hurt},t+1}$ reuses their gate scores), and optional early stopping is disabled in all reported runs.

4.5. Theoretical Guarantees

These results characterize the self-review objective on gated task buffers; held-out accuracy is evaluated empirically in §6. Proofs of Theorem 1, Theorem 2, and Proposition 1 appear in Appendix A. For Theorem 1, fix group $g \in \mathcal{G}$ with fixed Q_g: conditional on prior history, replicate errors $\{S(\tau_{q,r}) - \mu_q(W)\}_{q \in Q_g, r \leq N_{\text{rep}}}$ ($N_{\text{rep}} \geq 1$ under execution rollouts) are independent across q, r and σ-sub-Gaussian (mean $\bar{S}(W)$, expectation $\mu(W) = \frac{1}{|Q_g|} \sum_{q \in Q_g} \mu_q(W)$). Over $T_{\max}$ rounds, at most $n_{\text{eval}} = 1 + T_{\max}(M + \mathbb{I}_{\text{rep}})$ candidate harnesses are evaluated per group ($\mathbb{I}_{\text{rep}} \in \{0, 1\}$ for in-round `[iff]` repair) against recorded incumbent scores without re-running incumbents; let $\Delta_t = \bar{S}(W_t) - \bar{S}(W_{t-1})$.

Theorem 1 (Non-Regressive Gating and Concentration)**.** *Under the gate of* (12) *and the replicate setup above:*

(i) ***No downward drift:*** $\Delta_t > 0$ *if round t promotes and $\Delta_t = 0$ otherwise, so*

$$\bar{S}(W_T) = \bar{S}(W_0) + \sum_{t=1}^{T} \Delta_t \;\geq\; \bar{S}(W_0). \quad (14)$$

(ii) ***Regressions are paid for:*** *on every promoted round, with $Q_t^- = Q^-(W_t, W_{t-1})$ the tasks dropping by more than δ_{noise},*

$$\sum_{q \in Q_t^-} \big(S_q(W_{t-1}) - S_q(W_t)\big) \;<\; |Q_g|\, \Delta_t, \quad (15)$$

and the left-hand side is zero whenever $Q_t^- = \emptyset$ (in particular when $h_{\max} = 0$).

(iii) ***Concentration:*** *for any $\alpha \in (0,1)$, with probability at least $1-\alpha$, simultaneously for every group $g \in \mathcal{G}$ (with* log *the natural logarithm),*

$$\mu(W_T) \geq \mu(W_0) + \sum_{t=1}^{T} \Delta_t - 2\sigma\sqrt{\frac{2\log(2n_{\text{eval}}|\mathcal{G}|/\alpha)}{|Q_g|N_{\text{rep}}}}. \tag{16}$$

(iv) ***Factored speedup:*** *if group g promotes in each round independently with probability a_g, per-group gating advances a fraction $\frac{1}{|\mathcal{G}|}\sum_g a_g$ of groups per round vs. $\prod_g a_g$ under a global gate ($a^{1-|\mathcal{G}|}\times$ for $a_g = a$).*

The union bound in (iii) covers all n_{eval} evaluated harnesses per group (paying only $\log(M + \mathbb{I}_{\text{rep}})$ inside n_{eval} for within-group selection and repair) and applies to execution rollouts ($N_{\text{rep}} \geq 1$; main-text suites use $N_{\text{rep}} = 1$ and Appendix C uses $N_{\text{rep}} = 2$) instead of single critic calls. Here μ is the expected self-review score on Q_g, and Eq. (16) serves as an asymptotic sample-complexity scaling law ($|Q_g|N_{\text{rep}} = \Omega(\sigma^2 \log(n_{\text{eval}}|\mathcal{G}|/\alpha)/\Delta^2)$ samples keep the selection optimism of recorded incumbent scores, at most twice the deviation term in (16), below a target gain Δ) rather than a non-vacuous finite-sample certificate at our benchmark task counts. When incumbents are re-scored on growing episode histories (PastBench) or online task blocks shift across stages (EvoHarnessBench, Drift Gauntlet), (i)–(ii) hold within each round's gated task buffer rather than telescoping unconditionally across distribution shifts, which motivates pre-execution interface-signature routing and replay-buffer anchor tasks (§6.3).

Theorem 2 (Targeted Repair vs. Passive Discovery)**.** *Consider J invariants exercised by Q_g and violated by W_0 (an invariant is a task-independent rule that the playbook must encode). Compare two observability regimes: a* passive end-task *learner that observes an invariant only when a coarse task failure surfaces it with probability $p_j \leq \epsilon$, versus* dense pre-execution checklist *grading where every violated invariant is captured by a derived criterion evaluated each round. Let T_{passive} be the rounds until passive surfacing observes all J invariants, T_{target} the rounds until targeted repair resolves all J invariants, and $H_J = \sum_{\ell=1}^{J} 1/\ell$ ($H_0 = 0$).*

(i) ***Passive lower bound:*** *if, independently across invariants and rounds, invariant j surfaces in each round with probability $p_j \leq \epsilon$ for $\epsilon \in (0,1)$, then*

$$\mathbb{E}[T_{\text{passive}}] \geq \frac{H_J}{-\log(1-\epsilon)} \geq \frac{1-\epsilon}{\epsilon} H_J = \Omega\left(\frac{\log J}{\epsilon}\right), \tag{17}$$

where the $\Omega((\log J)/\epsilon)$ asymptotic form holds for $\epsilon \leq 1/2$.

(ii) ***Targeted upper bound:*** *if in each round (a) every unresolved invariant fails a derived criterion ($\text{Priority}_t(c) > 0$), (b) the winner addresses at least $\min(b_{\min}, u)$ of the u still unresolved invariants, for an integer $b_{\min} \geq 1$, (c) each is resolved and admitted independently with probability $\geq \rho > 0$, and (d) resolved invariants stay resolved, then*

$$\mathbb{E}[T_{\text{target}}] \leq \frac{J}{\rho\, b_{\min}} + 1 + \frac{H_{b_{\min}-1}}{\rho}. \tag{18}$$

Dense per-criterion grading (assumption (a)) removes the $1/\epsilon$ dependence regardless of the specific priority formula: any ranking assigning positive weight to failing criteria, including 1 – obey, satisfies (a), and under all-or-nothing candidate admission with joint bundle pass probability $\rho_{\text{bundle}} > 0$, $\mathbb{E}[T_{\text{target}}] \leq \lceil J/b_{\min} \rceil/\rho_{\text{bundle}} \leq \frac{J}{\rho_{\text{bundle}} b_{\min}} + \frac{\mathbb{I}[b_{\min}>1]}{\rho_{\text{bundle}}}$. What uniquely distinguishes paired priority is regression handling, formalized next.

Proposition 1 (Strict Regression Dominance under Multi-Rollout Grading)**.** *Let obedience (3) average the verdicts of $K_{\text{gr}} \geq 1$ graded rollouts, so that it may be fractional. When round $t-1$ promoted, write $o_c = \text{obey}(c; W_{t-1})$ and $\delta_c = \text{obey}(c; W_{t-2}) - o_c$ ($\delta_c := 0$ if $t = 1$ or $\mathcal{T}_c(W_{t-2}) = \emptyset$), so $\text{Priority}_t(c) = (1 - o_c) + \max(0, \delta_c)$ by (4). Let c be a regressed criterion ($\delta_c > 0$) and c' a non-regressed one ($\delta_{c'} \leq 0$).*

(i) ***Dominance at equal or higher obedience:*** *if* $o_{c'} \geq o_c$*, then* $\text{Priority}_t(c) - \text{Priority}_t(c') \geq \delta_c > 0$*, however small the regression (for example, one additional failing rollout).*

(ii) ***Exact trade-off:*** $\text{Priority}_t(c) > \text{Priority}_t(c')$ *if and only if* $o_c - o_{c'} < \delta_c$*. In the single-rollout binary case* ($K_{\text{gr}} = 1$*), a fresh regression (*$1 \to 0$*, priority* 2*) therefore outranks every non-regressed criterion (priority at most* 1*).*

(iii) ***Exact incumbent pinning set:*** *under* (4) *in any round (promoted or rejected),* $\text{Priority}_t(c) = 0$ *holds if and only if* $\text{obey}(c; W_{t-1}) = 1$ *and* $\Delta_{\text{hurt},t}(c) = 0$*; in particular, no regressed criterion and no criterion failed by incumbent* W_{t-1} *is ever pinned.*

Frequency ranking $1 - \text{obey}(c; W_{t-1})$ satisfies (iii) but fails (i), tying newly broken criteria with static gaps of equal obedience; conversely, signed difference $\text{obey}(c; W_{t-2}) - \text{obey}(c; W_{t-1})$ satisfies (i) but fails (iii), assigning 0 to unchanged failures ($0 \to 0$) alongside mastered invariants ($1 \to 1$). Paired priority satisfies both parameter-free. For $K_{\text{gr}} \geq 3$ rollouts (or $|\mathcal{T}_c| \geq 3$ tasks), a partial regression ($1 \to 2/3$, priority 2/3) scores below a $0 \to 0$ gap (1) by (ii), and tying occurs at $K_{\text{gr}} = 2$ ($1 \to 0.5$, priority 1) or post-rejection when $\text{obey}(c; W_{t-1}) = 1$ ($\Delta_{\text{hurt},t} = 1, \Delta_{\text{gap},t} = 0$). This trade-off is intended: a criterion that the incumbent always fails outranks one that regressed on a minority of rollouts, while the Regressed list still surfaces every regression to the proposer (§4.2).

5. Experimental Setup

To evaluate continual harness RSI against local overfitting, memory pollution, and catastrophic forgetting across escalating non-stationarity, we test `GUARD` on three primary public benchmarks: (1) PAST-Bench [50] (cross-session memory retention, precision, and SOP updates), (2) EvoHarnessBench (`EnterpriseOps-Gym`) [17] (stage-wise continual transfer across evolving Model Context Protocol [MCP] tools and skills), and (3) Drift Gauntlet [5] (continual learning without forgetting across hidden-boundary workflow, topic, domain, and compound drift), complemented by two rubric-dense professional suites and controlled component ablations (`Analyze`, `Propose`, and `Gate`) in Appendix C. All three primary benchmarks compare the unified six-arm suite (Appendix B.2): `No-Evolve` (W_0), native memory/skill persistence (Hermes+ [50] / MemToolAgent [17]), RRSI [45], GEPA [1], Meta-Harness [20], and `GUARD`. To enforce strict reference-free hygiene on public datasets, our runner strips all held-out evaluator fields (reference answers, test cases, verifiers, oracle skills, and rubrics), aborts if any survives, and snapshots initial sandboxes so adaptation sees only public task inputs and execution states; after adaptation, every arm freezes its promoted harness W_T and reports standard single-rollout (`pass@1`) held-out accuracy across all three primary benchmarks ($N_{\text{rep}} = 2$ frozen `DEV` trials on the Appendix C professional suites), with zero test-time candidate pooling, self-review, or retries.

Primary single-model `Gemini 4` RSI and applicability to teacher–student knowledge distillation. Our primary continual-RSI experiments use a single frozen `Gemini 4` model across every role (solver, self-reviewer, distiller, and evolver, with zero stronger teacher): this powers the EvoHarnessBench Official Hosted SDK (Table 2), all four hidden-boundary continual streams on Drift Gauntlet (Table 3), and the controlled dual-axis `Propose`-stage ablation (Table C.8), alongside two full 12-family model-identical PAST-Bench suites (pure `claude-opus-5-5` and pure `gemini-3.8-flash`; §6.1, Appendix B.2). Complementing this primary single-model regime, we deliberately include a cross-model teacher–student configuration on PAST-Bench (Table 1), the five-domain local gym (Tables C.5 and C.7), and the Appendix C professional suites, pairing a frozen `gemini-3.8-flash` student solver with a `claude-opus-5-5` teacher reviewer and distiller shared identically across all evolving baselines. Evaluating both regimes serves a deliberate purpose: the single-model `Gemini 4` suites

isolate strictly self-referential continual improvement without external teacher supervision, while the cross-model suites show that GUARD also applies to teacher–student knowledge-distillation harness pipelines, distilling stronger-model reasoning into compact student skills and memories with low memory pollution and no regressed domain.

PAST-BENCH and critic-scored selection. Across 12 families spanning memory (SM01–02, EP01), procedural reuse (PC01, PC03–04), proactive gathering (PG01–02), and updates (PC02, SM04, SM06–07; 1 round per teach/update episode), official grading reports *Task* (plus procedural SOP *Proc.*), *write precision* (*W.Prec.*), *Recall*, *update correctness* (*Upd.*), *retention* (*Ret.*), *pollution* (*Poll.*), *mechanism* (*Mech.*, mean of W.Prec., Recall, Upd., Ret., and 1 − Poll.), and Evolve (12-family mean of $\Delta_{\text{task},f} \times \text{Mech.}_f$). To avoid re-executing stateful multi-session histories, GUARD uses critic scoring (§4.3), evaluating the incumbent and candidate states in a single pass over accumulated session trajectories instead of re-running prior episodes (Appendix B.2); in place of per-task pre-execution checklists, the critic scores four fixed suite-level criteria (history grounding, SOP fidelity, [iff] scope, and superseded-rule replacement; Box B.1(c)) that overlap with the benchmark's write-path pillars, while baselines keep their native reflection or search prompts. Because a single critic pass scores both the incumbent and the candidates without re-execution, it drives both selection and gating (24/24 round winners admitted across the 12 families; execution-gated promotion fidelity is evaluated separately across EVOHARNESSBENCH, DRIFT GAUNTLET, and the professional suites in §6.4), while held-out official grading and paired per-family comparisons (Table B.4) evaluate downstream write-path quality.

EVOHARNESSBENCH and DRIFT GAUNTLET protocols. On EVOHARNESSBENCH, at stage $v_k \in \{\text{v1}, \ldots, v_V\}$, arms adapt on $\mathcal{D}_{\text{train}}^{(v_k)}$ via execution scoring (§4.1) and evaluate on $\mathcal{D}_{\text{test}}^{(v_1 \ldots v_k)}$, forming matrix $A_{k,j}$ ($k \geq j$). Following [17], we report final *Overall ACC*, strict *Solved*, *Cohort ACC* $A_{V,j}$, backward transfer $\text{BWT} = \sum_{j<V} n_j (A_{V,j} - A_{j,j}) / \sum_{j<V} n_j$, forward transfer $\text{FWT} = \sum_{k \geq 2} n_k (A_{k,k} - A_{k,k}^{\text{pre}}) / \sum_{k \geq 2} n_k$ (local gym uses unweighted cohort BWT and stage FWT $A_{k,k} - A_{k,k}^{W_0}$), and per-task tokens, calls, and latency on: (1) the *Official Hosted SDK* (evolving_tools/calendar, v1–v3, 18 train / 43 test tasks) under single-model Gemini 4 for all roles; and (2) the *Full 5-Domain* local gym (236 test / 97 train tasks, 16 stages across Calendar, Drive, ITSM v1–v4, CSM, HR), plus a controlled single-model Gemini 4 Propose-stage ablation across dual-axis ITSM (N=132 test tasks; Appendix C.4). On DRIFT GAUNTLET [5], agents evolve sequentially across $K \in \{2, 3, 4\}$ unannounced task blocks (30 online adaptation tasks/block, 30 held-out test tasks/group, 12 blocks total) under *single-model* Gemini 4 for every role across four continual drift scenarios spanning BFCL tool calling, BIRD text-to-SQL, and ScienceWorld interactive environments: *workflow drift* (BFCL-Miss → BFCL-Base, K=2), *topic drift* (BIRD-F1 → BIRD-Code → BIRD-TP, K=3), *domain drift* (ScienceWorld → BFCL-Base → BIRD-F1, K=3), and *compound drift* (BIRD-Code → BIRD-F1 → BFCL-Miss → BFCL-Base, K=4). From the $(K{+}1) \times K$ checkpoint matrix $A[t, k]$ ($t \in \{0, \ldots, K\}, k \in \{1, \ldots, K\}$) and seed vector $B_k = A[0, k]$, we report [5]: *Final Mean* $\bar{A}_K = \frac{1}{K} \sum_{k=1}^{K} A[K, k]$, terminal gain $\Delta_K = \bar{A}_K - \bar{B}$, average gain $\bar{\Delta} = \frac{1}{K} \sum_{t=1}^{K} (\bar{A}_t - \bar{B})$, within-block *plasticity* $P = \frac{1}{K} \sum_{k=1}^{K} (A[k, k] - A[k{-}1, k])$, *pre-block interference* $I = \frac{1}{K-1} \sum_{k=2}^{K} \frac{1}{k-1} \sum_{t=1}^{k-1} (A[t, k] - B_k)$, worst-case *post-block forgetting* $F = \frac{1}{K-1} \sum_{k=1}^{K-1} (\min_{t>k} A[t, k] - A[k, k])$, and *stability* $S = (I + F)/2$. GUARD keys playbook sections to each task's pre-execution observable environment/schema signature and renders only the matching section, whereas baselines keep one global playbook (lexically retrieved skills for HERMES+); for every arm, held-out rollouts are memoized per task and rendered playbook (Appendix B). Arms are compared strictly within each benchmark and model regime.

Table 1 | **PAST-BENCH results across 12 families (matched backbone, unified six-arm suite; metrics in §5).** All columns except Poll. are higher-is-better; bold marks the best persistent arm. GUARD leads Proc., Mech., Upd., W.Prec., and Evolve, stays within 0.001 of the Ret. ceiling, and ties the 0.167 Poll. floor, with positive mean gains on Mech., Upd., and W.Prec. against every persistent baseline and 0/48 family losses on W.Prec. or Poll. across all four baseline comparisons (paired differences and sign tests in Table B.4).

Arm	Task	Proc.	Mech.	Upd.	W.Prec.	Recall	Ret.	Poll.↓	Evolve
`No-Evolve` (W_0)	0.372	0.427	0.454	0.278	0.000	0.000	0.990	0.000	0.000
HERMES+ [50]	0.852	0.808	0.840	0.791	0.724	0.854	**1.000**	**0.167**	0.421
RRSI [45]	0.809	0.714	0.725	0.901	0.438	0.811	0.974	0.500	0.350
GEPA [1]	**0.858**	0.815	0.837	0.777	0.719	**0.858**	0.996	**0.167**	0.424
META-HARNESS [20]	0.812	0.677	0.739	0.903	0.479	0.814	0.998	0.500	0.360
GUARD (Ours)	0.854	**0.830**	**0.878**	**0.924**	**0.776**	0.856	0.999	**0.167**	**0.436**

6. Main Experimental Results

6.1. PAST-BENCH: Cross-Session Memory Write Path

While 5 of 12 families are single-rule tasks where every persistent arm reaches 0.984–1.000, write-path attribution and multi-step procedural SOPs (PC01, PC03, PC04) expose a sharp plasticity–pollution trade-off (Table 1, Figure 3; Tables B.3–B.4). Unscoped accumulators (HERMES+, GEPA) reach 80.8%–81.5% procedural task at 0.167 overall pollution, but retain stale clauses on updates (52.9% / 46.4% procedural; 0.791 / 0.777 overall). Conversely, aggressive rewriters (RRSI, META-HARNESS) lift update correctness to 0.901–0.903 (82.4%–84.1% procedural) while tripling pollution (0.500 overall, 1.000 procedural) and depressing write precision (0.438–0.479). Pairing regret ranking, rule pinning, and `[iff]` guards, GUARD leads procedural task (83.0%), procedural mechanism (90.0% vs. at most 78.0%), and update (90.1%) at 0.000 procedural pollution; across all 12 families (12/12 $>$ W_0), it leads update (0.924, +13.3pp vs. HERMES+, +14.8pp vs. GEPA), mechanism (0.878, p=0.016 vs. RRSI), write precision (0.776, p=0.031 vs. RRSI), and evolve score (0.436) within 0.004 of top task score at the 0.167 pollution floor (0.000 on 10/12 families, and 0/48 family losses on W.Prec. or Poll. across all four baselines; Table B.4). These gains persist without any teacher–student model gap (Appendix B.2): under model-identical all-`claude-opus-5-5` execution, GUARD ranks first overall on task score (0.870 vs. ≤ 0.866, +0.523 over W_0), recall (0.871), and evolve score (0.458) with the sole lowest memory pollution (0.167 vs. 0.208–0.250); under all-`gemini-3.8-flash`, it lifts task score from 0.371 to 0.852, ties for lowest pollution (0.167), and leads latent rule induction (PC03: 0.748 vs. 0.510–0.653 on reflective/search baselines, +9.5 to +23.8pp).

6.2. EVOHARNESSBENCH: Continual Evolution Across Tools and Skills

Unified `Gemini 4` hosted comparison, efficiency, and 5-domain local gym. On the hosted SDK (Table 2, Figure 4), GUARD achieves the highest accuracy (80.80%; +8.89pp over W_0, paired p = 0.008; +6.28pp over META-HARNESS, p = 0.11 on 43 tasks) and solve rate (58.14%, 25/43 vs. ≤ 22/43), leading Cohorts v1 (78.47%) and v3 (89.29%) and tying Cohort v2 (84.17%), while passing 182/221 official verifier checks (82.4% vs. 156–170/221 [70.6%–76.9%] on baselines) with the fewest zero-score tasks (3/43 vs. 5–7/43). Soft-HURT promoted the Stage-v1 playbook (2,315 chars, 83.33 → 100 source self-review, 0 hurt) and abstained at Stages v2–v3, so FWT = 0.00pp holds by abstention and the entire lift stems from that Stage-v1 playbook. GUARD is the only arm with no negative entry across BWT Score, BWT Pass, and FWT (+0.85 / +2.56 / 0.00pp), whereas RRSI drops −4.50pp on Cohort v2 at Stage v3 (77.67% → 73.17%, BWT Score −0.45pp) and GEPA (−3.97pp BWT), META-HARNESS (−13.21pp FWT), and MEM-

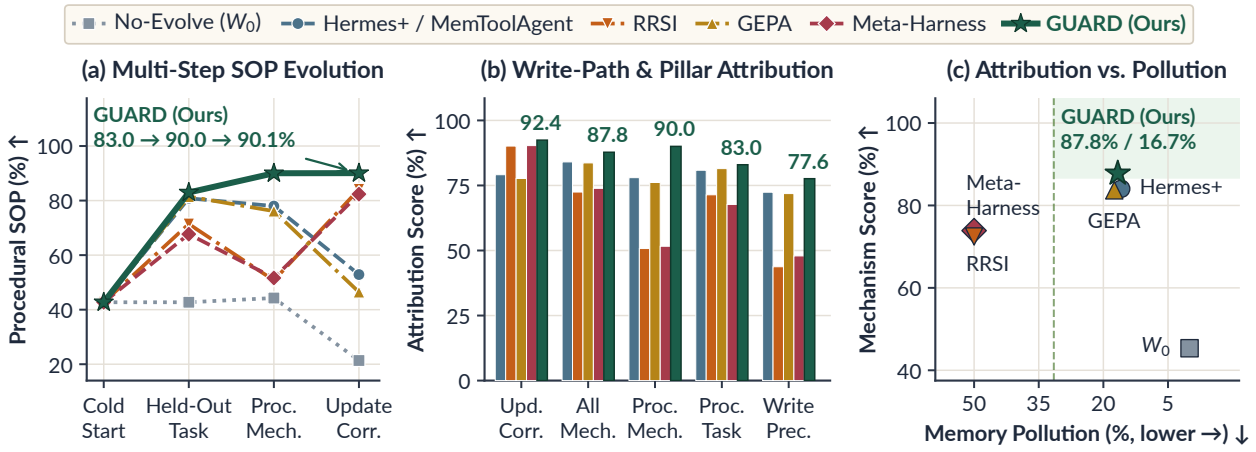


Figure 3 | **PAST-Bench procedural SOP and memory write-path attribution (12 families).** (a) Multi-step SOP families (PC01, PC03, PC04; horizontal axis lists metrics alongside cold-start 42.7%, not training stages). (b) Write-path and pillar breakdown. (c) Mechanism attribution vs. memory pollution (0.0% pollution on 10/12 families under GUARD).

Table 2 | **EvoHarnessBench Official Hosted SDK comparison under unified Gemini 4 (evolving_tools/calendar, stages v1–v3, 43 test tasks; §5).** Cohort sizes are n_1=29, n_2=10, n_3=4 (Cohort v3 reflects single-task granularity); rollout hygiene notes in Appendix B. The GUARD row is the final of five hosted configurations evaluated on this test split (mean 77.09%, range 75.57–80.80%; all five above every baseline). Paired over the 43 test tasks (exact sign-flip test), the ACC gain of GUARD is significant against W_0 (p = 0.008; 8 tasks better, 0 worse), GEPA (p = 0.023), and MemToolAgent (p = 0.010), and directional against RRSI (p = 0.094) and Meta-Harness (p = 0.11; 95% bootstrap CI [−0.4, +14.5]pp). Bold marks best per column ($^=$exact Cohort-v2 tie).

Method (Unified Gemini 4)	Adaptation Family	ACC	Solved	Cohort v1 / v2 / v3 ACC	BWT Score	BWT Pass	FWT	Tokens / Task↓	Calls↓	Latency↓
No-Evolve (W_0, Control)	Static ReAct Baseline	71.91%	48.84%	70.46% / 79.17% / 64.29%	−6.00pp	−5.13pp	–	131,032	14.47	221.6s
MemToolAgent [17]	Memory (#1 family)	68.00%	51.16%	61.95% / **84.17%**$^=$ / 71.43%	−4.02pp	+2.56pp	−4.98pp	82,426	15.84	239.2s
RRSI [45]	Reflective Harness Rewrite (M=1)	74.28%	51.16%	75.06% / 73.17% / 71.43%	−0.45pp	0.00pp	+1.79pp	84,519	13.19	613.1s
GEPA [1]	Prompt (#1 family)	70.69%	48.84%	68.45% / 71.17% / 85.71%	−3.97pp	−2.56pp	**+13.61pp**	85,156	17.05	**159.1s**
Meta-Harness [20]	Agent Code (#1 family)	74.52%	51.16%	73.42% / 81.83% / 64.29%	**+5.83pp**	**+5.13pp**	−13.21pp	78,705	14.19	220.0s
GUARD (Ours)	**Regret-Guided Harness RSI**	**80.80%**	**58.14%**	**78.47% / 84.17%**$^=$ **/ 89.29%**	+0.85pp	+2.56pp	0.00pp	**59,651** (−54.5%)	**11.21** (−22.5%)	233.8s

ToolAgent (−4.02 / −4.98pp) each suffer negative transfer. Together with two-layer gating (§6.4) and a compact playbook (2,315 vs. 4,538–10,013 chars), GUARD reaches 0/43 eight-step budget exhaustions (vs. 7/43 in W_0 and 6–10/43 on baselines), which cuts tokens/task by 54.5% (59,651 vs. 131,032) and MCP calls by 22.5% while solving 9.75 tasks per 1M tokens (1.50× the best baseline, 6.50 on Meta-Harness, and 2.61× W_0, 3.73). Every one of the five hosted GUARD configurations evaluated in our study (75.57%–80.80% ACC, mean 77.09%, 171–182/221 checks passed, 0–3/43 budget exhaustions; Appendix B) beats every external baseline (≤ 74.52% ACC, ≤ 170/221 checks, 6–10/43 exhaustions). Across all five local-gym domains (236 test tasks; Table C.5), GUARD improves upon or matches W_0 in final accuracy across domains (macro 65.81% → 66.46%, +1.00pp mean FWT), lifting multi-stage skill tracks such as ITSM (+1.63pp ACC, +8.47pp Solved) and CSM (+1.60pp ACC, with a minor trade-off on strict all-or-nothing solve rate) while safely abstaining where candidate edits do not clear Soft-HURT.

6.3. Drift Gauntlet: Continual RSI Under Hidden Stream Drift

Sustained multi-block progression and the plasticity–stability trade-off. While EvoHar-

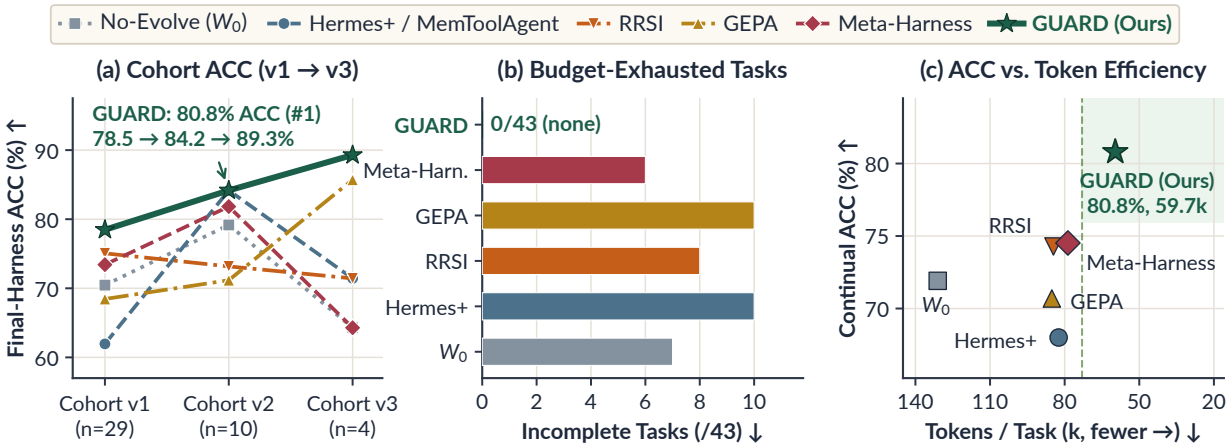


Figure 4 | **EvoHarnessBench Hosted SDK evaluation under unified `Gemini 4` (`evolving_tools/calendar`, v1–v3).** (a) Final-harness cohort progression $A_{3,j}$. (b) Tasks exhausting the eight-step budget (0/43 for `GUARD` vs. 6–10 on baselines). (c) Overall accuracy vs. mean tokens per task.

Table 3 | **Drift Gauntlet [5] continual RSI comparison under single-model `Gemini 4` across all four hidden-boundary drift streams (12 blocks, 360 adaptation / 360 held-out test rollouts per arm; metrics in §5, full 24-row breakdown in Table C.4).** *Block Drops* counts checkpoints $t \in \{1, \ldots, K\}$ where stream-mean accuracy $\bar{A}_t < \bar{A}_{t-1}$. Bold marks the best evolving arm per column ($^{=}$exact stream tie).

	Four-Stream Macro-Average							Per-Stream Final Mean (F)			
Method (Single-Model `Gemini 4`)	**Final Mean**	Δ_K	$\bar{\Delta}$	**Plast.** P	**Forgetting** F	**Stab.** S	**Block Drops↓**	**Workflow**	**Topic**	**Domain**	**Compound**
`No-Evolve` (W_0, Control)	69.92%	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	0 / 12	65.00% (0.0)	67.78% (0.0)	76.05% (0.0)	70.84% (0.0)
Hermes+ / MemToolAgent	70.31%	+0.40pp	+1.09pp	+2.67pp	−5.59pp	−2.66pp	5 / 12	65.00% (−13.3)	**70.00%**$^{=}$ (+1.7)	76.26% (−1.8)	70.00% (−8.9)
RRSI [45]	67.92%	−2.00pp	−0.47pp	+4.49pp	−8.11pp	−6.02pp	8 / 12	60.00% (−20.0)	66.67% (−3.3)	79.17% (−5.8)	65.84% (−3.3)
GEPA [1]	72.40%	+2.48pp	+2.18pp	+2.80pp	−1.09pp	**−0.35**pp	3 / 12	**70.00%**$^{=}$ (**+3.3**)	67.78% (−1.7)	**79.31%** (−2.7)	72.50% (−3.3)
Meta-Harness [20]	71.21%	+1.30pp	+0.15pp	+3.28pp	−0.07pp	−2.51pp	4 / 12	63.33% (0.0)	68.89% (−5.0)	79.29% (**+1.4**)	73.33% (**+3.3**)
`GUARD` (Ours)	**73.22%**	**+3.30**pp	**+2.71**pp	**+4.55**pp	**0.00**pp	−0.93pp	**0 / 12**	**70.00%**$^{=}$ (0.0)	**70.00%**$^{=}$ (0.0)	78.69% (0.0)	**74.17%** (0.0)

nessBench shifts tool schemas within one domain at a time, Drift Gauntlet [5] stress-tests whether self-evolving harnesses can accumulate capabilities across unannounced workflow, topic, domain, and compound shifts without boundary resets. Table 3 and Figure 5 (per-group trajectories in Figure C.2) show two main results under single-model `Gemini 4`. `GUARD` is the only arm whose stream-mean accuracy never decreases: across all 12 continual block transitions, `GUARD` incurs 0/12 stream-mean drops (65.00 → 70.00 → 70.00% on *workflow*; 67.78 → 67.78 → 68.89 → 70.00% on *topic*; 76.05 → 78.69 → 78.69 → 78.69% on *domain*; 70.84 → 71.67 → 71.67 → 74.17 → 74.17% on *compound*). Its stage-normalized macro gain rises at every checkpoint (0.00 → +2.12 → +2.40 → +3.30pp), ranking first on Final Mean (73.22%, 0.82pp above GEPA, about three task-equivalents of 360 held-out rollouts), Δ_K (+3.30pp, positive on all 4/4 streams), and $\bar{\Delta}$ (+2.71pp). Every evolving baseline suffers 3–8 block drops out of 12: Hermes+ peaks at Block 1 before decaying (+1.76 → +0.91 → +0.40pp), RRSI falls below W_0 (+0.80 → −0.65 → −2.00pp), GEPA dips −1.35pp at Block 2 (+2.94 → +1.59 → +2.48pp), and Meta-Harness drops −1.65pp below W_0 at Block 1 (−20.00pp pre-block interference on *workflow drift*). Consistent with [5], plasticity P and stability S are strongly negatively correlated across the 16 (baseline, stream) runs (Pearson $r = -0.82$, reported descriptively because P and S share checkpoint cells): high-plasticity RRSI (P = +4.49pp) suffers severe post-block forgetting (F = −8.11pp, −20.00pp on opposing `BFCL-Miss` →

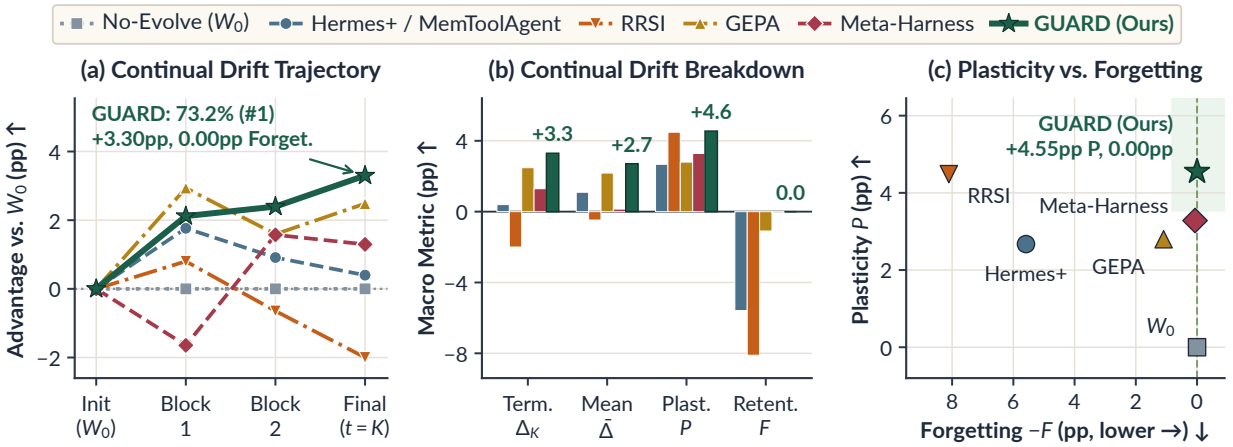


Figure 5 | **Drift Gauntlet [5] continual RSI under single-model `Gemini 4` (four drift streams).** (a) Stage-normalized macro gain trajectory across continual checkpoints (t=0, Block 1, Block 2, Final): `GUARD` is the sole method improving monotonically (0.00 → +2.12 → +2.40 → +3.30pp). (b) Macro continual metrics breakdown (Δ_K, $\bar{\Delta}$, plasticity P, post-block retention F, where 0.00pp denotes zero forgetting). (c) Plasticity P vs. post-block forgetting loss $-F$: `GUARD` lies above the baseline plasticity–stability trend, combining top plasticity (+4.55pp) with 0.00pp post-block forgetting.

`BFCL-Base`), whereas GEPA's monolithic 4,200-char prompt rewrite compresses prior-block character share on *compound drift* (`BIRD-Code` 100% → 28.5%; `BFCL-Miss` 49.4% → 22.4% once `BFCL-Base` arrives, dropping `BFCL-Miss` accuracy 73.33% → 66.67%). By keying sections to pre-execution observable tool/schema signatures (0 prior-signature characters overwritten) and gating updates against same-signature Vitter reservoir replay tasks (promoting 9/12 blocks, including in-round `[iff]` repair on `ScienceWorld` to reach 82.73%, and rejecting 3/12 opposing-drift overwrites where hurt equaled lift), `GUARD` achieves both top plasticity (+4.55pp) and 0.00pp post-block forgetting across all four streams, with only mild pre-block caution before opposing tool conventions are observed. On *workflow* and *compound* drift, where opposing blocks share the `bfcl` signature, Soft-HURT directly rejects both regressive overwrites (worked example in Box C.1). In a symmetric routing control (Appendix B), granting the same pre-execution signature routing to GEPA (71.55%, F = −2.78pp) and Meta-Harness (71.08%) leaves `GUARD` ahead across all four streams (73.22%), while a single-global-playbook `GUARD` ablation without signature routing still achieves 72.35% Final Mean, the highest global trajectory gain ($\bar{\Delta}$ = +2.25pp), and zero post-block forgetting (F = 0.00pp).

6.4. Self-Reviewer Calibration, Gate Fidelity, and Component Attribution

Table 4 (full audit across 4,173 EvoHarnessBench/PAST-Bench rollouts plus Drift Gauntlet block decisions in Table C.6) checks whether $S(\tau)$ and Soft-HURT track hidden verifiers. Layer-1 sandbox checks curb LLM-only optimism: on hosted EvoHarnessBench, a 1-layer reviewer passes (≥ 75%) 22.2% (12/54) of execution-broken traces vs. 0.0% (0/23, sanity check by Eq. (2)) under 2-layer review. The final hosted configuration (two-layer review, cumulative-source gating, and a completeness-first distiller instruction; Appendix B) improves on the initial one-layer run by +4.82pp ACC (75.98% → 80.80%) and +9.30pp Solved (48.84% → 58.14%), with budget exhaustions dropping from 2/43 to 0/43. Across 4,059 EvoHarnessBench test rollouts (1,920 hosted, 2,139 local; all arms), Layer 1 separates official ACC by +27.92pp (73.90% vs. 45.97%; 82.6% NPV); on PAST-Bench (N=114), Layer 2 has

Table 4 | **Summary of offline self-reviewer calibration and Soft-HURT gate fidelity** (zero oracle labels exposed during adaptation; NPV = negative predictive value; full 9-row audit in Table C.6).

Calibration / Gate Check (Sample N)	**Ref. / 1-L**	**GUARD**
Initial vs. final hosted ACC / Solved (43)	75.98/48.84%	**80.80**/**58.14**%
Official ACC: L1-fail vs. L1-pass (4,059 test)	45.97%	**73.90**% (82.6% NPV)
PAST playbook rank concord. (169 pairs)	50.0%	**82.2**% (75.6% strict)
Exec.-gated promotions $\Delta_{\text{test}} \geq 0$[†]	—	**19**/**19** (15 >0, 4 ties)
Same gate w/o `Analyze` (16 local stages)	4 prom., 2/5 ↓	**4** prom., **0**/**5** ↓
Bypassing Soft-HURT on `Calendar v3`	−6.98pp Solv.	**0.00**pp (rejected)

[†] Δ_{test}: held-out score of a promoted harness minus W_0 on the active cohort/group of the promoted stage/block (at H_{v_3} on hosted `Calendar`, where W_3=W_1; `CSM v2`'s delta over `CSM v1` is not separately tabulated). Execution-gated suites give 19/19 $\Delta_{\text{test}} \geq 0$ (15/19 >0): 9/9 Drift Gauntlet blocks (6 >0 [+3.33 to +16.67pp], 3 ties; Table C.4), 5/5 EvoHarness-Bench stages (hosted `v1` +8.01, `CSM v1/v2` +4.17/+3.84, `HR v1` +1.90pp, `ITSM v1` 0.00pp tie), and 5/5 prof. role promotions (PRBench `R1 Finance/Legal` +28.57/+5.77pp; FinRate `R1–R3` +6.25/+12.50/+12.50pp). Critic-scored PAST-Bench (24/24 admitted) lifts 12-family Task 0.372 → 0.854 (all 12/12 > W_0).

near-zero mean bias (−0.012, MAE = 0.131, r=0.369) and 82.2% within-family concordance (75.6% non-tied, r=0.432). Under active-cohort Δ_{test} vs. W_0, all 19/19 execution-gated Soft-HURT promotions across Drift Gauntlet (9/9), EvoHarnessBench (5/5), and the professional suites (5/5) satisfy $\Delta_{\text{test}} \geq 0$ (15/19 strictly >0 and 4/19 0.00pp ties; Table 4 note), whereas bypassing Soft-HURT on a rejected local `Calendar v3` candidate lowers held-out accuracy and solve rate. In the five-domain teacher–student `Analyze`/`Gate` component ablation (236 tasks; Table C.7), removing `Analyze` (1 – obey, no pinning or contrasts) drops macro ACC by −7.33pp (66.46% → 59.13%, promoting 4 stages that regress multi-stage skill domains), and removing Soft-HURT drops macro ACC by −10.35pp (66.46% → 56.11%, regressing 3/5 domains vs. 0/5). Under single-model `Gemini 4` on dual-axis EvoHarnessBench `ITSM` (N=132 test tasks, N_{rep}=3; Table C.8), parallel M=2 candidate pooling (81.08% ACC; BWT +0.61%) beats matched-budget sequential M=1, 2T search (78.85%, +2.23pp, p=0.0267; 64% fewer rejected gate rounds) and single-candidate M=1, T (77.20%, +3.88pp, p=0.0029).

7. Conclusion

Continually adapting an agent's harness around a frozen backbone model without external graders requires preventing local overfitting, catastrophic forgetting, and memory pollution. Drawing on curriculum regret from UED and group-relative scoring from reinforcement learning (GRPO), **GUARD** (*Grader-free UED with Adaptive Rule Distillation*) treats the text harness as a non-parametric policy component and updates it from task-derived checklists: `Analyze` separates edit-induced regressions (Δ_{hurt}) from unresolved gaps (Δ_{gap}) while pinning mastered criteria, `Propose` selects among parallel candidate edits via critic-free group-standardized fitness and contrasts the best and worst candidates to rewrite proposal prompts, and `Gate` (*Soft-HURT*) promotes only net-positive updates with formal non-regression guarantees on the self-review objective (Theorems 1–2, Proposition 1). At test time, it deploys a single frozen harness in one pass. Across four benchmark settings, `GUARD` improves upon or matches W_0 in held-out accuracy, achieves 19/19 non-regressive execution-gated promotions, and sustains monotonic `pass@1` gains with zero post-block forgetting under hidden stream drift while cutting hosted tokens per task by 54.5%.

References

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

[2] Afra Feyza Akyürek, Advait Gosai, Chen Bo Calvin Zhang, Vipul Gupta, Jaehwan Jeong, Anisha Gunjal, Tahseen Rabbani, Maria Mazzone, David Randolph, et al. PRBench: Large-scale expert rubrics for evaluating high-stakes professional reasoning. *arXiv preprint arXiv:2511.11562*, 2025.

[3] Yuzheng Cai, Siqi Cai, Yuchen Shi, Zihan Xu, Lichao Chen, Yulei Qin, Xiaoyu Tan, Gang Li, Zongyi Li, Haojia Lin, Yong Mao, Ke Li, and Xing Sun. Training-

free group relative policy optimization. *arXiv preprint arXiv:2510.08191*, 2025.

[4] Minghao Chen, Yihang Li, Yanting Yang, Shiyu Yu, Binbin Lin, and Xiaofei He. AutoManual: Constructing instruction manuals by LLM agents via interactive environmental learning. In *NeurIPS*, 2024.

[5] Xiaoyin Chen, Yizheng Huang, Yinglun Xu, Ziyun Zhao, Emmett McQuinn, Julie Liu, and Hope Yao. Are self-evolving agents continual learners? *arXiv preprint arXiv:2610.25858*, 2026.

[6] Jonathan Cook, Tim Rocktäschel, Jakob Foerster, Dennis Aumiller, and Alex Wang. TICKing all the boxes: Generated checklists improve LLM evaluation and generation. *arXiv preprint arXiv:2410.03608*, 2024.

[7] Michael Dennis, Natasha Jaques, Eugene Vinitsky, Alexandre Bayen, Stuart Russell, Andrew Critch, and Sergey Levine. Emergent complexity and zero-shot transfer via unsupervised environment design. In *NeurIPS*, 2020.

[8] Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. Promptbreeder: Self-referential self-improvement via prompt evolution. *arXiv preprint arXiv:2309.16797*, 2023.

[9] Yao Fu, Dong-Ki Kim, Jaekyeom Kim, Sungryull Sohn, Lajanugen Logeswaran, Kyunghoon Bae, and Honglak Lee. AutoGuide: Automated generation and selection of context-aware guidelines for large language model agents. In *NeurIPS*, 2024.

[10] Lucheng Fu, Ye Yu, Yiyang Wang, Yiqiao Jin, Haibo Jin, B. Aditya Prakash, and Haohan Wang. TextReg: Mitigating prompt distributional overfitting via regularized text-space optimization. *arXiv preprint arXiv:2605.21318*, 2026.

[11] Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. EvoPrompt: Connecting LLMs with evolutionary algorithms yields powerful prompt optimizers. In *ICLR*, 2024.

[12] Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. In *ICLR*, 2025.

[13] Chengsong Huang, Wenhao Yu, Xiaoyang Wang, Hongming Zhang, Zongxia Li, Ruosen Li, Jiaxin Huang, Haitao Mi, and Dong Yu. R-Zero: Self-evolving reasoning LLM from zero data. In *ICLR*, 2026.

[14] Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized level replay. In *ICML*, 2021.

[15] Minqi Jiang, Michael Dennis, Jack Parker-Holder, Jakob Foerster, Edward Grefenstette, and Tim Rocktäschel. Replay-guided adversarial environment design. In *NeurIPS*, 2021.

[16] Yidong Jiang, Junrong Chen, Eftychia Makri, Jialin Chen, Peiwen Li, Ali Maatouk, Leandros Tassiulas, Eliot Brenner, Bing Xiang, and Rex Ying. Fin-RATE: A real-world financial analytics and tracking evaluation benchmark for LLMs on SEC filings. *arXiv preprint arXiv:2602.07294*, 2026.

[17] Zixuan Ke, Vaidehi Patil, Haizhou Shi, Yang Li, Ye Liu, Sarath Shekkizhar, Anurag Koul, Jiayu Wang, Xuan Phi Nguyen, Semih Yavuz, Mohit Bansal, and Shafiq Joty. EvoHarnessBench: Can your agents keep pace with an evolving harness? *arXiv preprint arXiv:2609.04280*, 2026.

[18] Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. DSPy: Compiling declarative language model calls into self-improving pipelines. In *ICLR*, 2024.

[19] Joel Lehman and Kenneth O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.

[20] Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-Harness: End-to-end optimization of model harnesses. *arXiv preprint arXiv:2603.28052*, 2026.

[21] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *NeurIPS*, 2017.

[22] Xinghua Lou, Miguel Lázaro-Gredilla, Antoine Dedieu, Carter Wendelken, Wolfgang Lehrach, and Kevin P. Murphy. AutoHarness: Improving LLM agents by automatically synthesizing a code harness. *arXiv preprint arXiv:2603.03329*, 2026.

[23] Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. In *NeurIPS*, 2023.

[24] Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra, Peter Jansen, Oyvind Tafjord, Niket Tandon, Li Zhang, Chris Callison-Burch, and Peter Clark. CLIN: A continually learning language agent for rapid task adaptation and generalization. In *COLM*, 2024.

[25] Fanqing Meng, Lingxiao Du, Qiguang Chen, Ziqi Zhao, Haocheng Lu, Mengkang Hu, and Michael Qizhe Shieh. RSIBench-Data: Benchmarking data-centric research for recursive self-improvement. *arXiv preprint arXiv:2607.25886*, 2026.

[26] Fanqing Meng, Lingxiao Du, Haocheng Lu, Qiguang Chen, Ziqi Zhao, Zijian Wu, Jiayuan Zhuo, Mengkang Hu, and Michael Qizhe Shieh. RSIGym: A flexible environment for recursive self-improvement. *arXiv preprint arXiv:2610.10310*, 2026.

[27] Charles Moslonka, Amaury de Vitry, Arthur Garnier, Hicham Randrianarivo, and Emmanuel Malherbe. LEDGER: A long-context benchmark of corporate annual reports. *arXiv preprint arXiv:2606.13100*, 2026.

[28] Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

[29] Deepak Nathani, Lovish Madaan, Nicholas Roberts, Nikolay Bashlykov, Ajay Menon, Vincent Moens, Amar Budhiraja, Despoina Magka, Vladislav Vorotilov, Gaurav Chaurasia, Dieuwke Hupkes, Ricardo Silveira Cabral, Tatiana Shavrina, Jakob Foerster, Yoram Bachrach, William Yang Wang, and Roberta Raileanu. MLGym: A new framework and benchmark for advancing AI research agents. *arXiv preprint arXiv:2502.14499*, 2025.

[30] Siru Ouyang, Jun Yan, I-Hung Hsu, Yanfei Chen, Ke Jiang, Zifeng Wang, Rujun Han, Long T. Le, Samira Daruki, Xiangru Tang, Vishy Tirumalashetty, George Lee, Mahsan Rofouei, Hangfei Lin, Jiawei Han, Chen-Yu Lee, and Tomas Pfister. ReasoningBank: Scaling agent self-evolving with reasoning memory. *arXiv preprint arXiv:2509.25140*, 2025.

[31] Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.

[32] Jack Parker-Holder, Minqi Jiang, Michael Dennis, Mikayel Samvelyan, Jakob Foerster, Edward Grefenstette, and Tim Rocktäschel. Evolving curricula with regret-based environment design. In *ICML*, 2022.

[33] Yiwei Qin, Kaiqiang Song, Yebowen Hu, Wenlin Yao, Sangwoo Cho, Xiaoyang Wang, Xuansheng Wu, Fei Liu, Pengfei Liu, and Dong Yu. InFoBench: Evaluating instruction following ability in large language models. In *Findings of ACL*, 2024.

[34] Ben Rank, Hardik Bhatnagar, Ameya Prabhu, Shira Eisenberg, Karina Nguyen, Matthias Bethge, and Maksym Andriushchenko. PostTrainBench: Can LLM agents automate LLM post-training? *arXiv preprint arXiv:2603.08640*, 2026.

[35] Alexander Rutherford, Michael Beukman, Timon Willi, Bruno Lacerda, Nick Hawes, and Jakob Foerster. No regrets: Investigating and improving regret approximations for curriculum discovery. In *NeurIPS*, 2024.

[36] Rulin Shao, Shannon Zejiang Shen, Junjie Oscar Yin, Yuetai Li, Minheng Wang, Hamish Ivison, Radha Poovendran, Nathan Lambert, Teng Xiao, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, and Pang Wei Koh. Context language models. *arXiv preprint arXiv:2609.37725*, 2026.

[37] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

[38] Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *NeurIPS*, 2023.

[39] Mirac Suzgun, Mert Yuksekgonul, Federico Bianchi, Dan Jurafsky, and James Zou. Dynamic cheatsheet: Test-time learning with adaptive memory. *arXiv preprint arXiv:2504.07952*, 2025.

[40] Liyan Tang, Cyrus Rashtchian, Chun-Sung Ferng, Andrew Tomkins, Da-Cheng Juan, and Tu Vu. WikiSkill: Compiling agent experience into persistent knowledge for skill evolution. *arXiv preprint arXiv:2608.27454*, 2026.

[41] Vijay Viswanathan, Yanchao Sun, Shuang Ma, Xiang Kong, Meng Cao, Graham Neubig, and Tongshuang Wu. Checklists are better than reward models for aligning language models. In *NeurIPS*, 2025.

[42] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.

[43] Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. Agent workflow memory. *arXiv preprint arXiv:2409.07429*, 2024.

[44] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345, 2024.

[45] Peng Xia, Rujun Han, Zifeng Wang, Yanfei Chen, Yufan Zhuang, Yoonho Lee, Chengsong Huang, Han Yu, Zhongying CuiZhu, Yifei Ming, Huaxiu Yao, Burak Gokturk, Tomas Pfister, and Chen-Yu Lee. RRSI: Regularized recursive self-improvement of agent harnesses. *arXiv preprint arXiv:2609.24972*, 2026.

[46] Peng Xia, Kaide Zeng, Jiaqi Liu, Can Qin, Fang Wu, Yiyang Zhou, Caiming Xiong, and Huaxiu Yao. Agent0: Unleashing self-evolving agents from zero data via tool-integrated reasoning. In *COLM*, 2026.

[47] Yuqing Xie, Yi-An Lai, Yuanjun Xiong, Yi Zhang, and Stefano Soatto. Regression bugs are in your model! Measuring, reducing and analyzing regressions in NLP model updates. In *ACL*, 2021.

[48] Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. WizardLM: Empowering large pre-trained language models to follow complex instructions. In *ICLR*, 2024.

[49] Ruiyao Xu, Yanfei Chen, Zhongying CuiZhu, Bhavana Dalvi Mishra, Yifei Ming, Han Yu, Rujun Han, Chen-Yu Lee, and Tomas Pfister. HARNESS^2: Recursive agent harnessing for an open world. *arXiv preprint arXiv:2610.14224*, 2026.

[50] Shuhan Xue, Zixin Ding, Yichen Shen, Yinjie Wang, Zhenfei Yin, Yingcheng Wu, Yuxin Chen, Mengdi Wang, and Ling Yang. PAST-Bench: Benchmarking the foundations of recursive self-improvement in personal agents. *arXiv preprint arXiv:2608.04003*, 2026.

[51] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers. In *ICLR*, 2024.

[52] John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. SWE-agent: Agent-computer interfaces enable automated software engineering. In *NeurIPS*, 2024.

[53] Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and James Zou. TextGrad: Automatic "differentiation" via text. *arXiv preprint arXiv:2406.07496*, 2024.

[54] Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. STaR: Bootstrapping reasoning with reasoning. In *NeurIPS*, 2022.

[55] Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. AFlow: Automating agentic workflow generation. In *ICLR*, 2025.

[56] Qizheng Zhang, Changran Hu, Shubhangi Upasani, Boyuan Ma, Fenglu Hong, Vamsidhar Kamanuru, Jay Rainton, Chen Wu, Mengmeng Ji, Hanchen Li, Urmish

Thakker, James Zou, and Kunle Olukotun. Agentic context engineering: Evolving contexts for self-improving language models. *arXiv preprint arXiv:2510.04618*, 2025.

[57] Hangfan Zhang, Shao Zhang, Kangcong Li, Chen Zhang, Yang Chen, Yiqun Zhang, Lei Bai, and Shuyue Hu. Self-Harness: Harnesses that improve themselves. *arXiv preprint arXiv:2606.09498*, 2026.

[58] Jenny Zhang, Shengran Hu, Cong Lu, Robert Lange, and Jeff Clune. Darwin Gödel machine: Open-ended evolution of self-improving agents. In *ICLR*, 2026.

[59] Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. ExpeL: LLM agents are experiential learners. In *AAAI*, 2024.

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

[61] Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*, 2023.

[62] Yifei Zhou, Sergey Levine, Jason Weston, Xian Li, and Sainbayar Sukhbaatar. Self-challenging language model agents. In *NeurIPS*, 2025.

Supplementary Material & Technical Appendix
GUARD: Stabilizing Recursive Self-Improvement with Regret Guidance

Minghe Zhang* Xuetao Kang* Yifei Ming Yanfei Chen Han Yu
Yufan Zhuang Chen-Yu Lee

Google Cloud

{minghezhang,xuetaok,yifeim,yanfeichen,hanyyu,yufanzhuang,chenyulee}@google.com

*Equal contribution.

A. Proofs

A.1. Proof of Theorem 1 (Non-Regressive Gating and Concentration)

Proof. Part (i). If round t promotes under the parameterized Soft-HURT gate (12)–(13), whether directly ($W_t = \widetilde{W}_t$) or after in-round `[iff]` repair ($W_t = \widetilde{W}_t^{\mathrm{rep}}$ on its re-scored trajectories), the promoted playbook satisfies

$$\Delta_t \;>\; \mathrm{Hurt}(W_t, W_{t-1}) \;\geq\; 0,$$

where $\mathrm{Hurt}(W_t, W_{t-1}) = 0$ whenever $Q_t^- = \emptyset$. If round t rejects, $W_t = W_{t-1}$ and $\Delta_t = 0$. The incumbent score used by the gate is the score recorded when the incumbent was itself evaluated, so the increments telescope and give (14) with non-negative summands. A band of width $\delta_{\mathrm{band}} > 0$ relative to the incumbent instead admits any candidate with $\Delta_t \geq -\delta_{\mathrm{band}}$, so T consecutive admissions at $\Delta_t = -\delta_{\mathrm{band}}$ accumulate a loss of $T\delta_{\mathrm{band}}$; a band relative to the best score so far bounds the aggregate loss by δ_{band} but leaves per-task regressions unbounded.

Part (ii). By (11),

$$|Q_g| \cdot \mathrm{Hurt}(W_t, W_{t-1}) \;=\; \sum_{q \in Q_t^-} \big(S_q(W_{t-1}) - S_q(W_t)\big),$$

and promotion requires $\mathrm{Hurt} < \Delta_t$, which is (15); whenever $Q_t^- = \emptyset$ (in particular when $h_{\max} = 0$), the left-hand side vanishes.

Part (iii). Conditional on everything observed before a candidate harness evaluation, its $|Q_g|N_{\mathrm{rep}}$ errors $\{S(\tau_{q,r}) - \mu_q(W)\}_{q \in Q_g,\, r \leq N_{\mathrm{rep}}}$ are independent across tasks and replicates and σ-sub-Gaussian, so $\bar{S}(W) - \mu(W)$ is $\sigma/\sqrt{|Q_g|N_{\mathrm{rep}}}$-sub-Gaussian and

$$\Pr\big(|\bar{S}(W) - \mu(W)| > \xi\big) \;\leq\; 2\exp\left(-\frac{|Q_g|N_{\mathrm{rep}}\,\xi^2}{2\sigma^2}\right).$$

Over $T_{\max}$ rounds, each group evaluates W_0, M pool candidates per round, and at most $\mathbb{I}_{\mathrm{rep}} \in \{0, 1\}$ repaired candidate per round, giving at most $n_{\mathrm{eval}} = 1 + T_{\max}(M + \mathbb{I}_{\mathrm{rep}})$ harness evaluations per group and $n_{\mathrm{eval}}|\mathcal{G}|$ across all groups. Choosing $\xi = \sigma\sqrt{2\log(2n_{\mathrm{eval}}|\mathcal{G}|/\alpha)/(|Q_g|N_{\mathrm{rep}})}$ makes each tail probability at most $\alpha/(n_{\mathrm{eval}}|\mathcal{G}|)$. A union bound over the at most $n_{\mathrm{eval}}|\mathcal{G}|$ evaluations gives $|\bar{S} - \mu| \leq \xi$ for all of them with probability at least $1 - \alpha$; the union bound remains valid although later candidates and repairs depend adaptively on earlier noise, because each tail bound holds conditionally on the past. On this event,

$$\mu(W_T) \;\geq\; \bar{S}(W_T) - \xi \;=\; \bar{S}(W_0) + \sum_{t=1}^{T} \Delta_t - \xi \;\geq\; \mu(W_0) + \sum_{t=1}^{T} \Delta_t - 2\xi,$$

which is (16).

Part (iv). With per-group gating, group g advances in a round with probability a_g, so the expected fraction of advancing groups is $\frac{1}{|\mathcal{G}|}\sum_g a_g$. A single global gate advances only if every group's candidate passes, which under independence has probability $\prod_g a_g$. For equal $a_g = a$ the ratio is $a/a^{|\mathcal{G}|} = a^{1-|\mathcal{G}|}$ (e.g., 8× for $|\mathcal{G}| = 4$ and $a = 0.5$). □

A.2. Proof of Theorem 2 (Targeted Repair vs. Passive Discovery)

Proof. Part (i). The passive learner must observe every invariant, so $T_{\text{passive}} \geq \max_j G_j$, where $G_j \sim \text{Geom}(p_j)$, supported on $\{1, 2, \dots\}$, is the first round in which invariant j surfaces and the G_j are independent. Because $p_j \leq \epsilon < 1$, each G_j stochastically dominates $G'_j \sim \text{Geom}(\epsilon)$, and independence carries the dominance to the maxima. Couple $G'_j = \lceil E_j \rceil \geq E_j$ with $E_1, \dots, E_J$ i.i.d. exponential of rate $\nu = -\log(1-\epsilon)$; the coupling is exact since

$$\Pr\left(\lceil E_j \rceil > \ell\right) = e^{-\nu\ell} = (1-\epsilon)^{\ell}.$$

By the Rényi representation of exponential order statistics, $\max_j E_j \overset{d}{=} \sum_{\ell=1}^{J} X_\ell/(\ell\nu)$ with X_ℓ i.i.d. Exp(1), so for all $\epsilon \in (0,1)$,

$$\mathbb{E}[T_{\text{passive}}] \;\geq\; \mathbb{E}\Big[\max_j E_j\Big] \;=\; \frac{H_J}{-\log(1-\epsilon)} \;\geq\; \frac{1-\epsilon}{\epsilon}\, H_J,$$

using $-\log(1-\epsilon) \leq \epsilon/(1-\epsilon)$ on $(0,1)$ (giving $\Omega((\log J)/\epsilon)$ when $\epsilon \leq 1/2$). For example, $J = 10$ and $\epsilon = 0.05$ give $\mathbb{E}[T_{\text{passive}}] \geq 57.1$ rounds.

Part (ii). Let u_t be the number of invariants still unresolved after round t ($u_0 = J$), and let $\mathcal{F}_t$ be the history of the first t rounds. By (a) and (d), resolved invariants never displace unresolved ones in the ranking and stay resolved.

Phase 1 ($u_{t-1} \geq b_{\min}$). By (b) and (c), on the event $\{t_1 > t-1\} \in \mathcal{F}_{t-1}$ where $t_1 = \min\{t : u_t < b_{\min}\}$, at least $b_{\min}$ unresolved invariants are targeted and

$$\mathbb{E}\left[u_{t-1} - u_t \,\middle|\, \mathcal{F}_{t-1}\right] \;\geq\; \rho\, b_{\min}.$$

The stopping time t_1 is integrable because each active round resolves at least one invariant with probability at least ρ. Defining $Z_t = (J - u_t) - \rho\, b_{\min} t$, the stopped process $Z_{t\wedge t_1}$ is a bounded-increment submartingale; optional stopping yields

$$\rho\, b_{\min}\, \mathbb{E}[t_1] \;\leq\; J - \mathbb{E}[u_{t_1}] \;\leq\; J.$$

Phase 2 ($u_{t_1} < b_{\min}$). Conditional on $\mathcal{F}_{t_1}$, if $u_{t_1} = 0$ (in particular whenever $b_{\min} = 1$), zero further rounds are needed; if $1 \leq u_{t_1} \leq b_{\min} - 1$, every remaining invariant is targeted in every subsequent round and resolved independently with probability at least ρ, so the conditional remaining time is stochastically dominated by the maximum of $u_{t_1} \leq b_{\min} - 1$ independent Geom(ρ) variables. With the coupling of Part (i) at rate $\nu' = -\log(1-\rho) \geq \rho$, this maximum is at most $1 + \max_j E'_j$, whose conditional expectation is at most $1 + H_{b_{\min}-1}/\nu' \leq 1 + H_{b_{\min}-1}/\rho$.

Taking unconditional expectations and adding the two phases gives (18), which does not involve ϵ. (Under all-or-nothing bundle admission with joint pass probability $\rho_{\text{bundle}} > 0$ per round, each accepted round resolves $\min(b_{\min}, u_{t-1})$ invariants, requiring at most $\lceil J/b_{\min} \rceil$ accepted bundles of mean waiting time $1/\rho_{\text{bundle}}$ each.) □

A.3. Proof of Proposition 1 (Strict Regression Dominance under Multi-Rollout Grading)

Proof. When round $t-1$ promoted, (4) gives $\Delta_{\text{gap},t}(c) = 1 - o_c$ and $\Delta_{\text{hurt},t}(c) = \max(0, \delta_c)$ for any numbers of rollouts graded under W_{t-2} and W_{t-1}, so $\text{Priority}_t(c) = (1-o_c) + \max(0, \delta_c)$. For a regressed criterion c ($\delta_c > 0$) and a non-regressed criterion c' ($\delta_{c'} \le 0$, so $\Delta_{\text{hurt},t}(c') = 0$),

$$\text{Priority}_t(c) - \text{Priority}_t(c') \;=\; \big(1 - o_c + \delta_c\big) - \big(1 - o_{c'}\big) \;=\; \delta_c - (o_c - o_{c'}).$$

Part (i). If $o_{c'} \ge o_c$, the right-hand side is at least $\delta_c > 0$.

Part (ii). The right-hand side is positive if and only if $o_c - o_{c'} < \delta_c$. In the single-rollout binary case ($K_{\text{gr}} = 1$), a fresh regression has $o_c = 0$ and $\delta_c = 1$, so $\text{Priority}_t(c) = 2$, whereas every non-regressed criterion has $\text{Priority}_t(c') = 1 - o_{c'} \le 1$.

Part (iii). In any round (promoted or rejected), both $\Delta_{\text{hurt},t}(c)$ and $\Delta_{\text{gap},t}(c) = 1 - \text{obey}(c; W_{t-1})$ in (4) are non-negative, so $\text{Priority}_t(c) = 0$ holds if and only if $\text{obey}(c; W_{t-1}) = 1$ and $\Delta_{\text{hurt},t}(c) = 0$. Because each verdict $r(\tau, c)$ is binary, $\text{obey}(c; W_{t-1}) = 1$ holds exactly when every trajectory in $\mathcal{T}_c(W_{t-1})$ obeys c; hence no criterion failed by W_{t-1} and no regressed criterion ($\Delta_{\text{hurt},t}(c) > 0$) is ever pinned. □

B. Harness Instantiation Used in Our Experiments

This appendix lists the concrete constants and baseline details for §2 and §4–§5.

- **Soft-HURT parameters, pool reward, and ledger window:** Single-task noise floor $\delta_{\text{noise}} = 15$ points on $[0, 100]$ (fixed across all suites without per-domain tuning); default trade-off budget $h_{\max} = 1$ hurt task and margin $\mu_{\min} = 1.5$ points (on the professional suites, $h_{\max} = \text{round}(0.18\,|Q_g|)$, $\mu_{\min} = 3.0$, and at least two helped tasks; on multi-stage continual tool/skill and drift suites, unrepaired candidates use conservative threshold $h_{\max} = 0$ and trigger in-round `[iff]` repair up to $h_{\text{rep}} \in \{1, 2, 3\}$ with re-executed Soft-HURT evaluation at $h_{\max} = 1, \mu_{\min} = 0$); active deduplicated rejection-ledger read-back window $L_{\max} = 12$ entries per group on hosted EvoHarnessBench, $L_{\max} = 40$ in the reference candidate-level runner (Drift Gauntlet, local EvoHarnessBench, and PAST-Bench), and all active deduplicated criterion-removal entries in the professional-suite runner; $\text{Perf}_m = \bar{S}^{\text{strict}}(\widetilde{W}_t^{(m)}) + \frac{1}{4}\bar{S}(\widetilde{W}_t^{(m)})$ (estimated by the critic under critic scoring), $\beta = 3.5$, $\gamma = 1.5$, stabilizer 10^{-6}, contrast threshold $\delta_{\text{con}} = 10$ points; pool size $M = 2$ across all main-text, local-gym, and Fin-Rate & Ledger runs and $M = 4$ on PRBench-Legal & Harvey LAB (reaching 66.67% Official; an $M = 2$ prior-free run reaches 65.76%, and an exploratory $M = 4$ hosted run reaches 76.42%), initialized from s_0 across all M parallel candidates in Round 1 and evolved via m^+ vs. m^- ($m^- \ne m^+$) in rounds $t \ge 2$ (Appendix B.4).
- **Suite-level Soft-HURT validation buffers:** Hosted and local EvoHarnessBench gate candidates on cumulative visible stage tasks (5/5 stage promotions admitted cleanly under Soft-HURT: 1 hosted, 4 local); at each new stage, the local runner also re-checks the carried incumbent against W_0 with the same Soft-HURT rule on the expanded training split and returns to W_0 if it fails, which happened once (`HR`, Stage `v2`). Drift Gauntlet gates against a 24-task stratified validation buffer combining current-block UED frontier tasks with up to 10 Vitter Algorithm R reservoir replay samples sharing the pre-execution observable tool/schema signature `observable_domain_sig` (promoting 9/12 blocks, including in-round `[iff]` repair on `ScienceWorld` under *domain drift*, and rejecting 3/12 blocks where hurt equaled lift). PAST-Bench applies Soft-HURT ($\delta_{\text{noise}} = 15, h_{\max} = 1, \mu_{\min} = 1.5$) to per-task scores from the joint critic pass (24/24 round winners admitted across 12 families; §5). The two professional suites gate strictly on `TRAIN` pre-execution checklist self-reviews ($\delta_{\text{noise}} = 15$, in-round `[iff]` repair; PRBench-Legal gates $m^\star = \arg\max_m \hat{A}_m$ as in Algorithm 1, while Fin-Rate & Ledger evaluates across the candidate pool), while post-hoc `DEV` uses the sealed official

rubric (Table C.2).

- **Rounds and replicates:** Although the reference runner supports early stopping ($<$ 0.5-point macro gain with $<$ 2 net-helped tasks, no group promoted, or $T_{\max} = 5$), all reported experiments fix the round count: 1 per teach/update episode on PAST-BENCH; 1 per stage `adapt`(v_k) on EVOHARNESSBENCH (`v1`–`v3`, and `v1`–`v4` on `ITSM`); 1 per online task block $k \in \{1, \ldots, K\}$ on DRIFT GAUNTLET ($K \in \{2, 3, 4\}$); and 3 adaptation rounds (`R1`–`R3` from baseline `R0` = W_0) on the two professional suites. Main-text and local-gym suites use $N_{\text{rep}} = 1$ and the two professional suites use $N_{\text{rep}} = 2$ evaluation replicates per task.
- **Benchmark rollout notes:** Under unified `Gemini 4` (8-step agent budget/task on hosted EVOHARNESSBENCH), Layer 1 verifies a final answer and required state mutations while Layer 2 grades task-derived criteria on cumulative stage tasks (12, 16, and 18 at `v1`–`v3`), and an identical single-re-execution pass re-ran relay/tunnel infrastructure outages across *every* arm under the unchanged harness state before aggregating Table 2, leaving 0 errored records in all six final 43-task rows (6 infra-errored records occurred only in intermediate cell (Cohort `v1`, H_{v_2}), which enters no reported column). All five hosted `GUARD` configurations evaluated in our study (the reported M=2 main arm at 80.80% in Table 2, its original-distiller control at 75.57%, the initial one-layer run at 75.98% in Table 4, a fixed-persona M=2 variant at 76.69%, and an M=4 operator pool at 76.42%) beat the top external baseline (74.52%, META-HARNESS). On DRIFT GAUNTLET [5] (single-model `Gemini 4`, seed 42), every task exposes a pre-execution observable environment/schema signature (`bfcl`, `sciworld`, or `bird:<db_id>`) without any hidden block index or group label; each signature's section is bounded to 3,800 characters and placed at both the start and end of the system prompt. Because the three BIRD groups use different databases and *domain drift* spans three environments, signatures separate the groups of *topic* and *domain* drift, whereas `BFCL-Miss` and `BFCL-Base` share `bfcl` and test within-signature opposing shifts via the 10-sample replay buffer. Held-out evaluation draws one rollout per task (solver temperature 0.2) and memoizes it by task and rendered playbook for every arm, avoiding the artificial −0.59pp apparent forgetting induced when a constant W_0 is repeatedly re-sampled at temperature 0.2 (3-pass W_0 mean 69.86% ± 0.18pp). In a symmetric 2×3 routing control across all four streams, equipping GEPA and META-HARNESS with the same pre-execution signature routing lowers their macro means to 71.55% (F = −2.78pp, with −13.33pp forgetting on *workflow drift*) and 71.08%, whereas running `GUARD` with a single global playbook (no signature routing) still reaches 72.35% (Δ_K = +2.44pp, $\bar{\Delta}$ = +2.25pp, F = 0.00pp). This shows that `Analyze`, `Propose`, and Soft-HURT account for 74% of `GUARD`'s total gain over W_0. In the lightweight local five-domain gym (Table C.5), the semantic reviewer grades structured execution, entity-grounding, policy-completeness, and outcome-verification predicates directly with deterministic trace fallbacks.

B.1. Comparison with Prior RSI Frameworks and Environments

Table B.1 summarizes representative recursive self-improvement (RSI) methods and environments discussed in §3 across their optimization targets, closed-loop stability mechanisms, and adaptation and active-context scalability.

B.2. Baselines and Procedural Breakdown Across the Three Primary Benchmarks

No arm reads held-out sessions, test splits, drift boundaries, or grader files (per-arm configurations and hosted playbook sizes in Table B.2; procedural PAST-BENCH breakdown in Table B.3). On PAST-BENCH (all persistent arms on HERMES+ with `gemini-3.8-flash` solver and `claude-opus-5-5` optimizer/judge), RRSI [45] reflects on visible trajectories and rewrites memory and skill files after each teach or update episode (omitting its validation-split score-band gate because families expose

Table B.1 | **Optimization targets, closed-loop mechanisms, and scalability of recursive self-improvement (RSI) frameworks** (taxonomy adapted from RSIGym [26]). Checkmarks (✓), crosses (✗), and ✓ mark full presence, absence, or partial support; † marks baselines evaluated here. **Scalability** pairs adaptation cost (T rounds, Q tasks, M or K candidates) with active context/archive growth in T (§4.3, Appendix B.5).

Framework / Environment	Optimization Targets			Closed-Loop Features			Scalability	System Design & Update Mechanism
	Weight Updates	Harness & Skills	Text Memory	Regret Ranking	Annealed Budget	Non-Regr. Gate	(Work / Depth; T-Ctx.)	
Reflexion / ExpeL / AWM [38, 59, 43]	✗	✗	✓	✗	✗	✗	$O(TQ)$ / $O(TQ)$; $O(T)$ mem.	Episodic verbal buffers, insight extraction, & workflow memory
Voyager / MemGPT / WikiSkill [42, 31, 40]	✗	✓	✓	✗	✗	✗	$O(TQ)$ / $O(TQ)$; $O(T)$ lib.	Skill libraries & append-only wikis; skill edits gated on validation score
Dynamic Cheatsheet / ACE / ReasoningBank [39, 56, 30]	✗	✗	✓	✗	✗	✗	$O(TQ)$ / $O(TQ)$; $O(T)$ ctx.	Test-time cheatsheets, incremental context updates, & trace memory
DSPy / TextGrad / EvoPrompt [18, 53, 11]	✗	✗	✓	✗	✗	✗	$O(TKQ)$ / $O(TQ)$; $O(1)$ ctx.	Scalar-objective, textual-gradient, or evolutionary prompt search
TextReg [10]	✗	✗	✓	✗	✗	✗	$O(TQ)$ / $O(TQ)$; $O(T)$ bank	Regularized text updates filtering low-recurrence patches & length (requires GT)
GEPA†[f] [1]	✗	✗	✓	✗	✗	✗	$O(TKQ)$ / $O(TQ)$; $O(TKQ)$ front.	Reflective prompt evolution with per-instance Pareto archive (no `[iff]`)
Hermes+ (PAST-Bench)† [50]	✗	✓	✓	✗	✗	✗	$O(TQ)$ / $O(TQ)$; $O(T)$ mem.	Cross-session personal-agent memory & skill persistence
MemToolAgent (EvoHarnessBench)† [17]	✗	✓	✓	✗	✗	✗	$O(TQ)$ / $O(TQ)$; $O(T)$ skills	Continual skill & tool-memory accumulation across evolving stages
MLGym / PostTrainBench[a] [29, 34]	✓	✗	✗	✗	✗	✗	GPU train; $O(T)$	Containerized ML workspaces & autonomous post-training
Absolute Zero / R-Zero / Agent0[d] [60, 13, 46]	✓	✗	✗	✓	✗	✗	RL grads; $O(T)$	Zero-data RL self-play via code-executor or majority-vote rewards
Context Language Models (CLM) [36]	✓	✗	✓	✗	✗	✗	RL grads; $O(1)$ ctx.	RL-trained intra-episode context-file editing with efficiency gating
ADAS / AFlow / DGM[e] [12, 55, 58]	✗	✓	✗	✗	✗	✗	$O(TKQ)$ / $O(TQ)$ GT; $O(T)$	Archive / MCTS agent-code evolution scored on validation ground truth
AutoHarness[b] / Self-Harness [22, 57]	✗	✓	✗	✗	✗	✓	$O(TKQ)$ / $O(TQ)$; $O(T)$ tests	Code-harness synthesis; regression-test gating in Self-Harness
Meta-Harness† [20]	✗	✓	✓	✗	✗	✗	$O(TKQ)$ / $O(TQ)$; $O(TQH)$ log	Agentic search over harness code & memory using prior traces
RSIBench-Data / RSIGym[c] [25, 26]	✓	✓	✗	✗	✗	✗	Shared-$ budget	Everything-as-a-Service (EaaS) under a shared dollar budget
RRSI† [45]	✗	✓	✓	✗	✓	✗	$O(TQ)$ / $O(TQ)$; $O(1)$ ctx.	Annealed edit budget, critic & pruner, score-band gate (within δ_{band} of the best score)
Harness^2 [49]	✗	✓	✓	✗	✗	✗	$O(KQ)$ probe; $O(Q)$ lib.	Per-task online Propose–Probe–Compose from W_0 + continual improver library
GUARD (Ours)[g]	✗	✓	✓	✓	✗	✓	$O(TMQ)$ / $O(TQ)$; $O(1)$ **ctx.**	**Non-parametric regret-guided RSI: $\Delta_{\text{hurt},t}+\Delta_{\text{gap},t}$ priority, `[iff]` guards, m^+/m^- pool, Soft-HURT gate**

[a]Optimizes model weights/hyperparameters under fixed agent code. [b]Synthesizes action-blocking code harnesses around frozen weights; Self-Harness [57] gates on regression tests. [c]Exposes training, rollout, evaluation, and sandboxes under a shared budget. [d]Also Self-Challenging Agents [62]; curriculum RL targets solver uncertainty or code-verified tasks instead of per-criterion Δ_{hurt} across playbooks. [e]Mutates Python agent workflows or coding-agent repositories (SWE-bench/Polyglot in DGM) against ground-truth validation scores. [f]GEPA maintains an instance-Pareto archive without an incumbent non-regression promotion gate. [g]`GUARD` updates zero model weights (though its abstract interface accommodates adapter/data engines); every reported run uses reference-free checklist self-review, paired criterion regret ranking $\Delta_{\text{hurt},t} + \Delta_{\text{gap},t}$, fixed un-annealed budgets, and Soft-HURT gating.

single-trajectory sessions without a validation split), GEPA [1] evolves instructions via $K = 2$ reflective candidates with per-episode Pareto tracking, and Meta-Harness [20] proposes $K = 2$ candidate rewrites of both files from accumulated trajectories and keeps the higher self-scored state. To confirm that `GUARD`'s gains on PAST-Bench do not rely on a stronger teacher optimizer (`claude-opus-5-5` distilling for a `gemini-3.8-flash` solver), we also executed two complete 12-family, six-arm *model-identical* suites in which solver, distiller, critic, and grader all use the exact same model: (1) under

pure `claude-opus-5-5`, GUARD (M=2 prior-free) ranks first on overall Task (0.870 vs. 0.866 RRSI, 0.861 HERMES+, 0.861 META-HARNESS, 0.858 GEPA, and 0.347 W_0; +0.523 lift over W_0), Recall (0.871), and Evolve (0.458 vs. 0.431–0.454), and achieves the sole lowest memory pollution among all persistent arms (0.167 vs. 0.208 on HERMES+/GEPA and 0.250 on RRSI/META-HARNESS; 0.000 procedural pollution); and (2) under pure `gemini-3.8-flash`, GUARD lifts Task from 0.371 (W_0) to 0.852 (with Mechanism 0.840 vs. 0.829–0.835 on META-HARNESS/GEPA/RRSI), ties for the lowest memory pollution (0.167 vs. 0.194 on RRSI; 0.000 procedural pollution), and leads multi-session latent rule induction (PC03: 0.748 vs. 0.653 RRSI, 0.645 META-HARNESS, and 0.510 GEPA; +9.5 to +23.8pp). On hosted EVOHARNESSBENCH and DRIFT GAUNTLET (both under single-model `Gemini 4` for every role), RRSI performs single-candidate recursive self-critic diagnosis (`Failure Mode Diagnosis`) and rewrites the harness at each stage or continual block without paired priority ranking or Soft-HURT gating, while GEPA (K = 2, 4,200-char prompt cap on DRIFT GAUNTLET) and META-HARNESS (K = 2) evaluate candidates on each stage/block's visible adaptation tasks.

Table B.2 | **Per-arm adaptation configuration across the unified six-arm suite.** All evolving arms share identical solver and optimizer/reviewer models within each benchmark (§5). Final hosted chars report the Stage-v3 playbook/memory length on hosted EVOHARNESSBENCH.

Arm	Cands./Rd.	Adapt. Rollouts	Selection & Gate Rule	Adaptation Feedback Signal	Hosted Chars
`No-Evolve` (W_0)	0	0	Static base harness	None (zero-shot ReAct / base harness)	0
HERMES+ / MEMTOOLAGENT	1	1× online	Unconditional append/edit	Agent-initiated memory & skill writes	6,989
RRSI [45]	1	1× online / stage	Unconditional rewrite (no val. split)	Self-critic `Failure Mode Diagnosis` + pruner	10,013
GEPA [1]	K=2	2× (exec.) / 1× (critic)	Per-instance Pareto frontier argmax	Reflective trace critique on non-dominated tasks	5,123
META-HARNESS [20]	K=2	2× (exec.) / 1× (critic)	Best self-scored candidate (ungated)	Proposer–judge search over raw trace archive	4,538
GUARD (Ours)	**M=2**	**2× (exec.) / 1× (critic)**	**arg max $\hat{A}_m$ + Soft-HURT gate**	**2-layer checklist self-review (exec. or critic), $\Delta_{\text{hurt}}+\Delta_{\text{gap}}$, pinning, m^+/m^- pool**	**2,315**

Table B.3 | **PAST-BENCH procedural SOP sub-table across the three multi-step procedural families (PC01, PC03, PC04; Figure 3(a)).** Bold marks the best persistent arm per column.

Arm	Proc. Task	Proc. Mech.	Proc. Upd.	Proc. Poll.↓
`No-Evolve` (W_0)	0.427	0.443	0.214	0.000
HERMES+	0.808	0.780	0.529	**0.000**
RRSI [45]	0.714	0.508	0.841	1.000
GEPA [1]	0.815	0.761	0.464	**0.000**
META-HARNESS [20]	0.677	0.516	0.824	1.000
GUARD (Ours)	**0.830**	**0.900**	**0.901**	**0.000**

B.3. Reproducibility, Hyperparameters, and Prompt Templates

Model identifiers, decoding parameters, and budgets. Hosted EVOHARNESSBENCH (8-step rollout cap, B_W = 2,600 characters), DRIFT GAUNTLET (B_W = 3,800 characters per observable interface signature), and the controlled dual-axis `ITSM Propose`-stage ablation (25-turn cap, B_W = 4,201 characters, 32,768 solver and 65,536 distiller/reviewer output tokens; Table C.8) all use our

Table B.4 | **Paired per-family differences on PAST-BENCH, GUARD minus each persistent baseline (12 families; better/tied/worse counts; mean differences computed from unrounded per-family scores before rounding to three decimals).** GUARD achieves positive mean gains on Mech., Upd., and W.Prec. against every baseline with zero W.Prec. or Poll. family losses; exact two-sided sign tests reach unadjusted $p < 0.05$ on Mech. ($p = 0.016$) and W.Prec. ($p = 0.031$) vs. RRSI, with none surviving Holm correction across all 18 comparisons.

Comparison	Task	Mech.	Upd.	W.Prec.	Poll.↓
vs. HERMES+	+0.002 (4/5/3)	+0.037 (5/5/2)	+0.133 (6/5/1)	+0.052 (1/11/0)	+0.000 (0/12/0)
vs. RRSI	+0.045 (6/5/1)	+0.153 (7/5/0)	+0.023 (6/5/1)	+0.339 (6/6/0)	−0.333 (4/8/0)
vs. GEPA	−0.004 (1/6/5)	+0.041 (4/6/2)	+0.148 (4/6/2)	+0.057 (3/9/0)	+0.000 (0/12/0)
vs. META-HARNESS	+0.042 (5/5/2)	+0.139 (6/5/1)	+0.021 (5/5/2)	+0.297 (5/7/0)	−0.333 (4/8/0)

primary single-model `Gemini 4` configuration for every role. To evaluate GUARD as a drop-in stabilizer for teacher–student knowledge-distillation harnesses (§5), PAST-BENCH, the five-domain local EVOHARNESSBENCH gym (15-turn cap), and the two professional suites pair a `gemini-3.8-flash` student solver with `claude-opus-5-5` on Vertex AI as reviewer, distiller, and ContrastEvolve model (2,048–4,096 output tokens for pre-execution checklist derivation, review, and ContrastEvolve; 4,096–8,192 for distillation and repair). No temperature reaches the paired teacher–student endpoints or the hosted EVOHARNESSBENCH relay, whereas the DRIFT GAUNTLET runner forwards explicit temperatures to `Gemini 4` (solver rollouts 0.2; pre-execution checklist derivation and Layer-2 grading 0.0; GUARD distiller candidates 0.65, with 0.8–0.9 on empty or duplicate retries; in-round `[iff]` repair 0.25; ContrastEvolve 0.2; baseline proposers 0.2–0.5). B_W is 2,500 characters in the local gym, 2,400 per skill body (1,900 per memory entry) on PAST-BENCH, and 4,200 characters of playbook body excluding the `Output Contract` (5,600 after an `[iff]` repair) in the professional suites. Boxes B.1–B.2 summarize the remaining prompt templates (complementing s_0 and ContrastEvolve output in Boxes B.3–B.4); we plan to release reference harness code, pre-execution checklist self-reviewers, and prompt templates, subject to approval.

Box B.1: Pre-execution checklist derivation, Layer-2 semantic grader, and PAST-Bench critic templates

(a) Pre-execution atomic checklist derivation (hosted EvoHarnessBench, Drift Gauntlet, and professional suites; cached once before execution from public inputs only):
You are a deterministic rubric architect. Given ONLY the user prompt, visible tool declarations, and public schema (no trajectories or verifiers), decompose the request into concrete atomic binary sub-requirements (`req_1`, `req_2`, ...) covering clean execution, valid schema identifiers, idempotent mutations, and explicit prompt constraints. Return JSON `[{"id": "req_k", "desc": "..."}]`.

(b) Layer-2 per-criterion semantic self-reviewer (an unparseable verdict counts as failed or falls back to deterministic trace checks):
You are a strict, reference-free Self-Reviewer with no access to hidden verifiers. Using ONLY `user_prompt`, `schema_summary`, `tool_trace` / sandbox row deltas, and `final_response`, grade each pre-registered criterion in `frozen_sub_requirements`. Reply strictly with JSON `{"passed": true|false, "why": "<evidence from trace and output>"}`.

(c) PAST-Bench four-criterion critic (one joint pass over the incumbent and $M = 2$ candidates):
Score the incumbent and each candidate memory/skill state on $[0, 100]$ across four criteria: (1) grounding in session history without hallucinated preferences, (2) exact format/SOP step fidelity, (3) explicit scope/recipient `[iff]` boundaries avoiding episodic pollution, and (4) complete replacement of superseded rules. Return per-family strict and continuous scores plus helped/hurt indicators.

Box B.2: Distill **request wrapper and** ContrastEvolve **meta-prompt**

(a) Distill$(s_{t,m}, \mathcal{B}_{t,g}, W_{t-1,g})$ **request wrapper** ($\leq B_W$ chars, active ledger window $L_{\max}$):
`### DISTILLER STRATEGY PROMPT: <`$s_{t,m}$`> | ### INCUMBENT PLAYBOOK: <`$W_{t-1,g}$`>`
`### REGRET & CONTRAST BUNDLE:` (1) `wrap_in_iff_guard` ($\Delta_{\text{hurt},t} > 0$), (2) `add_method_and_check` ($\Delta_{\text{hurt},t} = 0, \Delta_{\text{gap},t} > 0$), (3) `pin_mastered` ($\text{Priority}_t = 0$, pinned invariants), (4) grounded (τ^+, τ^-) trace contrasts, and (5) active deduplicated `rejection_ledger` entries (do not re-emit verbatim).

(b) ContrastEvolve **meta-prompt** (all task/episode IDs scrubbed to `<TASK_ID>`):
Given the strategy prompt, playbook excerpt, and scrubbed metrics $(R_m, |Q_m^+|, |Q_m^-|)$ of the BEST candidate (m^+) and WORST candidate (m^-) in round t, extract causal `learnings` explaining why m^+ outperformed m^- and rewrite both prompts into `{"learnings": [...], "updated_best_prompt": "...", "updated_worst_prompt": "..."}` so candidate 1 amplifies m^+'s scoped structure while preserving pinned rules and candidate 2 repairs m^-'s over-broad triggers into strict `[iff <precondition>]` guards.

B.4. Shared Seed Prompt Example

In Round 1, every candidate in the pool receives the same strategy prompt s_0 (`SHARED_SEED_PROMPT`, Eq. (6)), shown in Box B.3. s_0 is written once, before any task is seen, and names no benchmark, domain, tool, or task; in its wording, the *champion* is the incumbent playbook W_{t-1}. We reproduce the prompts and the evolver output exactly, setting Markdown code spans in `typewriter` type and Markdown lists as lists.

Box B.3: Shared seed prompt s_0

Diagnose the champion's dominant behavioral failure mode from its execution traces and failure report before designing anything. Then propose exactly ONE harness mechanism (prefer a skill: a SKILL.md workflow plus at most one small script) that directly removes that failure mode. Keep the mechanism narrow and guarded with explicit `[iff <precondition>]` scope so it cannot fire on tasks the champion already solves. Validate it on a small fix slice plus a guard slice before emitting it.

One evolution step. Box B.4 shows the output of the contrastive evolver of §4.3 after the first round of one of our runs ($M = 2$). In that run, each candidate is generated by a separate distiller call in which s_0 fills the strategy-prompt field, while the incumbent playbook, the regret bundle, and the rest of the request are shared; the two Round-1 requests were therefore byte-identical (the runner asserts this), and the two candidates differed only through sampling. Under execution scoring on the round's four source tasks, the incumbent averaged $\bar{S} = 75.00$ points. Candidate m_1 averaged $\bar{S}_1 = 93.75$ ($\bar{S}_1^{\text{strict}} = 84.375$) with one helped and no hurt tasks (pool reward $R_1 = (84.375 + \frac{1}{4} \cdot 93.75) - 3.5(0) + 1.5(1) = 109.31$, Eq. (8)), and candidate m_2 averaged $\bar{S}_2 = 68.75$ ($\bar{S}_2^{\text{strict}} = 46.875$) with one helped and three hurt tasks ($R_2 = (46.875 + \frac{1}{4} \cdot 68.75) - 3.5(3) + 1.5(1) = 55.06$). The gate promoted m_1 under Soft-HURT with zero hurt tasks (net gain 18.75 points), and the evolver contrasted $m^+ = m_1$ with $m^- = m_2$. Its two rewritten prompts became the candidate strategy prompts of Round 2. Both keep the structure of s_0 (diagnose, propose one mechanism, validate on a fix slice and a guard slice): s_2^+ keeps the opening instruction of s_0 verbatim and adds the scoping and content rules that separated m_1 from m_2, while s_2^- narrows the edit to the single most frequently failing tool call and lists what to stop doing.

Box B.4: Contrastive evolver output after Round 1

learnings:

- Trigger scope decides the outcome. The winner attached every rule to a concrete tool call or a literal request phrase, such as 'calling create_new_hr_case' or 'creating it if necessary'. The loser opened with a broad 'request names people, groups, services…' trigger and an unconditional '[always, at final response]' rule. Those fire on nearly every task and hurt tasks the champion already solved.
- Unconditional end-of-task rules that tell the agent to take more actions are dangerous. 'If a required write is still pending, perform it before replying' can cause spurious or duplicate writes. Final-response guidance should only interpret results, never trigger new actions.
- Precision in field formats matters. The winner said 'numeric service id, not the name' and 'leave the default unless the request states a status'. The loser said 'resolved service id' and invented a draft-or-work_in_progress rule with extra synonyms. Vague wording leaves the agent guessing, and invented defaults override correct behavior.
- Complete the data chain each rule needs. The winner told the agent to resolve each matched profile to a user's first and last name before writing, because the write field needs a name. The loser listed the matches but never said how to get the field value.
- Order conditional lookups explicitly: look up the service first, and fetch topic details only if it is not found. Ambiguous order invites wasted or duplicate calls.
- Error-recovery wording should be narrow ('fix only the named field's format and retry'), not generic.
- Notification rules should name the payload (affected case numbers and outcome) and require that the send call actually succeeded.

updated_best_prompt (s_2^+, candidate 1 in Round 2):
Diagnose the champion's dominant behavioral failure mode from its execution traces and failure report before designing anything. Then propose exactly ONE harness mechanism (prefer a skill: a SKILL.md workflow plus at most one small script) that directly removes that failure mode.

Scope rules:
– Gate every rule with an explicit `[iff <precondition>]`.
– Each precondition must be either a specific tool call about to be made or a literal phrase in the request.
– Never use broad triggers like 'request mentions entities'.
– Never write `[always]` rules.
– Put grounding steps inside the write-tool trigger they serve.

Content rules:
– For every field the failure touches, state the exact expected format (for example, display name vs numeric id) and where to obtain that value.
– Do not invent defaults. Say 'leave the default unless the request states otherwise'.
– Spell out complete data chains:
  – lookup, then resolve to the exact value the write field needs;
  – one write per item;
  – never stop early.
– Order conditional calls explicitly: check existence first, and fetch prerequisites only if the item is missing. Never create duplicates.
– Keep error recovery narrow: fix only the field named in the error, then retry.
– Success-interpretation rules must only define what counts as success or failure. They must never instruct the agent to take new actions at the end of a task.
– For notifications, name the payload and require that the send call actually succeeded before the agent claims it.

Validate on a small fix slice plus a guard slice of tasks the champion already solves. If any guard task changes behavior, narrow the trigger until it doesn't, then emit.

updated_worst_prompt (s_2^-, candidate 2 in Round 2):
Do not write a broad playbook. Diagnose the champion's single most frequent failing tool call from its traces: the write call whose arguments most often fail or are wrong. Build ONE mechanism that fixes only that call: a SKILL.md plus at most one small validation script that checks the call's arguments before submission.

Requirements:

– The skill fires only `[iff about to call <that exact tool>]`. No global grounding rule, no `[always]` rule, and no final-response rule that triggers extra actions or 'pending' writes.
– For each argument of that tool, list:
  – the exact format (display name vs numeric id, casing, enum values);
  – the exact lookup tool and returned field that supplies it.
– If an argument requires a value derived from another record (for example, a person's name from a profile), state the resolution step explicitly.
– Do not invent defaults or synonym mappings. Preserve the tool's default unless the request explicitly specifies a value.
– On a validation error, change only the named field and retry.

Stop doing:

– vague phrases like 'resolved id';
– unordered lookup lists;
– catch-all summaries.

Validate on a fix slice where that tool failed, plus a guard slice of solved tasks. Emit only if the guard slice is unchanged.

B.5. Time and Space Complexity Analysis

This subsection derives the asymptotic time and space complexity of GUARD across adaptation and deployment, comparing it against the unified baseline suite (No-Evolve W_0, HERMES+ / MEMTOOLAGENT, RRSI, GEPA, and META-HARNESS; Table B.5). Let $|\mathcal{G}|$ be the number of independent groups, $Q = \sum_{g\in\mathcal{G}} |Q_g|$ the total source tasks across groups, $T \le T_{\max}$ the adaptation rounds, $M \ge 2$ the candidate pool size (K for multi-candidate baselines, with $M = K = 2$ in matched runs), $C = \max_q |C(q)|$ the pre-execution checklist size (4–8 predicates, so $C = O(1)$), H the per-task rollout step horizon, B_W the bounded playbook size (2.3k–4.2k characters, $B_W = O(1)$ in T), $L_{\max}$ the bounded active rejection-ledger read-back window ($\le$ 12–40 entries per group, $L_{\max} = O(1)$ in T), and $B_{\mathcal{B}}$ the regret-bundle payload size passed to Distill (our runners cap rendered (τ^+, τ^-) contrasts to ≤ 6 pairs and clip failure rationales to a fixed prompt budget so $B_{\mathcal{B}} = O(1)$ in T; without capping, $B_{\mathcal{B}} = O(|Q_g|(C + H))$ under execution scoring, and critic scoring grows as $O(t|Q_g|H)$ if accumulated trajectories are unwindowed).

Table B.5 | **Asymptotic adaptation, proposer-prefill, and deployment complexity across the unified six-arm suite.** Here T is adaptation rounds, Q source tasks ($|\mathcal{G}|$ groups), M (or K) candidates per round (M=K=2 in matched runs), H rollout step horizon, and $B_W, L_{\max}, B_{\mathcal{B}} = O(1)$ in T the bounded playbook, active ledger-window, and capped regret-bundle sizes.

Method	Adaptation Rollouts (Total)	Sequential Rollout Depth	Proposer Prefill Context (Round t)	Test-Time Rollouts & Active Context
No-Evolve (W_0)	0	0	—	1× rollout; $O(1)$ ctx.
HERMES+ / MEMTOOLAGENT	$O(TQ)$	$O(TQ)$	$O(t)$ accumulated mem./skills	1× rollout; $O(T)$ ctx.
RRSI (M=1)	$O(TQ)$	$O(TQ)$	$O(B_W) = O(1)$	1× rollout; $O(B_W) = O(1)$ ctx.
GEPA (K cands.)	$O(TKQ)$	$O(TQ)$	$O(B_W) + O(tKQ)$ Pareto	1× rollout; $O(B_W) = O(1)$ ctx.
META-HARNESS (K cands.)	$O(TKQ)$	$O(TQ)$	$O(tQH)$ raw trace archive	1× rollout; $O(B_W) = O(1)$ ctx.
GUARD (Execution scoring)	$\boldsymbol{O(TMQ)}$ (M=K=2)	$\boldsymbol{O(TQ)}$ (M-way par.)	$\boldsymbol{O(B_W + L_{\max} + B_{\mathcal{B}}) = O(1)}$	**1× rollout; $O(B_W) = O(1)$ ctx.**
GUARD (Critic scoring)	$\boldsymbol{O(TQ)}$ (1× online)	$\boldsymbol{O(TQ)}$	$\boldsymbol{O(B_W + L_{\max} + B_{\mathcal{B}})}^{\star}$	**1× rollout; $O(B_W) = O(1)$ ctx.**

$^{\star}O(1)$ in T when the episode window is capped; $O(t|Q_g|H)$ when critic scoring inspects all unwindowed episodes.

Stage-by-stage adaptation time complexity. Decomposing Algorithm 1 across its stages isolates LLM rollout calls from local CPU bookkeeping:

1. *Pre-execution checklist derivation (§4.1):* Deriving the fixed atomic checklist $C(q)$ requires a single pre-execution model pass per source task once at $t = 0$ ($O(Q)$ single-pass calls total; benchmark field stripping in §5 is $O(Q)$ local CPU work), cached unchanged across all T rounds and M parallel candidates.
2. *Solver rollouts (§4.3):* Multi-turn environment rollouts dominate wall-clock adaptation cost. Under *execution scoring* on a fixed task set Q_g (Theorem 1), `GUARD` executes one baseline pass under W_0 plus M candidate passes per round, reusing recorded incumbent scores without re-running incumbents (plus at most $\mathbb{I}_{\text{rep}} \in \{0, 1\}$ conditional re-score when in-round `[iff]` repair fires on $1 \le |Q^-_{m^\star}| \le h_{\text{rep}}$, which adds at most one extra Q-task sweep in that round, i.e., worst-case per-round rollout depth $2Q$), totaling $(1 + T(M + \mathbb{I}_{\text{rep}}))Q = O(TMQ)$ rollouts ($O(TMQH)$ environment steps). Because the M candidates within round t depend only on $(W_{t-1,g}, \mathcal{B}_{t,g}, \mathcal{P}_{t,g})$, they execute in parallel across workers and across the $|\mathcal{G}|$ groups, yielding $O(TQ)$ sequential rollout depth. This matches single-candidate RRSI/Hermes+ in asymptotic wall-clock depth and matches K-candidate GEPA/Meta-Harness ($O(TKQ)$) in total work at $M = K = 2$. (On stage-expanded continual suites such as hosted EvoHarnessBench, where stage t adds q tasks under drifted tools and re-scores the incumbent and M candidates on the cumulative tq tasks, total rollouts scale as $\sum_{t=1}^{T}(M + 1)tq = O(T^2Mq)$ with $O(T^2q)$ sequential depth.) Under *critic scoring* on stateful cross-session benchmarks (PAST-Bench), `GUARD` executes $O(TQ)$ (1×) online episode rollouts, adding zero environment rollouts over Hermes+, RRSI, or Meta-Harness.
3. *Two-layer self-review (§4.1):* Layer 1 grades execution-verifiable predicates deterministically in the local sandbox ($O(H)$ CPU per rollout, zero model calls). Layer 2 grades semantic predicates one at a time in $O(C)$ localized single-pass judgments per rollout, adding $O(TMQC)$ unbatched single-pass calls under execution scoring ($O(TQC)$ under critic scoring, with $C = O(1)$) without multi-turn tool interaction.
4. *Three closed-loop stages (`Analyze`, `Propose`, `Gate`):* In `Analyze` (§4.2), computing obey$(c; W)$, $\Delta_{\text{hurt},t}(c)$, and $\Delta_{\text{gap},t}(c)$ and sorting QC criteria by $(\mathbb{I}[\Delta_{\text{hurt},t} > 0], \text{Priority}_t(c), \Delta_{\text{hurt},t}(c))$ takes $O(QC \log(QC))$ pure CPU time, and extracting same-task (τ^+, τ^-) contrasts takes $O(QM)$ CPU time (*zero* LLM calls). In `Propose` (§4.3), compiling M candidates via Distill and evolving (s^+_{t+1}, s^-_{t+1}) via ContrastEvolve costs $M + 1$ single-pass distiller calls per group and round ($O(T|\mathcal{G}|M)$ total), while within-group standardized fitness normalization costs $O(T|\mathcal{G}|M)$ CPU time; because each distiller call reads bounded $(W_{t-1,g}, \mathcal{B}_{t,g})$ and at most $L_{\max}$ active ledger entries, cumulative proposer prefill tokens over T rounds grow *linearly* in T as $O(T|\mathcal{G}|(B_W + L_{\max} + B_{\mathcal{B}}))$, whereas raw-trace archive search (Meta-Harness) grows per-round trace prefill as $O(tQH)$ and cumulative prefill quadratically as $\sum_{t=1}^{T} O(tQH) = O(T^2QH)$. In `Gate` (§4.4), evaluating Eq. (12) and updating $\mathcal{L}_g$ takes $O(Q)$ CPU arithmetic (*zero* LLM calls unless conditional `[iff]` repair triggers).

Deployment (test-time) complexity and empirical runtime. At deployment, $\mathcal{P}_{t,g}$, the two-layer self-reviewer, and $\mathcal{L}_g$ are discarded: each held-out task q_{test} runs **a single (1×) rollout** conditioned on $W_{T,g}$, incurring zero Best-of-M sampling or reviewer calls and $O(B_W) = O(1)$ prompt-prefix overhead independent of T. Empirically, compact `[iff]`-guarded rules shorten test-time trajectories instead of slowing them down: on hosted EvoHarnessBench under unified `Gemini 4` (43 test tasks; Table 2), `GUARD`'s final playbook (2,315 characters vs. 4,538 in Meta-Harness, 5,123 in GEPA, 6,989 in MemToolAgent, and 10,013 in RRSI), combined with two-layer gating (§6.4), reaches 0/43 eight-step budget exhaustions (vs. 6–10/43 on baselines), cutting inference tokens per task to 59,651 (−54.5% vs. W_0's 131,032, and −24.2% to −30.0% below evolving baselines) and keeping 43-task test evaluation wall-clock time (167.6 min, 233.8s/task) within 6% of W_0 (158.8 min) and

2.6× faster than RRSI (439.4 min).

Space and active-context complexity. Memory scales along three distinct tiers: (i) *active deployment prompt context* is $O(B_W) = O(1)$ in T (and progressive `wiki_read` loads only referenced sections on multi-file suites), avoiding the $O(T)$ prompt growth and 0.500 memory pollution of uncompacted buffers; (ii) *active proposer prefill context* is $O(B_W + L_{\max} + B_{\mathcal{B}})$ ($O(1)$ in T under capped contrasts/rationales, vs. $O(|Q_g|(C+H))$ uncapped), avoiding both the $O(TQH)$ raw-trace archive of META-HARNESS and the $O(TKQ)$ instance-Pareto score matrix of GEPA; and (iii) *host working RAM and disk* hold $O(QC)$ cached pre-execution predicates plus $O(|\mathcal{G}|MB_W)$ active pool strings in RAM, while the deduplicated on-disk rejection log is windowed to $L_{\max} = O(1)$ active entries whenever read back into model context.

C. Additional Experiments

This appendix evaluates `GUARD` across two public rubric-dense professional suites (PRBENCH-LEGAL & HARVEY LAB [2] and FIN-RATE & LEDGER [16, 27]; §C.2), reports the complete 24-row per-stream continual breakdown and per-group learning dynamics on DRIFT GAUNTLET [5] (§C.3), and concludes with the five-domain EVOHARNESSBENCH breakdown, matched `Analyze/Gate` leave-one-out ablation, and nine-row calibration audit (§C.4).

C.1. Setup and Active Mechanisms

Scoring, splits, and models. On the two rubric-dense professional suites, a judge grades each trajectory τ on a weighted rubric $\{c_j\}$ ($\sum_j w_j = 100$, score $v_j \in [0,1]$, critical flag $\kappa_j \in \{0,1\}$), giving *Uncapped Score* $V(\tau) = \sum_j w_j v_j$ and *Official Capped Score* (capped at 30 if any critical criterion fails). Both suites evolve over three adaptation rounds (`R1`–`R3` from baseline `R0` $= W_0$) on 24 `TRAIN` tasks and evaluate post-hoc on 16 frozen `DEV` tasks ($N_{\text{rep}} = 2$ single-rollout trials/task under the frozen round harness, with zero test-time candidate pooling or self-review), with all distillation, in-round `[iff]` repair, and Soft-HURT gating ($\delta_{\text{noise}} = 15$) operating strictly reference-free on `TRAIN` self-reviews in hermetic workspaces with progressive `wiki_read`. PRBENCH-LEGAL & HARVEY LAB [2] is a public benchmark spanning statutory/contractual verification, redline reconciliation, compliance triage, and financial-legal covenants across `Legal` and `Finance` (1 `Finance` and 15 `Legal` `DEV` tasks); FIN-RATE & LEDGER pairs public SEC 10-K financial analytics QA from FIN-RATE [16] (`fin_rate`) with corporate annual-report KPI extraction from LEDGER [27] (`ledger`), stratified across `Finance`, `Legal`, `HR`, and `StratOps` (3.125-point Official steps) under a unified 100-point weighted verification rubric kept sealed from all adaptation loops. Across both professional suites and the five-domain `Analyze/Gate` ablation, `gemini-3.8-flash` serves as agent and `claude-opus-5-5` (Vertex AI) serves as distiller, self-reviewer, critic, and rubric judge, with strict within-suite model parity across all arms (Tables C.1–C.2). Paired W_0 vs. W_{final} comparisons on the professional suites use exact two-sided sign-flip permutation tests, task-level bootstrap confidence intervals, and pooled within-task replicate standard deviation σ_w.

C.2. Public Professional Suites (PRBENCH-LEGAL and FIN-RATE & LEDGER)

Table C.3 and Figure C.1 show two main results under `TRAIN`-only reference-free gating. `GUARD` eliminates the late-round regressions of unguarded `RRSI`: on public PRBENCH-LEGAL & HARVEY LAB, `RRSI` drops 2.12pp in `R3` as statutory exception rules collide (63.30%, $\sigma_w = 7.11$) and `Meta-Harness` reaches 66.11% ($\sigma_w = 7.40$), whereas `GUARD` ($M = 4$ prior-free) reaches 66.67% Official (+7.19pp

Arm	Experience source and optimizer
`No-Evolve` (W_0)	Base agent executing without a learned playbook.
`RRSI` [45]	Multi-round failure-driven reflective playbook rewrite on `TRAIN` self-reviews.
`Meta-Harness` [20]	Counterfactual diagnostic harness search over full `TRAIN` execution histories.
`GUARD` (Ours)	Paired priority $\Delta_{\text{hurt}} + \Delta_{\text{gap}}$, rule pinning, `[iff]` guards, prior-free contrastive prompt evolution, and Soft-HURT gating (§4).

Table C.1 | **Arm taxonomy for the professional-suite experiments.** All arms within each benchmark share identical agent and optimizer/judge models (§C.1); main-text six-arm details appear in §5.

Experiment (rounds; splits)	Arm	`Analyze` stage	`Propose` stage	`Gate` rule	Reported in
PRBench-Legal & Harvey LAB and Fin-Rate & Ledger (3; 24 train, 16 frozen DEV×2)	`No-Evolve` (W_0) / `RRSI` / `Meta-Harness`	none / `TRAIN` self-review / trace history	– / reflective rewrite / harness search	– / rewrite / trace search	Tab. C.3, Fig. C.1
	`GUARD` (Ours)	paired $\Delta_{\text{hurt}} + \Delta_{\text{gap}}$, rule pinning, (τ^+, τ^-) contrasts	$M \in \{2, 4\}$ pool, `[iff]` guards, m^+/m^- evolution	`TRAIN` Soft-HURT	Tab. C.3, Fig. C.1
EvoHarnessBench 5-domain `Analyze`/`Gate` ablation (16 stages; 97 source, 236 test)	`No-Evolve` (W_0) / `w/o Analyze` / `w/o Gate`	– / 1 – obey (no pinning or contrasts) / full `Analyze`	– / M=2 pool (m^+/m^-) / M=2 pool (m^+/m^-)	– / Soft-HURT / unconditional	Tab. C.7
	`GUARD` (Ours)	paired $\Delta_{\text{hurt}} + \Delta_{\text{gap}}$, rule pinning, (τ^+, τ^-) contrasts	M=2 pool, `[iff]` guards, m^+/m^- evolution	Soft-HURT	Tab. C.7

Table C.2 | **Active closed-loop stages and reference-free gate rules across Appendix C.** Every arm adapts strictly without reference solutions, oracle skills, or held-out rubrics; `TRAIN Soft-HURT` gates on `TRAIN` pre-execution checklist self-reviews ($\delta_{\text{noise}} = 15$, in-round `[iff]` repair; §4.4).

Suite (24 `TRAIN` → 16 Frozen `DEV`)	Arm	R0	R1	R2	R3 (Final)	Δ Official vs. W_0 [95% CI]	Final Uncapped	Δ Uncapped	Hurt	Sub-domain Breakdown (Official)
PRBench-Legal & Harvey LAB (`gemini-3.8-flash` agent, `claude-opus-5-5` distiller/judge)	`No-Evolve` (W_0)	59.48%	59.48%	59.48%	59.48%	—	59.77%	—	—	Finance (n=1): 48.05% | Legal: 60.24%
	`RRSI`	59.48%	64.91%	65.42%	63.30%	+3.82pp [−0.21, +8.06]	63.59%	+3.82pp	**0**	Finance (n=1): 72.08% | Legal: 62.72%
	`Meta-Harness`	59.48%	63.81%	65.49%	66.11%	+6.63pp [+2.52, +11.04]	66.40%	+6.63pp	**0**	Finance (n=1): **76.62%** | Legal: 65.41%
	`GUARD` (Ours)	59.48%	**66.67%**	**66.67%**	**66.67%**	**+7.19pp** [+1.85, +12.85]*	**66.96%**	**+7.19pp**	**0**	Finance (n=1): **76.62%** | Legal: **66.01%**
Fin-Rate & Ledger Suite (`gemini-3.8-flash` agent, `claude-opus-5-5` distiller/judge)	`No-Evolve` (W_0)	37.50%	37.50%	37.50%	37.50%	—	70.58%	—	—	`fin_rate`: 12.50% | `ledger`: 62.50%
	`RRSI`	37.50%	**46.88%**	40.62%	40.62%	+3.12pp [−6.25, +12.50]	71.50%	+0.92pp	1	`fin_rate`: **18.75%** | `ledger`: 62.50%
	`Meta-Harness`	37.50%	37.50%	40.62%	40.62%	+3.12pp [+0.00, +9.38]	73.54%	+2.96pp	**0**	`fin_rate`: **18.75%** | `ledger`: 62.50%
	`GUARD` (Ours)	37.50%	43.75%	**50.00%**	**50.00%**	+12.50pp [+0.00, +31.25]	**82.96%**	**+12.38pp**	**0**	`fin_rate`: 12.50% | `ledger`: **87.50%**

Table C.3 | **Multi-round frozen `DEV` results (`R0` → `R3`) on the two professional suites** (24 `TRAIN` tasks, reference-free gating; 16 frozen `DEV` tasks × 2 trials). *Exact two-sided sign-flip $p = 0.0251$ ($M = 4$ prior-free, four helped and zero hurt tasks; $M = 2$ prior-free reaches 65.76%, +6.28pp); on Fin-Rate & Ledger ($M = 2$, $n = 16$), the 95% CI [+0.00, +31.25] touches 0 (two helped `ledger` tasks, exact sign-flip $p = 0.50$).

over W_0, 95% CI [+1.85, +12.85], exact sign-flip $p = 0.0251$; 66.96% Uncapped) with zero hurt tasks (4 helped / 0 hurt) and the lowest replicate variance among evolving arms ($\sigma_w = 5.56$). It lifts Tier-1 critical pass rates from 54.90% to 62.10% and Tier-2 from 63.85% to 69.38%, advancing both `Legal` (60.24% → 66.01%, $n = 15$) and the single `Finance` `DEV` task (48.05% → 76.62%, $n = 1$). On Fin-Rate & Ledger [16, 27], where `RRSI` drops 6.26pp after `R1` ($\sigma_w = 30.62$, 1 hurt task) and `Meta-Harness` stalls at 40.62%, `GUARD` reaches 50.00% Official (+12.50pp over W_0, 95% CI [+0.00, +31.25] touching 0 at $n = 16$, 0 hurt, $\sigma_w = 0.00$) and 82.96% Uncapped, lifting SEC 10-K `ledger` extraction from 62.50% to 87.50% (7/8 `DEV` tasks solved across both replicates while `fin_rate` remains at 12.50%).

C.3. Per-Stream Continual Breakdown and Dynamics on Drift Gauntlet

Table C.4 reports all 24 (stream, arm) evaluations underlying Table 3 (§6.3), and Figure C.2 plots the 5×4 per-group continual advantage trajectories $A[t, k] - B_k$ (modeled on Figure 1 of Chen et al. [5]). Across every scenario, `GUARD` achieves $F = 0.00$pp worst-case post-block forgetting and

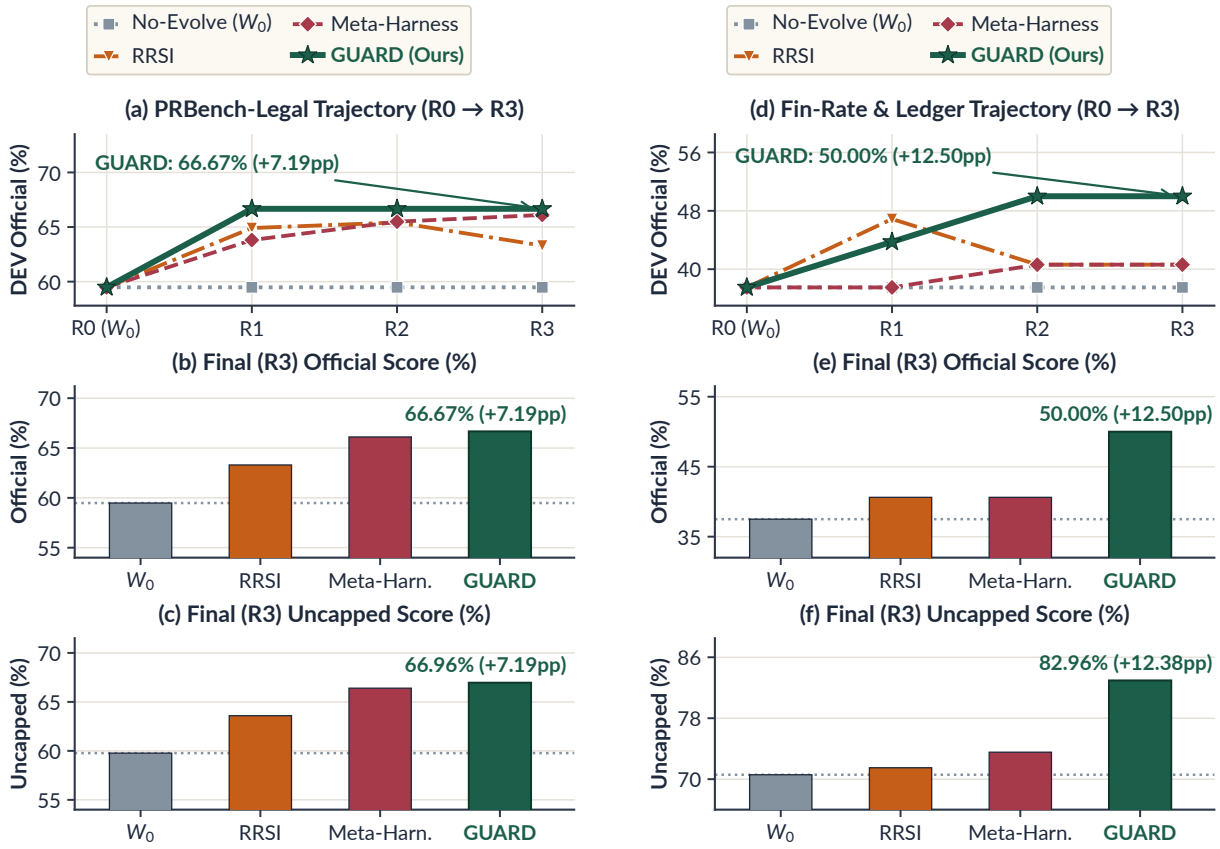


Figure C.1 | **Frozen DEV trajectories (R0 → R3, top), final Official (capped) scores (middle), and final Uncapped scores (bottom) on PRBench-Legal & Harvey LAB (panels (a)–(c), left) and Fin-Rate & Ledger (panels (d)–(f), right).** GUARD is non-decreasing with zero hurt tasks on both suites (66.67% Official / 66.96% Uncapped on PRBench-Legal; 50.00% Official / 82.96% Uncapped with $\sigma_w = 0.00$ on Fin-Rate & Ledger).

zero downward block transitions in $\bar{A}_t$: on opposing *workflow drift* (BFCL-Miss → BFCL-Base), Block 1 lifts BFCL-Miss +16.67pp (53.33% → 70.00%) and Soft-HURT rejects Block 2's over-broad rewrite (where 10 sampled bfcl reservoir replay tasks detected hurt = lift = 6.25pp), avoiding the −20.00pp and −13.34pp collapses of RRSI and Hermes+; on *topic drift* (BIRD-F1 → BIRD-Code → BIRD-TP), GUARD promotes clean database-scoped sections across all three blocks ($I = F = S = 0.00$pp); on cross-environment *domain drift*, in-round [iff] repair narrows Block 1 (ScienceWorld 74.80% → 82.73%, +7.93pp) while Soft-HURT rejects Block 2; and on four-block *compound drift*, GUARD leads Final Mean (74.17%), Δ_K (+3.33pp), $\bar{\Delta}$ (+2.08pp), and plasticity (+5.00pp) with $F = 0.00$pp.

Box C.1: An admitted scoped rule and a rejected overwrite on Drift Gauntlet

(a) Admitted scoped rule (*compound drift*, Block 1, section keyed to signature `bird:codebase_community`; validation net lift +5.21pp with no hurt task; held-out BIRD-Code 76.67% → 80.00%). One of the promoted Tier-2 rules, verbatim:

`[iff the prompt or Evidence expresses year or date predicates such as YEAR(col) = YYYY in a SQLite environment]` Replace unsupported `YEAR()` calls with valid SQLite date extraction (`STRFTIME('%Y', col) = 'YYYY'` or `col LIKE 'YYYY%'`) and enforce the year filter on every date column specified in the Evidence.

Table C.4 | **Complete 24-row per-stream continual RSI results on DRIFT GAUNTLET [5] under single-model `Gemini 4` (30 adaptation tasks/block, 30 held-out test tasks/group; summarized in Table 3).** $A[K,\cdot]$ lists final checkpoint accuracy across the stream's K task groups in arrival order. Bold marks the best evolving arm within each stream.

Continual Drift Stream (K blocks)	Arm	Final Mean	Δ_K	$\bar{\Delta}$	Plast. P	Interf. I	Forget. F	Stab. S	Final Group Vector $A[K,\cdot]$ (%)
Workflow Drift (K=2)	No-Evolve (W_0)	65.00%	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	[53.33, 76.67]
BFCL-Miss → BFCL-Base	HERMES+ / MEMTOOLAGENT	65.00%	0.00pp	+2.50pp	+8.34pp	−3.34pp	−13.34pp	−8.34pp	[53.33, 76.67]
	RRSI [45]	60.00%	−5.00pp	−1.67pp	**+11.67**pp	−13.34pp	−20.00pp	−16.67pp	[50.00, 70.00]
	GEPA [1]	**70.00%**	**+5.00**pp	**+5.00**pp	+3.33pp	**0.00**pp	**+3.34**pp	**+1.67**pp	[66.67, 73.33]
	META-HARNESS [20]	63.33%	−1.67pp	−5.84pp	+8.33pp	−20.00pp	0.00pp	−10.00pp	[53.33, 73.33]
	GUARD (Ours)	**70.00%**	**+5.00**pp	**+5.00**pp	+8.34pp	−6.67pp	0.00pp	−3.34pp	[**70.00**, 70.00]
Topic Drift (K=3)	No-Evolve (W_0)	67.78%	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	[76.67, 76.67, 50.00]
BIRD-F1 → BIRD-Code	HERMES+ / MEMTOOLAGENT	**70.00%**	**+2.22**pp	+0.74pp	−1.11pp	+1.66pp	**+1.66**pp	+1.66pp	[76.67, **80.00**, 53.33]
→ BIRD-TP	RRSI [45]	66.67%	−1.11pp	+1.85pp	−2.22pp	+4.17pp	−3.34pp	+0.42pp	[76.67, 73.33, 50.00]
	GEPA [1]	67.78%	0.00pp	+1.48pp	−2.22pp	**+5.00**pp	−1.66pp	**+1.67**pp	[76.67, 70.00, **56.67**]
	META-HARNESS [20]	68.89%	+1.11pp	**+3.70**pp	+1.11pp	**+5.00**pp	−4.99pp	0.00pp	[76.67, 76.67, 53.33]
	GUARD (Ours)	**70.00%**	**+2.22**pp	+1.11pp	**+2.22**pp	0.00pp	0.00pp	0.00pp	[76.67, **80.00**, 53.33]
Domain Drift (K=3)	No-Evolve (W_0)	76.05%	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	[74.80, 76.67, 76.67]
ScienceWorld → BFCL-Base	HERMES+ / MEMTOOLAGENT	76.26%	+0.21pp	+0.31pp	+0.96pp	**+1.66**pp	−1.78pp	−0.06pp	[75.43, 76.67, 76.67]
→ BIRD-F1	RRSI [45]	79.17%	+3.12pp	+0.42pp	**+9.32**pp	−4.17pp	−5.76pp	−4.96pp	[77.50, 76.67, **83.33**]
	GEPA [1]	**79.31%**	**+3.26**pp	+2.03pp	+8.41pp	−4.17pp	−2.72pp	−3.45pp	[77.93, **83.33**, 76.67]
	META-HARNESS [20]	79.29%	+3.24pp	+2.10pp	+4.53pp	−3.34pp	**+1.40**pp	−0.97pp	[81.20, 80.00, 76.67]
	GUARD (Ours)	78.69%	+2.64pp	**+2.64**pp	+2.64pp	0.00pp	0.00pp	**0.00**pp	[**82.73**, 76.67, 76.67]
Compound Drift (K=4)	No-Evolve (W_0)	70.84%	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	0.00pp	[76.67, 76.67, 53.33, 76.67]
BIRD-Code → BIRD-F1	HERMES+ / MEMTOOLAGENT	70.00%	−0.84pp	+0.83pp	+2.50pp	**+1.11**pp	−8.89pp	−3.89pp	[**80.00**, 70.00, 50.00, **80.00**]
→ BFCL-Miss → BFCL-Base	RRSI [45]	65.84%	−5.00pp	−2.50pp	−0.83pp	−2.41pp	−3.33pp	−2.87pp	[76.67, 76.67, 40.00, 70.00]
	GEPA [1]	72.50%	+1.66pp	+0.21pp	+1.67pp	+0.74pp	−3.33pp	−1.30pp	[73.33, **83.33**, 56.67, 76.67]
	META-HARNESS [20]	73.33%	+2.50pp	+0.62pp	−0.84pp	−1.48pp	**+3.33**pp	**+0.93**pp	[73.33, 76.67, 63.33, **80.00**]
	GUARD (Ours)	**74.17%**	**+3.33**pp	**+2.08**pp	**+5.00**pp	−0.74pp	0.00pp	−0.37pp	[**80.00**, 76.67, **70.00**, 70.00]

(b) Rejected overwrite (*workflow drift*, Block 2, shared signature `bfcl`). After `BFCL-Base` arrived, the pool winner rewrote the shared section. On the 24-task validation buffer it lifted the mean score by 6.25pp but regressed four tasks (three `BFCL-Miss` reservoir replay tasks and one `BFCL-Base` task) with a hurt penalty of 6.25pp, so Soft-HURT kept W_1 and logged the candidate to the rejection ledger. Held-out `BFCL-Miss` stayed at 70.00%, whereas the ungated RRSI rewrite lost 20.00pp on that group.

C.4. Five-Domain EVOHARNESSBENCH Breakdown, Component Ablations (`Analyze`, `Propose`, `Gate`), and Full Calibration Audit

Table C.5 reports the per-domain continual evolution breakdown across all five local-gym EVOHARNESSBENCH domains alongside the hosted `Calendar` track, and Table C.6 gives the complete nine-row offline self-reviewer calibration and Soft-HURT gate audit summarized in Table 4 (§6.4). We evaluate component-level ablations across both model regimes introduced in §5: Table C.7 tests leave-one-out removal of `Analyze` and `Gate` across all five local-gym EVOHARNESSBENCH domains (`Calendar`, `Drive`, `CSM`, `HR`, and `ITSM`; 16 continual stages, 97 visible adaptation tasks, and 236 held-out test tasks) in the teacher–student distillation setting (`gemini-3.8-flash` solver paired with `claude-opus-5-5` self-reviewer, skill distiller, and prompt mutator, $M = 2$, seed s_0), while Table C.8 isolates the `Propose` stage under our primary single-model `Gemini 4` configuration across both EVOHARNESSBENCH `ITSM` evolution axes (`evolving_tools`, T=3 stages, and `evolving_skills`, T=4 stages; pooled N=132 held-out test tasks $\times N_{\text{rep}}$=3 replicates across four controlled configurations, strictly isolated rollouts with zero cache sharing and train-only self-review, distillation, and Soft-HURT gating at $\delta_{\text{noise}} = 15$):

- **Necessity of the `Analyze` stage (paired regret ranking, rule pinning, and contrasts).** Jointly ablating $\text{Priority}_t(c) = \Delta_{\text{hurt},t}(c) + \Delta_{\text{gap},t}(c)$, rule pinning, and paired (τ^+, τ^-) trace contrasts in favor of single-round failure frequency $1 - \text{obey}(c; W_{t-1})$ (`GUARD w/o Analyze`), while keeping the identical $M = 2$ candidate pool and Soft-HURT promotion gate, reduces five-domain macro accuracy by −7.33pp (66.46% → 59.13%, −6.68pp below static W_0), strict solve rate by −5.39pp (38.31% → 32.92%), and mean forward transfer to −7.57pp (vs. +1.00pp under `GUARD`). The Soft-

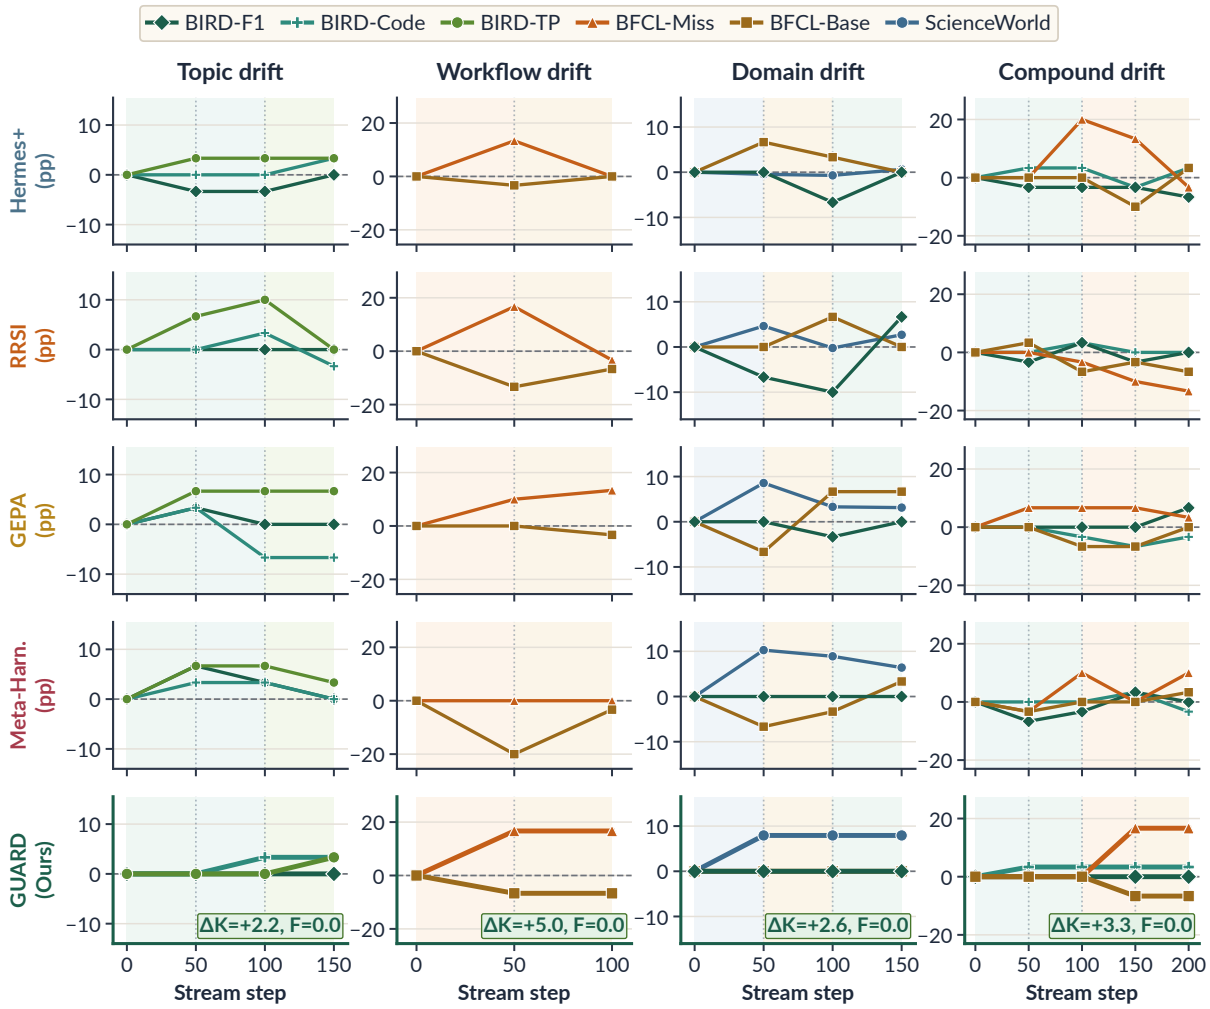


Figure C.2 | **Per-group continual learning dynamics $A[t,k] - B_k$ across all five evolving methods and four DRIFT GAUNTLET scenarios under single-model `Gemini 4`** (shaded spans mark each group's active online adaptation block; layout following Figure 1 of [5]). Baselines exhibit sharp post-block drops once subsequent blocks arrive, whereas `GUARD` (bottom row) retains 100% of post-block gains (F = 0.00pp on all four streams).

HURT gate alone cannot rescue an unscoped proposer: it still promotes four stages on visible source tasks whose frequency-distilled rules overfit early stages and regress both multi-stage skill domains on held-out cohorts (`CSM` drops 61.40% → 38.94%, −22.46pp; `ITSM` drops 60.28% → 46.12%, −14.16pp).

- **Necessity of the Soft-HURT promotion gate (`Gate` stage).** Retaining paired regret priority, rule pinning, and contrastive candidate generation while unconditionally promoting the within-group standardized-fitness winner across all 16 stages drops five-domain macro accuracy by −10.35pp (66.46% → 56.11%), strict solve rate by −10.04pp (38.31% → 28.27%), mean forward transfer to −5.62pp, and weighted backward transfer to −6.90pp. Ungated promotion causes negative transfer vs. W_0 on 3/5 domains (`Calendar` −1.89pp, `HR` −7.96pp, and `ITSM`, where an unvetted Stage-v4 candidate instructs a hallucinated MCP tool call `get_current_user` and collapses cumulative test accuracy to 7.61%), whereas full `GUARD` maintains 0/5 regressed domains.
- **Controlled `Propose`-stage attribution under single-model `Gemini 4` (parallel M=2 pooling**

vs. equal-TM sequential search). Holding `Analyze` and Soft-HURT fixed across four controlled arms on dual-axis `ITSM` (Table C.8), all three `GUARD` configurations beat W_0 (72.61% pooled task ACC / 73.39% macro ACC) with 0/2 regressed axes. At matched total candidate-evaluation budget $TM = 2T$, parallel M=2 candidate pooling with group-standardized fitness selection (81.08% pooled / 81.26% macro ACC) outperforms single-candidate sequential M=1, $2T$ search (78.85%, +2.23pp, $p = 0.0267$; +3.26pp, $p = 0.0495$ on `evolving_tools`) while cutting rejected gate rounds by 64% (4 vs. 11 ledger entries), beats single-candidate M=1, T (77.20%) by +3.88pp pooled ($p = 0.0029$) and +6.64pp on `evolving_tools` ($p = 0.0033$), and achieves positive backward transfer (BWT +0.61% vs. −1.28% and −1.65% under M=1).

Table C.5 | **Full 5-domain EvoHarnessBench continual evolution summary (**236 **local-gym test tasks across** 16 **stages plus the hosted `Calendar` track).** Hosted and local-gym runners follow their respective BWT/FWT definitions (§5), ordered Score / Pass as in Table 2. On local-gym tracks where `GUARD` abstains ($W_{\text{final}} = W_0$ on `Calendar` and `Drive`), raw cohort BWT reflects cross-stage environment drift of W_0 itself (ΔBWT vs. W_0 = 0.00pp).

Domain Track	Models	Stages (Test/Train)	W_0 (ACC / Solved)	GUARD (ACC / Solved)	Stage Highlights (W_0 → GUARD)	BWT (Score / Pass)	FWT
`Calendar` (Tools, Hosted SDK)	`Gemini 4`	v1–v3 (43 / 18)	71.91% / 48.84%	**80.80%** / **58.14%**	Cohort v1 (H_{v_3}): 70.46% → **78.47**% (+**8.01**pp); v3: 64.29% → **89.29**%	+**0.85**pp / +**2.56**pp	**0.00**pp (Abstained v2–v3)
`Calendar` (Tools, Local Gym)	`flash+opus`	v1–v3 (43 / 18)	**77.15%** / **48.84%**	**77.15%** / **48.84%**	Stages v1 / v2 / v3: 78.22%/85.67%/**89.29**% ACC (75% Pass)	−4.10pp / −8.45pp	**0.00**pp ($W_{\text{final}} = W_0$)
`Drive` (Tools, Local Gym)	`flash+opus`	v1–v3 (45 / 19)	**71.78%** / **46.67%**	**71.78%** / **46.67%**	Stages v1 / v2 / v3: 77.70%/77.90%/**80.56**% ACC	−7.41pp / **0.00**pp	**0.00**pp ($W_{\text{final}} = W_0$)
`CSM` (Skills, Local Gym)	`flash+opus`	v1–v3 (36 / 14)	59.80% / **30.56%**	**61.40%** / 25.00%	Stage v1: 51.20% → **55.37**% (+**4.17**pp); v2: 66.03% → **69.87**% (+**3.84**pp)	−2.08pp / **0.00**pp	+**2.91**pp Mean (+**4.17**pp v1)
`HR` (Skills, Local Gym)	`flash+opus`	v1–v3 (53 / 22)	**61.67%** / **32.08%**	**61.67%** / **32.08%**	Stage v1: 49.62% → **51.52**% (+**1.90**pp); v2 / v3: 63.03%/**66.80**%	−0.95pp / **0.00**pp	+**0.63**pp Mean (+**1.90**pp v1)
`ITSM` (Skills, Local Gym)	`flash+opus`	v1–v4 (59 / 24)	58.65% / 30.51%	**60.28%** / **38.98%**	Stage v2: 53.92% → **59.67**%; v4: 66.54% → **71.21**% (+**18.19**pp Pass)	**0.00**pp / **0.00**pp	+**1.44**pp Mean (+**5.75**pp v2)

Table C.6 | **Full offline calibration of the reference-free two-layer self-reviewer $S(\tau)$ and Soft-HURT promotion gate against hidden official benchmark graders** (summarized in Table 4, §6.4). All statistics are computed post-hoc over cached evaluation trajectories and stage/block decisions with zero oracle labels, test cases, or rubrics exposed during adaptation.

Calibration / Gate Dimension	Benchmark & Sample (N)	Unguarded / Ref.	GUARD (Ours)
False-pos. pass ($\geq$ 75%) on Layer-1-failed traces*	EvoHarnessBench `Gemini 4` (54 vs. 23 L1-fail)	22.2% (1-layer)	**0.0**% (0/23, −100%)
Initial (1-layer) vs. final hosted configuration ACC / Solved	EvoHarnessBench `Calendar` (43 test)	75.98% / 48.84%	**80.80**% / **58.14**%
Official verifier ACC (Layer-1-only fail vs. pass)	EvoHarnessBench pooled (4,059 test: 1,920 hosted, 2,139 local)	45.97% (L1 fail)	**73.90**% (+27.92pp, 82.6% NPV)
Episode semantic bias ($S - V^\star$) & corr.	PAST-Bench (114 learn episodes)	0.853 (official $V^\star$)	**−0.012** (r=0.369, r_s=0.377)
Within-family playbook ranking concordance	PAST-Bench (72 playbooks, 169 pairs)	50.0% (random)	**82.2**% (75.6% non-tied; r=0.432)
Exec.-gated promotions ($\Delta_{\text{test}} \geq 0$)†	Drift Gauntlet (9) + EvoHarness (5) + prof. (5)	—	**19**/**19** (**15** >0, 4 ties)
Same Soft-HURT gate w/o `Analyze` stage	EvoHarnessBench local (16 stages, 5 dom.)	4 prom., 2/5 dom. regr.	**4** prom., **0**/**5** dom. regr.
Bypassing Soft-HURT on `Calendar v3`	EvoHarnessBench local (43 test)	−6.98pp Solved / −2.40pp ACC	**0.00**pp loss (rejected)
Tracks regressed in final ACC vs. W_0	Drift Gauntlet (4), EvoHarness (1 host., 5 loc.)	5/16 DG base. runs; 3/5 w/o Gate	**0**/**4** DG, **0**/**6** EvoH.

*Sanity check: Eq. (2) zeroes S^{strict} and Layer 1 fails execution-broken traces by construction, whereas the 1-layer baseline lacks Layer 1; row 3 provides the non-tautological Layer-1-only validation, and row 2 compares the initial and final hosted configurations, which also differ in gating set and distiller instruction (Appendix B). †Δ_{test} = held-out score of a promoted harness minus W_0 on the active cohort/group of the stage/block/round where it was promoted (evaluated at H_{ν_3} on hosted `Calendar`, where W_3=W_1; `CSM v2`'s incremental delta over the `CSM v1` playbook is not separately tabulated). Execution-gated suites yield 19/19 $\Delta_{\text{test}} \geq 0$ (15/19 strictly > 0): 9/9 Drift Gauntlet block promotions across 12 blocks (*workflow* b1 +16.67pp; *topic* b1 +0.00, b2 +3.33, b3 +3.33pp; *domain* b1 +7.93pp via in-round `[iff]` repair, b3 +0.00pp; *compound* b1 +3.33, b2 +0.00, b3 +16.67pp; Table C.4), 5/5 EvoHarnessBench promotions (hosted `Calendar v1` +8.01pp, local `CSM v1` +4.17pp, `CSM v2` +3.84pp, `HR v1` +1.90pp [returned to W_0 at Stage `v2` by the stage-shift check; Appendix B], and `ITSM v1` +0.00pp tie [+1.63pp final ACC]; Table C.5), and 5/5 professional-suite role promotions across 4 promoted rounds (PRBench-Legal `R1 Finance` +28.57pp and `R1 Legal` +5.77pp [combined `R1` +7.19pp; `R2`–`R3` retained the `R1` incumbents], and Fin-Rate `R1`/`R2`/`R3` +6.25/+12.50/+12.50pp, where `R3` promotes the `StratOps` playbook with unchanged `DEV` Official; Table C.3). Separately, on critic-scored PAST-Bench (12 families, 24/24 round winners admitted; no gate-discrimination signal), `GUARD` lifts 12-family mean Task by +0.482 over W_0 (0.372 → 0.854, Table 1; all 12 families > W_0, with 0/48 W.Prec. or Poll. family losses vs. persistent baselines in Table B.4).

Table C.7 | **Matched five-domain leave-one-out ablation of `Analyze` and `Gate` across all** 236 **local EvoHarnessBench test tasks (**16 **continual stages,** 97 **visible source tasks).** All arms share identical models (`gemini-3.8-flash` agent, `claude-opus-5-5` self-reviewer/distiller), pool size (M = 2), and seed prompt (s_0), with zero oracle skills or held-out verifier leakage. Bold marks the best value per column.

Arm	Ablated Stage & Mechanism	Macro ACC	Δ vs. GUARD	Macro Solved	Mean FWT	Regressed Domains↓	Promoted / Rejected
`No-Evolve` (W_0, Control)	Static base harness (no adaptation)	65.81%	−0.65pp	37.73%	0.00pp	**0 / 5**	0 / 16
GUARD (Ours)	**Full harness (`Analyze` + `Propose` + Soft-HURT `Gate`)**	**66.46%**	—	**38.31%**	**+1.00pp**	**0 / 5**	4 / 12
`GUARD w/o Analyze`	Ablates `Analyze` (1 − obey(c; W_{t-1}), no pinning or contrasts)	59.13%	−7.33pp	32.92%	−7.57pp	2 / 5	4 / 12
`GUARD w/o Soft-HURT Gate`	Ablates `Gate` (promotes within-group winner unconditionally)	56.11%	−10.35pp	28.27%	−5.62pp	3 / 5	16 / 0

Table C.8 | **Controlled single-model `Gemini 4` `Propose`-stage and equal-budget (TM) ablation on EvoHarnessBench `ITSM` (7 continual stages across `evolving_skills` [T=4, N=59] and `evolving_tools` [T=3, N=73]; pooled N=132 held-out test tasks ×3 replicates = 396 final-stage rollouts/arm).** All evolving arms run in isolated MCP sandboxes with zero shared rollouts and share the identical paired-regret `Analyze` stage, rule pinning, and Soft-HURT `Gate` (δ_{noise}=15) strictly on per-stage `train` splits; to isolate pooling, every arm samples candidates from the shared seed s_0 in every round, without cross-round ContrastEvolve updates. Two-sided paired sign-flip p-values are computed over all N=132 task-level deltas (†: p=0.0267 vs. `GUARD Sequential` [M=1, $2T$] and p=0.0029 vs. `GUARD w/o Pool` [M=1, T]); all three `GUARD` configurations have 0/2 regressed axes, and bold marks best per column.

Arm (`ITSM`, Unified `Gemini 4`)	`Propose`-Stage Configuration	Skills ACC	Tools ACC	Task ACC (Macro)	Δ vs. W_0 (p)	Macro Solved	Macro FWT / Task BWT	Prom. / Rej. (Ledger)
`No-Evolve` (W_0, Control)	Static base harness (M=0)	80.70%	66.08%	72.61% (73.39%)	—	51.36%	+0.00pp / −3.44%	0 / 7 (0)
`GUARD w/o Pool` (M=1, T)	Single candidate (TM=T)	82.46%	72.95%	77.20% (77.70%)	+4.59pp (0.0009)	53.25%	+3.58pp / −1.65%	1 / 6 (6)
`GUARD Sequential` (M=1, $2T$ seq.)	2 sequential rounds/stage (TM=$2T$)	81.96%	76.34%	78.85% (79.15%)	+6.24pp (0.0006)	56.09%	+5.50pp / −1.28%	2 / 5 (11)
`GUARD` (Ours)	**Parallel M=2 pool + $\arg\max_m \hat{A}_m$ (TM=$2T$)**	**82.93%**	**79.59%**	**81.08%† (81.26%)**	+8.47**pp** (1.2×10^{-5})	**56.61%**	+6.69**pp** / +0.61%	**3** / **4** (4)