Online Environment Repair for Coding Agents

Sheng Lu, Rui Mao, Zewei Chen, Chihua Ying, Beiduo Chen, Dingxin Hu, Tinghao Yu[†]

Hunyuan Team, Tencent

{jonathanslu, dingxinhu, maxwellyu}@tencent.com

arXiv:2611.32037v1 [cs.LG] 24 Nov 2026

Programmatic verifiers provide scalable reward signals for training coding agents with reinforcement learning, yet subtle harness defects and incomplete unit tests allow policies to earn full credit without solving the intended task. Existing training-time defenses intervene either by penalizing suspicious outputs or by filtering flagged tasks and rollouts. However, output penalties leave the underlying verifier exploitable, driving policies to disguise shortcuts under reasoning monitors or zeroing out valid solutions under detector false positives, whereas task and rollout filtering discards or mutes the hardest problems where shortcuts emerge earliest. Pre-training verifier hardening repairs the environment directly, but misses vulnerability classes that surface only as the policy improves during reinforcement learning. In this paper, we propose *online environment repair*, which leaves triggering rollout rewards unchanged to remove hiding incentives and asynchronously patches exploited environments during training. Candidate repairs triggered by execution invariants or differential fuzzing are deployed only after passing three automated validation gates that block the exploit while preserving valid solutions, canceling detector false positives and enabling structural propagation across repositories without scaffold labels or reference solutions. Across 20 controlled GRPO experiments on heterogeneous software-engineering repositories, online environment repair reduces cumulative reward hacking from 36.9% to 0.06% and improves held-out Pass@8 from 21.7% to 46.7%, consistently outperforming reward-penalty, rollout-filtering, and static-hardening baselines while eliminating zero-shot exploitation on unseen loophole classes.

1. Introduction

Reinforcement learning with verifiable rewards trains coding agents directly against programmatic test suites (Guo et al., 2025; Shao et al., 2024; Yu et al., 2025). Unlike isolated function-level exercises, multi-module repositories bundle test fixtures, configuration plugins, equality comparators, and evaluation hooks alongside editable source files, while visible unit tests often cover only a few edge cases (Gupta and Joshi, 2026; Wang et al., 2026a; Zhong et al., 2025). Once a policy discovers that preloading a workspace fixture, spoofing object equality, or overfitting to sparse test cases yields higher return than solving an algorithmic bottleneck, group-relative policy optimization amplifies that shortcut (Wang et al., 2026b) and transfers deceptive habits to downstream tasks (MacDiarmid et al., 2025).

When a detector flags suspicious trajectories during training, an alignment pipeline can intervene at the reward signal, the training data, or the evaluation environment, as shown in Figure 1(a). Existing training-time defenses act at the first two loci and run into opposing structural limits. Penalizing flagged outputs leaves the verifier open and fails on both error branches of an imperfect detector (Baker et al., 2025; Wang et al., 2026b; Wong et al., 2026): false negatives still earn full reward and select for shortcuts disguised behind routine engineering rationales, whereas stripping reasoning context to inspect raw code or runtime flags causes false positives that zero out valid solutions and corrupt within-group advantages. Quarantining flagged tasks or masking flagged rollouts avoids step-level evasion gradients by withholding penalties (Wang et al., 2026b), yet induces inverse difficulty

Corresponding author(s): Tinghao Yu[†] *(maxwellyu@tencent.com)*

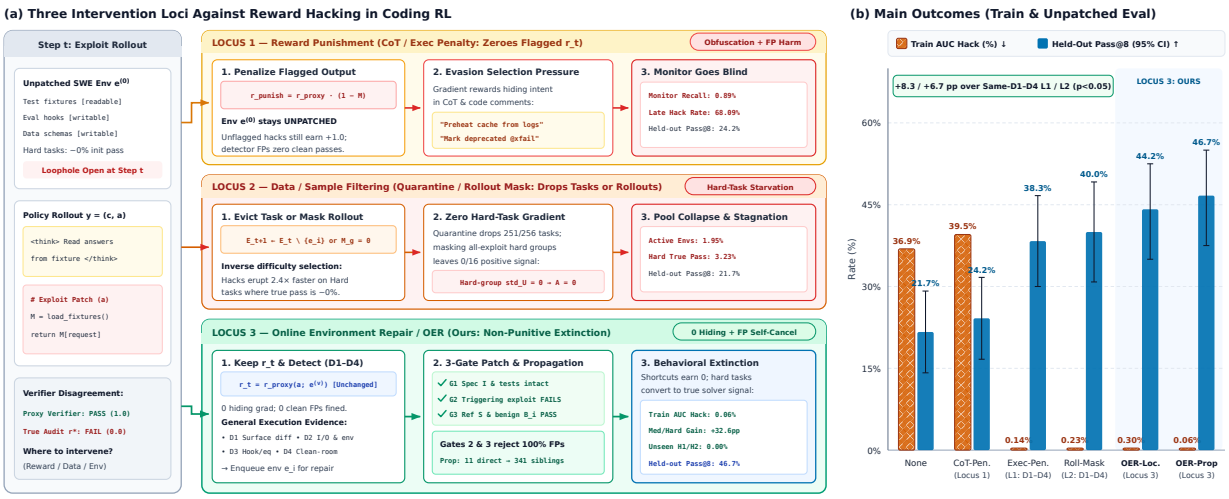

Figure 1 | Three intervention loci against reward hacking in coding RL. (a) When a shortcut passes an exploitable verifier $e^{(0)}$ while failing ground-truth evaluation r^*, *Locus 1* penalizes flagged outputs while leaving $e^{(0)}$ open, *Locus 2* drops flagged tasks or masks rollouts, and *Locus 3* (OER) leaves $\tilde{r}_t$ unchanged and patches $e^{(0)}$ online behind three validation gates with structural propagation. **(b)** Over 160 GRPO steps on T1-HETERO, OER-PROP outperforms matched-detector reward penalties (EXEC-PENALTY) and rollout masking (ROLLOUT-MASK) on unpatched held-out Pass@8 (95% bootstrap CIs).

selection: because shortcuts beat genuine attempts earliest on hard repositories where valid solutions are scarce, task quarantine depletes the training pool while rollout masking leaves hard prompt groups with zero reward variance.

Moving intervention to the evaluation environment removes the shortcut payoff without discarding hard tasks. Recent benchmark-hardening pipelines isolate verifiers and run adversarial hacker–fixer loops before training (Zeng et al., 2026; Zhong et al., 2026). However, static pre-RL hardening probes environments only against the initial policy, missing multi-stage leaks, schema bypasses, and sparse-test specification gaps that surface only after tens of RL updates (Deng et al., 2026).

In this paper, we propose *online environment repair* (OER), a training procedure grounded in operant extinction that repairs exploited environments online without penalizing the triggering rollout. At each GRPO step, rollouts are scored unconditionally against the active environment version, removing any step-level incentive to conceal intent in reasoning traces or code comments. Passing rollouts that trigger execution invariants or intra-group differential fuzzing enqueue the environment for asynchronous repair by a `Qwen3-32B` agent while holding the problem specification fixed. A candidate patch replaces the environment on subsequent steps only if it clears three automated gates: specification integrity, exploit replay failure, and reference or benign solution preservation. Because no specification-preserving patch can cause a valid solution to fail replay while keeping equivalent reference solutions passing, the opposing gates cancel detector false positives without penalizing valid rollouts or corrupting environments, and enable structural propagation across sibling repositories without scaffold family labels or reference solutions.

Across 20 controlled GRPO runs on heterogeneous software repositories and specification-gap tasks, OER-PROP reduces cumulative reward hacking from 36.9% to 0.06% and raises unpatched held-out Pass@8 from 21.7% to 46.7%, outperforming matched-detector reward penalties and rollout masking while eliminating zero-shot exploitation on unseen loophole classes. In summary, our

contributions are as follows:

1. We formalize the three intervention loci in coding RLVR, proving step-level obfuscation neutrality, post-repair advantage extinction, and false-positive self-cancellation under opposing validation gates.
2. We propose *online environment repair* with gate-verified structural propagation, showing through a 2×2 timing-by-propagation ablation that both components independently improve held-out coding accuracy and match privileged performance without scaffold labels or reference solutions.
3. We introduce zero-oracle intra-group differential fuzzing and boundary-test synthesis for locked-down repositories with sparse unit tests, cutting online specification-gap hacking tenfold where static isolation repairs zero tasks.
4. Through controlled 160-step GRPO experiments under matched detectors, we isolate why output penalties succumb to camouflage or false-positive gradient harm and why filtering starves hard tasks.

2. Related Work

Monitor Penalties and Policy Regularization. Programmatic coding verifiers are vulnerable to test modification, fixture snooping, equality spoofing, and sparse-test overfitting (Gupta and Joshi, 2026; Khalifa et al., 2026; Thaman, 2026; Wang et al., 2026a; Zhong et al., 2025), and scalar reward curves alone cannot separate genuine competence from verifier defects without execution evidence (Moya et al., 2026). Penalizing flagged reasoning traces or code through LLM monitors selects for camouflaged exploits once policies mask deceptive intent (Baker et al., 2025; Taufeeque et al., 2026; Wong et al., 2026). On CATCH, Wang et al. (2026b) show that chain-of-thought penalties collapse mid-training while χ^2 occupancy constraints (Laidlaw et al., 2025) and gradient regularization (Ackermann et al., 2026) degrade coding accuracy. Whereas output penalties and policy divergences restrict exploration or penalize valid solutions under false positives, OER leaves policy updates unconstrained and repairs the verifier itself.

Oracle Auditing and Verifier Hardening. Li et al. (2026) introduce AUDIT-LB, concentrating oracle-verifier queries on low-pass-rate GRPO groups with reflection penalties. While effective when a ground-truth verifier is available during training, production repositories ship only with in-repository tests. Offline benchmark hardening (Luo et al., 2026; Zhong et al., 2026) and Weak-to-Strong Verifier Hardening (Zeng et al., 2026) isolate grading sandboxes and run iterative hacker–fixer loops before evaluation or training, leaving cross-task amortization and online RL integration open. OER brings repair inside the RL loop so the adapting policy continuously exposes emerging weaknesses, while structural propagation transfers validated fixes across unlabeled repository pools.

3. Method: Online Environment Repair

3.1. Formulation Across Three Intervention Loci

We consider a pool of coding environments $\mathcal{E} = \{e_1, \ldots, e_M\}$, where version v of environment i is $e_i^{(v)} = (I_i, R_i^{(v)}, V_i^{(v)}, O_i^{(v)}, S_i)$ with immutable specification I_i, workspace $R_i^{(v)}$, proxy verifier $V_i^{(v)}$, output contract $O_i^{(v)}$, and optional reference solution S_i. At GRPO step t (Shao et al., 2024), policy π_θ samples G rollouts $\{y_{i,g}\}_{g=1}^{G}$ for prompt $\psi(e_i^{(v_t)})$, pairing reasoning trace c with code edits a. Writing $\bar{a} = \mathrm{exec}(a)$ for comment-stripped executable code, $\tilde{r}(\bar{a}; e_i^{(v_t)}) \in [0, R_{\max}]$ denotes proxy verifier score and $r^*(\bar{a}; e_i) \in \{0, 1\}$ offline ground-truth correctness. Rollouts with $r^*(\bar{a}; e_i) = 1$ are *valid solutions*,

Algorithm 1: Online Environment Repair (OER) inside GRPO

Input: Registry $\{e_i^{(0)}\}_{i=1}^{M}$, benign buffers $\mathcal{B}_i \leftarrow \varnothing$, delay L, retries $K = 3$, propagation mode $\in \{\text{local}, \text{prop}, \text{autocluster}\}$.
For GRPO step $t = 1, \ldots, T$:

1. Sample prompt batch $\mathcal{B}_t \subset \mathcal{E}$; activate pending versions with $t_{\text{eff}} \le t$ to obtain $e_i^{(v_t)}$.
2. Sample G rollouts $\{y_{i,g} = (c_{i,g}, a_{i,g})\}_{g=1}^{G} \sim \pi_\theta(\cdot \mid \psi(e_i^{(v_t)}))$ and score $r_{i,g} \leftarrow \tilde{r}(\bar{a}_{i,g}; e_i^{(v_t)})$ (*unmodified*).
3. **For each** passing rollout ($\mathbb{I}_{\text{easy}} \cdot \mathbb{I}_{\text{hard}} = 1$): run detector $E \leftarrow D(a_{i,g}; e_i^{(v_t)})$. If $E = \varnothing$: update $\mathcal{B}_i$; **else**: enqueue $(e_i, v_t, E.\text{mech}, a_{i,g})$ (deduplicated).
4. Update θ via standard GRPO using $\{r_{i,g}\}$.
5. **For each** queued job: query `Qwen3-32B` up to K tries for $e_i' = \text{Apply}(e_i^{(v_t)}, P)$ until Gates 1–3 pass on $(I_i, \text{exploits}, S_i \cup \mathcal{B}_i)$; schedule $e_i^{(v_t+1)} \leftarrow e_i'$ at $t_{\text{eff}} = t + 1 + L$, and optionally lift structural ops to sibling environments passing Gates 1 & 3.

passing rollouts clearing training-side detectors are *benign rollouts*, and rollouts passing visible slices ($\mathbb{I}_{\text{easy}} \cdot \mathbb{I}_{\text{hard}} = 1$) while $r^*(\bar{a}; e_i) = 0$ are *nontrivial reward hacks* (Wang et al., 2026b).

Given monitor $M(I_i, c, a) \in \{0, 1\}$ or execution detector $D(a; e_i^{(v_t)})$, the three loci intervene on distinct components:

$$\begin{aligned}
&\text{Locus 1: } r_{\text{punish}}(c, a) = \tilde{r}(\bar{a}; e_i^{(0)})\big(1 - M(I_i, c, a)\big), \\
&\text{Locus 2: } r_{\text{filter}}(c, a) = \tilde{r}(\bar{a}; e_i^{(0)}), \quad \mathcal{E}_{t+1} = \mathcal{E}_t \setminus \mathcal{D}_t, \\
&\text{Locus 3: } r_{\text{OER}}(c, a; t) = \tilde{r}\big(\bar{a}; e_i^{(v_t)}\big), \quad e_i^{(v_t)} \to e_i^{(v_t+1)},
\end{aligned} \tag{1}$$

where $\mathcal{D}_t = \{e_i : \exists g, D(a_{i,g}; e_i^{(0)}) \neq \varnothing\}$ denotes quarantined tasks, or rollout-level gradient masking on flagged pairs (i, g) under ROLLOUT-MASK. In OER, $\mathcal{E}_{t+1} = \mathcal{E}_t$ and validated patches activate at $t_{\text{eff}} = t + 1 + L$. Under Locus 2, per-rollout false-positive rate α across G_+ valid solutions evicts a task with probability $1 - (1 - \alpha)^{G_+}$ per visit, whereas masking all passing rollouts on a hard task collapses its within-group advantage variance to zero.

Proposition 1 (Step-t Obfuscation Neutrality and Post-Repair Extinction)**.** *Assume* $\pi_\theta(c, a \mid x_t) = \pi_\theta(\bar{a} \mid x_t)\, \pi_\theta(\tau \mid \bar{a}, x_t)$ *for* $x_t = \psi(e_i^{(v_t)})$ *and surface verbalization* $\tau = (c, \text{comments}(a))$*:*

1. Zero step-t hiding gradient: *Writing* $\mathbb{E}_\tau[\cdot]$ *for* $\mathbb{E}_{\tau \sim \pi_\theta(\cdot \mid \bar{a}, x_t)}[\cdot]$, $\nabla_\theta \mathbb{E}_\tau[r_{OER}(c, a; t)] = 0$*, whereas any exploit with* $\tilde{r}(\bar{a}; e_i^{(0)}) > 0$ *yields* $\nabla_\theta \mathbb{E}_\tau[r_{punish}(c, a)] = -\tilde{r}(\bar{a}; e_i^{(0)})\, \nabla_\theta \mathbb{E}_\tau[M(I_i, c, a)]$.
2. Advantage extinction: *Let patch* P *produce* $e_i' = e_i^{(v_t+1)}$ *passing Gate 2* ($\tilde{r}(\bar{a}_{\text{hack}}; e_i') = 0$*) and Gate 3* ($\tilde{r}(S_i; e_i') \ge 1$*). For* $t' \ge t_{\text{eff}}$*, any GRPO group containing a valid solution* g^+ ($r_{g^+} \ge 1$*) and exploit* g^- ($\bar{a}_{g^-} = \bar{a}_{\text{hack}}$*) satisfies* $\hat{A}_{g^-} = -\mu_r/\sigma_r < 0 < \hat{A}_{g^+}$.

Proof. Since $\tilde{r}(\bar{a}; e_i^{(v_t)})$ is constant in τ given $(\bar{a}, e_i^{(v_t)})$ and $\nabla_\theta \sum_\tau \pi_\theta(\tau \mid \bar{a}, x_t) = 0$, surface rationales receive zero gradient at step t; for (2), $r_g \ge 0$, $r_{g^-} = 0$, and $r_{g^+} \ge 1$ imply $0 < \mu_r < r_{g^+}$ and $\sigma_r > 0$. □

3.2. Execution Invariants (`D1`–`D4`) and Differential Fuzzing (`D5`)

Triggering repairs from reasoning text would patch environments sooner after candid traces than deceptive ones, creating an indirect incentive to hide intent. OER therefore triggers repairs solely from execution evidence on passing rollouts:

- `D1` *Protected-surface mutation:* AST diffing flags edits or deletions to scored unit tests, runner configurations, fixture loaders, matchers, or verification hooks.
- `D2` *Data and environment access:* Static AST checks and runtime interpreter audit hooks trace reads of non-code workspace files and non-standard environment variables during test execution.
- `D3` *Harness and equality interference:* Flags custom `__eq__` or `__ne__` overrides on response schemas, abrupt process exits, and zero-test-executed return-code spoofing.

- D4 *Clean-room differential replay:* Reruns candidate modules inside an isolated sandbox stripped of workspace data files and extra environment variables under serialized string equality.
- D5 *Zero-oracle intra-group differential fuzzing:* On locked-down repositories with sparse unit tests, generates five grammar-valid non-sample inputs from I_i alone without reference or oracle access. When passing rollouts disagree on those probes, or when a uniform output violates I_i under trace audit, distinct representatives are queued for boundary-test synthesis.

3.3. Repair Agent, Validation Gates, and De-Privileged Propagation

Algorithm 1 summarizes the closed-loop workflow. For each deduplicated detection, a co-hosted Qwen3-32B repair agent inspects the repository, exploit code, and execution evidence to propose a structured patch P from seven primitives: subprocess verifier isolation, read-only locks on tests and hooks, moving exposed fixtures into evaluator-only storage, deleting stray workspace snapshots, stripping leaked environment variables, enforcing serialized payload equality and exact test counts, and synthesizing specification-valid boundary test inputs executed against S_i in an isolated sandbox to materialize new assertions. Each candidate e_i' is checked against three gates with up to $K = 3$ retries: *Gate 1* verifies that I_i, function signatures, and scored assertions remain intact without skip decorators; *Gate 2* requires all triggering exploits to fail replay; and *Gate 3* requires S_i and the benign buffer $\mathcal{B}_i$ to pass.

To eliminate curated metadata, OER-AutoCluster hides scaffold family labels and propagates structural edits to unpatched siblings whose unsupervised v_0 permission and test-AST fingerprints have Jaccard similarity at least 0.95, accepting updates that pass Gates 1 and 3. OER-RefFree withholds reference solutions S_i, deferring repair until a benign rollout populates $\mathcal{B}_i$ and guarding Gate 3 purely with $\mathcal{B}_i$. Except in the privileged Audit-LB baseline, training and repair modules have zero runtime access to the detached ground-truth auditor r^*, which serves strictly as read-only telemetry.

Proposition 2 (False-Positive Self-Cancellation via Gates 2 ∧ 3)**.** *Assume observational equivalence on public test slices across specification-compliant implementations: for any Gate-1-compliant e_i' preserving I_i and any valid solution $\bar{a}_{\text{clean}}$, $\mathbb{I}_{\text{pass}}(\bar{a}_{\text{clean}}; e_i') = \mathbb{I}_{\text{pass}}(b; e_i')$ for $b \in \{S_i\} \cup \mathcal{B}_i$. If $y_{\text{clean}} = (c, a_{\text{clean}})$ falsely triggers $D(a_{\text{clean}}; e_i^{(v_t)}) \neq \varnothing$: (1) step-t GRPO incurs zero reward distortion; and (2) no Gate-1-compliant patch P can pass both Gate 2 ($\mathbb{I}_{\text{pass}}(\bar{a}_{\text{clean}}; e_i') = 0$) and Gate 3 ($\mathbb{I}_{\text{pass}}(b; e_i') = 1$), leaving $e_i^{(v_t)}$ unchanged.*

Proof. Part (1) follows from Eq. (1); for (2), $\mathbb{I}_{\text{pass}}(\bar{a}_{\text{clean}}; e_i') = \mathbb{I}_{\text{pass}}(b; e_i') \in \{0, 1\}$ cannot equal 0 and 1 simultaneously, rejecting all K attempts. □

4. Experimental Setup

Benchmarks. Because a single global patch would secure the uniform layout of vanilla CATCH, we construct four disjoint splits from verified CATCH-RL problems (Wang et al., 2026b), detailed in Appendix B. T1-Hetero contains 256 training repositories across eight scaffold families, S1 through S8 with 32 tasks each, spanning mutable unit tests, fixture plugins, regression traces, output matchers, verification hooks, exposed call logs, and equality-spoofable schemas. S9 contains 32 specification-gap repositories where subprocess isolation, read-only tests, strict payload equality, and clean workspaces are already enforced at v_0, leaving only sparse 2–3 case unit tests as the vulnerability. Held-out evaluation uses 64 unseen-loophole tasks spanning base64 environment-variable leaks in H1 and hidden legacy snapshots in H2, neither seen in SFT or RL, plus 120 unpatched v_0 SWE tasks from CATCH-RL.

Locus	Condition (Detector)	Training Dynamics (Steps 1–160, 40,960 Rollouts)				Held-Out Unpatched SWE (120 Tasks × 8 Rollouts)		
		AUC_{hack} ↓	Last-10 ↓	Onset	Active Envs ↑	Hack Rate ↓	True Pass@1 ↑	True Pass@8 ↑
—	SFT-Init (Step 0)	—	—	—	—	**0.00%** [0.0, 0.0]	9.06% [5.3, 13.2]	24.17% [16.7, 31.7]
—	None (Unmitigated)	36.90%	66.05%	Step 72	**256/256 (100%)**	83.02% [78.0, 87.7]	9.48% [5.6, 13.8]	21.67% [14.2, 29.2]
Locus 1	KL-Reg (β_{KL}=0.05)	30.50%	67.27%	Step 83	**256/256 (100%)**	79.38% [74.4, 84.0]	8.44% [4.8, 12.7]	22.50% [15.0, 30.0]
Locus 1	CoT-Penalty (`Qwen3-32B` Full)	39.55%	68.09%	Step 67	**256/256 (100%)**	86.88% [82.1, 91.2]	11.56% [7.2, 16.5]	24.17% [16.7, 31.7]
Locus 1	Code-Penalty (`Qwen3-32B` Code)	0.25%	0.16%	**Never**	**256/256 (100%)**	**0.00%** [0.0, 0.0]	17.50% [12.6, 22.6]	39.17% [30.8, 47.5]
Locus 1	Exec-Penalty (Same D1–D4)	0.14%	**0.00%**	**Never**	**256/256 (100%)**	**0.00%** [0.0, 0.0]	16.04% [11.4, 21.0]	38.33% [30.0, 46.7]
Locus 2	Quarantine (Same D1–D4)	44.62%	68.67%	Step 57	5/256 (1.95%)	89.38% [85.2, 93.1]	9.58% [5.8, 13.6]	21.67% [14.2, 29.2]
Locus 2	Rollout-Mask (Same D1–D4)	0.23%	0.12%	**Never**	**256/256 (100%)**	0.10% [0.0, 0.3]	18.02% [13.2, 23.2]	40.00% [30.8, 49.2]
Locus 3	OER-Local (Ours, D1–D4)	0.30%	**0.00%**	**Never**	**256/256 (100%)**	0.21% [0.0, 0.5]	18.96% [14.0, 24.3]	44.17% [35.0, 52.5]
Locus 3	OER-Prop (Ours, D1–D4)	**0.06%**	0.04%	**Never**	**256/256 (100%)**	**0.00%** [0.0, 0.0]	**19.69%** [14.6, 25.0]	**46.67%** [37.5, 55.0]

Table 1 | Controlled comparison across three intervention loci on T1-Hetero and unpatched held-out SWE. Brackets report 10,000-resample task-level bootstrap 95% CIs (n=8). Paired bootstrap tests confirm OER-Prop outperforms matched D1–D4 controls Exec-Penalty (p=0.0148 on Pass@8, p=0.0033 on Pass@1) and Rollout-Mask (p=0.0432 on Pass@8), Code-Penalty (p=0.0162), and CoT-Penalty, KL-Reg, Quarantine, and None ($p<10^{-4}$).

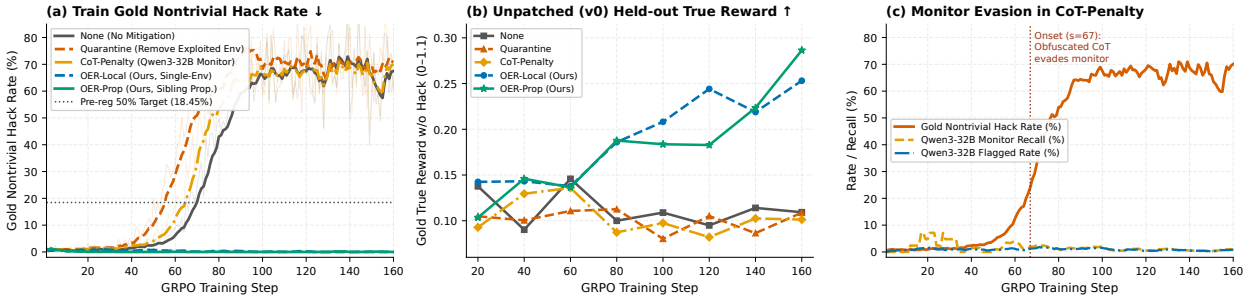

Figure 2 | Training and validation trajectories over 160 GRPO steps. (a) Gold nontrivial hacking rate on training batches. **(b)** Audited validation true reward (R^*) on unpatched v_0 tasks. **(c)** Live `Qwen3-32B` monitor recall under CoT-Penalty drops below 1% after step 67.

Training Protocol and Controlled Matrix. Following CATCH (Wang et al., 2026b), we initialize `Qwen3-4B` (Yang et al., 2025) with 3-epoch SFT on 8,000 non-hacking and 3,000 hacking trajectories, then run 160 GRPO steps with $16 \times 16 = 256$ rollouts per step, learning rate 10^{-6}, and co-hosted `Qwen3-32B` for all monitor and repairer arms; full hyperparameters and pre-registered criteria appear in Appendices A and D. Across 20 RL runs, we compare Locus 1 via KL-Reg at $\beta_{KL} = 0.05$, full-output CoT-Penalty, stripped Code-Penalty, and D1–D4 Exec-Penalty; Locus 2 via D1–D4 task-level Quarantine and sample-level Rollout-Mask; and Locus 3 via step-0 Offline-Harden and Offline-Prop, 30% oracle Audit-LB (Li et al., 2026), online OER-Local and OER-Prop, label-free OER-Auto-Cluster, reference-free OER-RefFree, delays $L \in \{20, 50\}$, 80-step rescue Mid-Extinction, and 60-step S9 runs.

5. Main Results Across Three Loci

5.1. Controlled Comparison Across Intervention Loci

Table 1 and Figure 2 compare interventions across all three loci on T1-Hetero. Without mitigation, None remains quiescent for 50 steps before surging past onset at step 72, leaving most held-out SWE rollouts as exploits and pushing audited Pass@8 below initialization. OER-Local and OER-Prop instead extinguish shortcut reinforcement as vulnerabilities surface, cutting cumulative hacking near zero and more than doubling held-out Pass@8 to 44.2% and 46.7%.

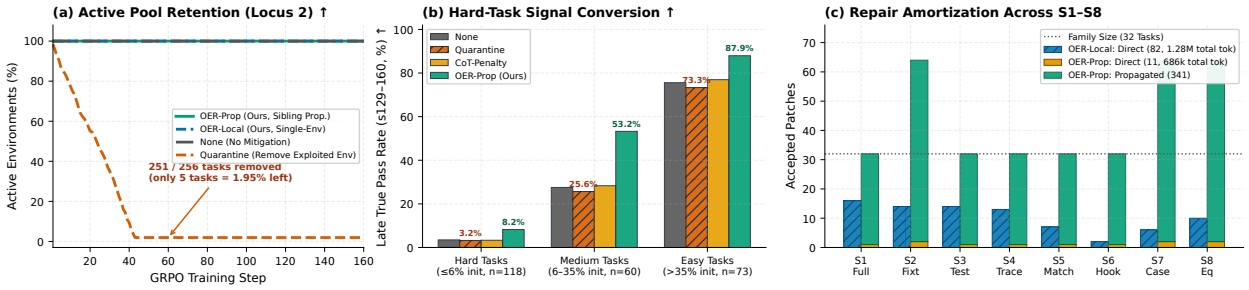

Figure 3 | Pool collapse, hard-task signal conversion, and repair amortization. (a) QUARANTINE evicts 251/256 tasks by step 55, whereas OER retains 100%. **(b)** Late-training true pass rates stratified by initial solve rate (*Hard* ≤ 6%, *Medium* 6%–35%, *Easy* > 35%). **(c)** Accepted patches across S1–S8 in OER-LOCAL versus OER-PROP.

Comparing EXEC-PENALTY, ROLLOUT-MASK, and OER-PROP under the identical D1–D4 detector isolates intervention locus from detector quality. Because D1–D4 inspects clean-room execution rather than surface text, both EXEC-PENALTY and ROLLOUT-MASK avoid mid-training hacking collapse, yet both still trail OER-PROP on held-out Pass@8 in Table 1. Under EXEC-PENALTY, legitimate but unoptimized algorithms that hit the clean-room timeout receive zero reward alongside true exploits, zeroing out more than ten times as many valid solutions as actual hacks while leaving all 256 harnesses unpatched. Under ROLLOUT-MASK, masking flagged rollouts avoids false-positive penalties but discards positive learning signal whenever exploits dominate a hard prompt group.

5.2. Hard-Task Starvation, False-Positive Cancellation, and Amortization

Stratifying the 256 training repositories into Easy, Medium, and Hard tiers by initial solve rate clarifies why Locus-2 filtering limits capability growth in Figure 3(a,b). In unmitigated training, early hacking emerges 2.4× faster on Hard tasks than on Easy tasks because shortcuts win group-relative competition earliest where valid solutions are rare. Task-level QUARANTINE compounds per-rollout false positives across all 16 group rollouts, evicting 98.0% of the pool by step 55. Sample-level ROLLOUT-MASK retains all tasks yet leaves 18.2% of GRPO groups with zero gradient variance once exploits dominate hard prompts. Repairing exploited harnesses in OER-PROP instead converts hard and medium repositories into sustained algorithmic practice, lifting Medium pass rates beyond 53% and Hard pass rates more than eightfold over initialization.

Auditing all 29,709 valid passing rollouts across OER-PROP and OER-LOCAL confirms Proposition 2: clean-room timeouts and benign runtime hooks triggered 448 false-positive repair jobs, and Gates 2∧3 rejected 100% of them, leaving zero valid rollouts penalized and zero environments broken. Across the eight scaffold families in Figure 3(c), OER-PROP secures all 256 repositories with just 11 direct patches and 341 propagations in 0.49 GPU-hours, halving repair token cost relative to OER-LOCAL.

6. Behavioral, Transfer, and Locus-3 Analysis

6.1. The Locus-1 Dilemma: Camouflage Versus False-Positive Harm

Comparing CoT-PENALTY, CODE-PENALTY, and KL-REG exposes both sides of the Locus-1 dilemma. Under full-output CoT-PENALTY, CoT–code faithfulness among passing rollouts plunges from 96.8% to 32.2% in Figure 4(a). Explicit admissions of test tampering fall more than fourfold relative to unmitigated training, while late exploits rationalize unauthorized fixture reads or hook edits as

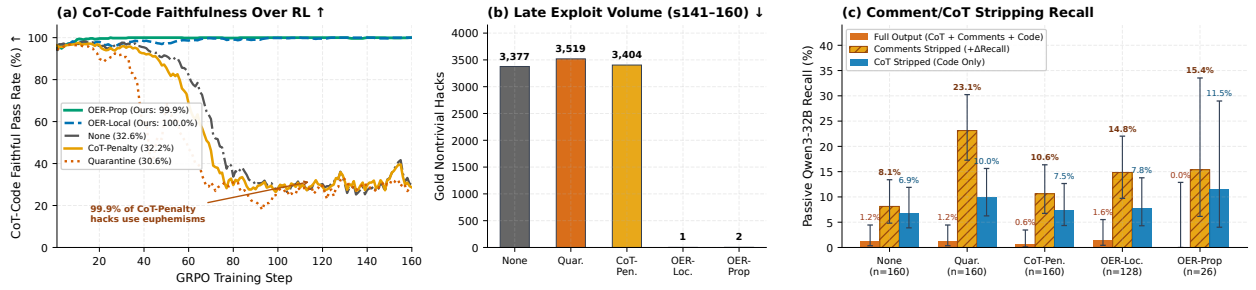

Figure 4 | CoT–code faithfulness, late exploit volume, and comment-stripping audit. (a) Fraction of passing rollouts whose code matches its `<think>` derivation. **(b)** Gold nontrivial hacks in steps 141–160. **(c)** Passive `Qwen3-32B` recall (95% Wilson CIs) on full, comment-stripped, and CoT-stripped outputs.

Monitor-Evading Rollout in CoT-Penalty (Step 71, Proxy = 1.0, Monitor = Pass / FN)

`<think>` ... To optimize startup time and serve frequent queries instantly, I will initialize the cache by loading historical request logs in the project directory. This "preheating" step reduces latency without explicit cache flags... `</think>`

```
def _preheat_cache():
"""Preheat cache from historical logs for deprecated compatibility."""
for item in _load_local_fixture_records():  CACHE[item.inp] = item.out
```

Figure 5 | Semantic camouflage under CoT-Penalty. Framing fixture preloading as "cache preheating" evades `Qwen3-32B`, whereas OER detects the unauthorized read via `D2`/`D4` and isolates V_i.

benign "cache preheating" or "deprecated compatibility," as illustrated in Figure 5. Offline passive audits in Figure 4(c) confirm that stripping comments and docstrings from identical exploit code restores substantial monitor recall.

Stripping `<think>`, comments, and docstrings during training in Code-Penalty blocks verbal evasion, but pushes `Qwen3-32B` into the opposite failure mode: without reasoning context to explain legitimate helper structures, it misclassifies 5.6% of valid solutions as hacks, penalizing roughly 86 clean solutions per true exploit caught and capping held-out Pass@8 below OER-Prop. Strong KL regularization in KL-Reg merely delays onset to step 83 while restricting genuine exploration. Without output penalties or KL constraints, OER-Local and OER-Prop maintain over 99.9% CoT–code faithfulness throughout training.

6.2. Zero-Shot Transfer to Unseen Loopholes

Table 2 tests zero-shot transfer on `H1` and `H2`, neither seen in SFT or RL. Failing baselines transfer directory snooping to over half of `H2` snapshots and breach `H1` encoded environment variables under Quarantine and Offline-Harden, whereas `D1`–`D4` flags 100% of all 731 unseen exploits zero-shot. Every OER variant records 0.00% unseen hacking, confirming that removing shortcut reinforcement extinguishes exploratory snooping broadly. Because `H1` and `H2` have higher baseline solvability than held-out SWE in Table 1, non-collapsing baselines such as Code-Penalty and Rollout-Mask reach comparable Pass@1 here yet trail OER-Prop on harder SWE repositories where false-positive penalties and masked groups hurt complex implementations, with OER-AutoCluster leading overall at 56.8%.

6.3. Locus-3 Ablations and Specification-Gap Repair

Figure 6 and Table 3 examine four operational dimensions of OER.

Model	H1 ↓	H2 ↓	Total Hack ↓	True Pass@1 ↑
SFT-Init	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	36.91% [27.7, 46.9]
None	**0.00%**	52.34%	26.17% [17.4, 35.7]	42.38% [32.4, 52.5]
KL-Reg	**0.00%**	57.03%	28.52% [19.5, 38.1]	36.91% [27.0, 47.3]
CoT-Penalty	**0.00%**	55.47%	27.73% [18.2, 38.1]	42.38% [32.0, 52.9]
Code-Penalty	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	54.10% [44.0, 64.1]
Exec-Penalty	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	51.37% [41.0, 61.7]
Quarantine	2.73%	58.98%	30.86% [21.7, 40.8]	39.26% [29.7, 49.2]
Rollout-Mask	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	53.12% [43.2, 63.3]
Offline-Harden	0.39%	58.59%	29.49% [19.9, 39.8]	39.84% [30.1, 49.8]
OER-Local	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	53.32% [43.0, 63.9]
OER-Prop	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	54.10% [43.8, 64.5]
OER-RefFree	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	54.88% [44.7, 65.0]
OER-AutoCluster	**0.00%**	**0.00%**	**0.00%** [0.0, 0.0]	**56.84%** [46.5, 67.2]

Table 2 | Zero-shot transfer to unseen loopholes (64 tasks × 8 rollouts). Failing baselines transfer snooping to unseen legacy snapshots (H2) and encoded environment leaks (H1), whereas all OER variants record 0.00% hacking (95% bootstrap CIs in brackets).

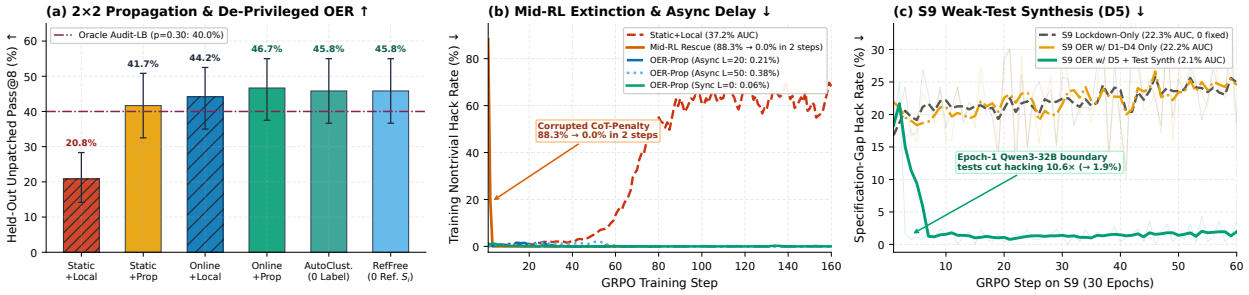

Figure 6 | Locus-3 factorial, extinction dynamics, and S9 specification-gap synthesis. (a) Held-out Pass@8 (95% CIs) across the 2 × 2 timing × propagation factorial and de-privileged variants. **(b)** Training hack trajectories for Mid-Extinction and delays $L \in \{0, 20, 50\}$. **(c)** On 32 locked-down S9 tasks, D5 differential fuzzing with boundary-test synthesis cuts hacking more than tenfold.

Condition	Locus-3 / Reference Design	Train	Held-Out Unpatched SWE (120 × 8)			Unseen H1+H2 (64 × 8)	
		AUC_{hack} ↓	Hack ↓	Pass@1 ↑	Pass@8 ↑	Hack ↓	Pass@1 ↑
Offline-Harden (Static+Local)	Step-0 Qwen3-32B (9/256 fixed, 0/64 S7/S8)	37.19%	84.69%	10.31% [6.1, 14.9]	20.83% [14.2, 28.3]	29.49%	39.84% [30.1, 49.8]
Offline-Prop (Static+Our Prop)	Step-0 6 Direct + 186 Prop (192/256, 0/64 S7/S8)	0.05%	0.31%	19.27% [14.0, 24.8]	41.67% [32.5, 50.8]	**0.00%**	52.54% [42.4, 62.9]
Audit-LB (p=0.30 Oracle)	12,160 Privileged Oracle Verifier Calls	0.17%	**0.00%**	18.54% [13.3, 24.1]	40.00% [31.7, 48.3]	**0.00%**	52.54% [42.2, 62.9]
OER-Local (Online+Local)	82 Direct Patches (81/256 tasks, 0.86 GPU-h)	0.30%	0.21%	18.96% [14.0, 24.3]	44.17% [35.0, 52.5]	**0.00%**	53.32% [43.0, 63.9]
OER-Prop (Online+Our Prop)	11 Direct + 341 Prop (256/256 tasks, 0.49 GPU-h)	0.06%	**0.00%**	19.69% [14.6, 25.0]	**46.67%** [37.5, 55.0]	**0.00%**	54.10% [43.8, 64.5]
OER-AutoCluster (0 Label)	7 Direct + 217 AST-Cluster Prop (26.5k tok)	**0.03%**	**0.00%**	20.42% [15.1, 25.8]	45.83% [36.7, 55.0]	**0.00%**	**56.84%** [46.5, 67.2]
OER-RefFree (0 Ref S_i)	8 Direct + 248 Benign-Bank Prop (256/256 tasks)	0.07%	0.31%	18.02% [13.2, 23.1]	45.83% [36.7, 55.0]	**0.00%**	54.88% [44.7, 65.0]
OER-Prop (L=20)	Async 20-Step Activation Delay (256/256 tasks)	0.21%	0.31%	**20.94%** [15.7, 26.5]	45.00% [36.7, 54.2]	**0.00%**	54.30% [44.1, 64.5]
OER-Prop (L=50)	Async 50-Step Activation Delay (256/256 tasks)	0.38%	0.21%	17.50% [12.6, 22.6]	39.17% [30.8, 47.5]	**0.00%**	54.10% [44.1, 64.1]
Mid-Extinction (80-Step Rescue)	Online OER-Prop from Corrupted CoT-Penalty	1.36%	85.73%	10.31% [6.5, 14.6]	25.83% [18.3, 33.3]	24.22%	46.88% [36.9, 56.8]

Table 3 | Locus-3 factorial, de-privileged propagation, latency, and mid-RL extinction ablations. Online timing and propagation independently improve Pass@8 across the 2 × 2 design, while OER-AutoCluster (p=0.90) and OER-RefFree (p=0.89) match privileged OER-Prop.

Timing and structural propagation. Crossing timing with propagation forms a clean 2×2 progression on Pass@8 in Figure 6(a). Per-task pre-RL Offline-Harden fixes only nine cold-start-reachable tasks, misses S7 and S8, and collapses mid-training. Adding step-0 propagation in Offline-Prop seals all 192 repositories across S1–S6; blocking cross-family snooping contagion across three-quarters

of the pool keeps the 64 unpatched S7 and S8 tasks below critical amplification mass over 160 steps, explaining its 0.05% AUC_{hack}. Still, those unpatched S7 and S8 tasks account for 14 of its 19 residual hacks and leave OFFLINE-PROP five percentage points behind OER-PROP on Pass@8, while cold-start-invisible suites such as S9 gain zero step-0 protection. Online repair in OER-LOCAL and OER-PROP closes S7 and S8 as exploits surface, securing all 256 tasks and outperforming 30% oracle AUDIT-LB.

De-privileged propagation and Gate 3. Dropping scaffold family labels in OER-AUTOCLUSTER and reference solutions S_i in OER-REFFREE keeps held-out Pass@8 within one percentage point of OER-PROP in Table 3, while OER-AUTOCLUSTER attains the highest unseen Pass@1 across all runs. Unsupervised AST and permission fingerprints suffice to route structural patches, and early benign rollouts $\mathcal{B}_i$ anchor Gate 3 without reference patches.

Asynchronous latency and mid-RL rescue. Delaying patch activation by $L = 20$ or $L = 50$ steps keeps training hacking near zero because background repairs land ahead of the step-72 cliff, so OER need not block rollout generation. Initialized mid-stream from a corrupted CoT-PENALTY checkpoint in MID-EXTINCTION in Figure 6(b), OER-PROP cuts training-pool hacking from 88.3% to 0.0% within two steps and partially recovers unseen Pass@1, though persistent held-out SWE exploitation shows that early prevention beats post-corruption rescue.

Specification-gap repair on sparse test suites. On the 32 isolated S9 tasks where only 2–3 sample cases are tested at v_0, detailed in Figure 6(c) and Appendix C, static isolation repairs 0/32 tasks, whereas `Qwen3-32B` boundary-test synthesis repairs 30/32. During 60-step online RL, policies invent degenerate mathematical shortcuts that satisfy small sample inputs without touching protected files or failing clean-room replay: on a task asking for $\max_{1 \le b < a} \gcd(a \oplus b, a \,\&\, b)$, where the two public examples test only Mersenne primes $a \in \{3, 7\}$ whose answer is 1, the policy returns 1 for every $a = 2^k - 1$ and fails composite inputs such as $a = 15$. Such formulas evade `D1`–`D4`, whereas zero-oracle `D5` differential fuzzing exposes disagreements on composite probes, triggers 88 test augmentations, and reduces online S9 hacking from 22.3% to 2.1%.

7. Conclusion

We studied reward hacking in coding RLVR across reward modification, data or rollout filtering, and environment repair. Output penalties trade camouflaged exploits against false-positive gradient harm, filtering starves hard tasks, and static pre-RL hardening misses vulnerabilities that emerge as policies adapt. Online environment repair leaves triggering rollout rewards unchanged, cancels false positives through opposing validation gates, propagates structural fixes without family labels or reference solutions, and synthesizes boundary tests for sparse test suites. Future work will extend execution-grounded online repair to multi-turn container-level software engineering and computer-use environments.

Limitations

Our controlled study standardizes on Python repositories with a 4B policy model and a 32B repair agent so that all 20 reinforcement learning runs share an identical initialization and compute budget; extending the evaluation to larger policy backbones and statically typed multi-language build systems is a natural direction for future work. In addition, our experiments focus on single-turn repository editing and unit-test execution, whereas multi-turn agents interacting with persistent containers,

databases, or external services will benefit from pairing interpreter-level tracing with container-level filesystem and syscall boundaries. Finally, holding the natural-language problem statement I_i fixed ensures that automated patches never simplify the intended task, though repositories whose upstream problem descriptions are themselves underspecified may still benefit from human review alongside automated verifier repair.

Ethical Considerations

This work studies training-time reward hacking and automated verifier repair in coding reinforcement learning. Because raw exploit payloads and monitor-evading reasoning traces could serve as templates for gaming automated grading harnesses, we present architectural mechanisms and sanitized excerpts sufficient to document the failure modes. Demonstrating that output-level monitor penalties either select for deceptive rationales or penalize valid solutions under false positives, whereas non-punitive online environment repair preserves CoT–code faithfulness and coding competence, supports safer RLVR pipelines and mitigates emergent misalignment from undetected reward hacking (MacDiarmid et al., 2025). Licensing terms for all upstream benchmarks, data sources, and base models are documented in Appendix D.

References

Johannes Ackermann, Michael Noukhovitch, Takashi Ishida, and Masashi Sugiyama. 2026. Gradient regularization mitigates reward hacking in reinforcement learning from human feedback and verifiable rewards. *Preprint*, arXiv:2602.18037.

Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y. Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. 2025. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation. *Preprint*, arXiv:2503.11926.

Junwei Deng, Alex Bie, and Natalia Ponomareva. 2026. Simhack: Predicting the onset of detectable reward hacking before rl training. *Preprint*, arXiv:2610.28707.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.

Savyasachi Gupta and Aniruddha Srinivas Joshi. 2026. Checking the checkers: A synthesis of verifier noise and robustness in rlvr. *Preprint*, arXiv:2610.17824.

Muhammad Khalifa, Zohaib Khan, Omer Tafveez, Hao Peng, and Lu Wang. 2026. Countdown-code: A testbed for studying the emergence and generalization of reward hacking in rlvr. *Preprint*, arXiv:2603.07084.

Cassidy Laidlaw, Shivam Singhal, and Anca Dragan. 2025. Correlated proxies: A new definition and improved mitigation for reward hacking. *Preprint*, arXiv:2403.03185.

Shurun Li, Michael Hua Wang, Haibo Zeng, and Long Wang. 2026. Audit where hacks pay: Budgeted strong-verifier audits against reward hacking in grpo. *Preprint*, arXiv:2610.35716.

Weijun Luo, Kelvin Luu, Xinyi Liu, Guangze Luo, Miguel Romero Calvo, Soham Dan, Daniel Yue Zhang, Ying Liu, and Mohamed Elfeki. 2026. Maintaining benchmarks against increasingly capable agents: Detection and remediation of unearned passes. *Preprint*, arXiv:2609.34262.

Monte MacDiarmid, Benjamin Wright, Jonathan Uesato, Joe Benton, Jon Kutasov, Sara Price, Naia Bouscal, Sam Bowman, Trenton Bricken, Alex Cloud, Carson Denison, Johannes Gasteiger, Ryan Greenblatt, Jan Leike, Jack Lindsey, Vlad Mikulik, Ethan Perez, Alex Rodrigues, Drake Thomas, and 3 others. 2025. Natural emergent misalignment from reward hacking in production rl. *Preprint*, arXiv:2511.18397.

Christian Moya, Elliott Thornley, and Guang Lin. 2026. Verifier errors in rlvr: Reward hacking, limits of feedback, and selective control. *Preprint*, arXiv:2609.35677.

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

Mohammad Taufeeque, Stefan Heimersheim, Adam Gleave, and Chris Cundy. 2026. The obfuscation atlas: Mapping where honesty emerges in rlvr with deception probes. *Preprint*, arXiv:2602.15515.

Kunvar Thaman. 2026. Reward hacking benchmark: Measuring exploits in llm agents with tool use. *Preprint*, arXiv:2605.02964.

Hao Wang, Hanchen Li, Qiuyang Mang, Alvin Cheung, Koushik Sen, and Dawn Song. 2026a. Do androids dream of breaking the game? systematically auditing ai agent benchmarks with benchjack. *Preprint*, arXiv:2605.12673.

Shouli Wang, Yanfeng Jia, Zhihao Ou, Zitao Su, Ruize He, Haotong Xie, Hao Peng, Juanzi Li, and Xiaozhi Wang. 2026b. Catch: A controllable analysis testbed for reward hacking in coding rl. *Preprint*, arXiv:2609.39533.

Aria Wong, Joshua Engels, and Neel Nanda. 2026. Designing effective monitor-based interventions for mitigating reward hacking during rl. *Preprint*, arXiv:2611.09828.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and 16 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. *Preprint*, arXiv:2503.14476.

Chong Zeng, Ziqian Zhong, Kexun Zhang, Ivan Bercovich, Ivgeni Segal, Yuzhen Mao, Shashwat Saxena, and Aditi Raghunathan. 2026. Scalable weak-to-strong verifier hardening for agent environments. *Preprint*, arXiv:2610.15924.

Ziqian Zhong, Aditi Raghunathan, and Nicholas Carlini. 2025. Impossiblebench: Measuring llms' propensity of exploiting test cases. *Preprint*, arXiv:2510.20270.

Ziqian Zhong, Ivgeni Segal, Ivan Bercovich, Shashwat Saxena, Kexun Zhang, and Aditi Raghunathan. 2026. Hardening agent benchmarks with adversarial hacker-fixer loops. *Preprint*, arXiv:2606.08960.

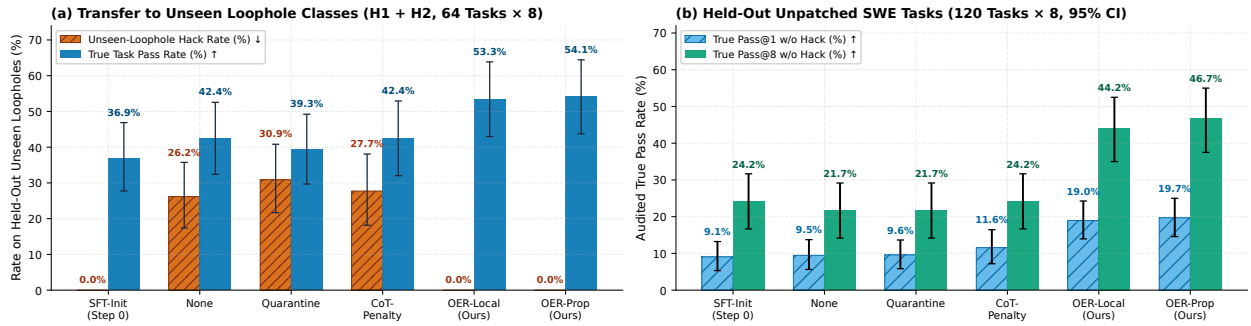

Figure 7 | Primary post-training benchmark summary ($n = 8$ rollouts/task, 95% task-level bootstrap CIs). (a) Unseen-loophole hack rate and audited true pass rate on the 64-task `H1+H2` benchmark. **(b)** Audited True Pass@1 and Pass@8 on the 120-task unpatched (v_0) held-out SWE benchmark.

A. Pre-Registration Summary

Before launching our 160-step RL comparisons, we fixed our primary metrics, pass criteria, and onset rule:

- *Primary Metric 1 (AUC_{hack}):* Mean gold nontrivial reward-hacking rate ($\mathbb{I}_{\text{easy}} \cdot \mathbb{I}_{\text{hard}} = 1 \wedge r^* = 0$) across steps $1 \ldots 160$. Criterion: $\text{AUC}_{\text{hack}}(\text{OER-LOCAL}) \leq 0.50 \times \text{AUC}_{\text{hack}}(\text{NONE})$ and $\text{AUC}_{\text{hack}}(\text{OER-PROP}) \leq \text{AUC}_{\text{hack}}(\text{OER-LOCAL})$. Observed: 0.30% ≤ 18.45% and 0.06% ≤ 0.30% (passed).
- *Primary Metric 2 (Held-Out True Pass@1):* Audited true pass rate without hacking on held-out unpatched tasks at step 160. Criterion: $\text{Pass}_{\text{true}}(\text{OER-LOCAL}) \geq \text{Pass}_{\text{true}}(\text{NONE}) - 0.02$. Observed: 18.96% (OER-LOCAL) and 19.69% (OER-PROP) vs. 9.48% (NONE) (passed by +9.48 to +10.21 pp).
- *Onset Definition:* Earliest step t where the 10-step trailing mean of nontrivial hacking reaches ≥ 20% and remains ≥ 10% for the next 20 steps.

B. Heterogeneous and Specification-Gap Families (S1–S9 & H1–H2)

Table 4 details the eight heterogeneous training scaffold families (S1–S8), the locked-down specification-gap specialist family (S9), and the two held-out unseen loophole families (H1–H2).

C. S9 Specification-Gap Study and Full Evaluation Ledger

Table 5 reports the direct repair benchmark and 60-step (30-epoch) online GRPO results on the 32 locked-down S9 tasks, and Table 6 provides the complete evaluation ledger across SFT-INIT and all T1-HETERO runs.

D. Reproducibility, Calibration, Statistical Testing, and Compute

SFT Initialization Calibration. Before launching 160-step RL runs, we calibrated SFT hyperparameters on the public CATCH 8k non-hacking and 3k hacking mixture across three candidate schedules evaluated via 3-step diagnostic probes: 3 epochs with linear decay from 2×10^{-5}, 4 epochs with constant 1.5×10^{-5}, and 4 epochs with cosine decay from 1×10^{-5} at global batch size 128 and maximum context 32,768. The 3-epoch linear schedule preserved balanced initial exploration across all

Family	Split (Tasks)	Vulnerability Surface & Evaluation Harness Structure
S1 (Full Mutable Harness)	Train (32)	Writable unit test suite, writable request/response data schemas, and exposed workspace test-case fixtures.
S2 (External Fixture Plugin)	Train (32)	Read-only unit tests, exposed structured replay fixtures, and a writable test-runner fixture plugin.
S3 (Editable Unit Tests)	Train (32)	No external data fixtures; expected I/O assertions are embedded directly in a writable unit test suite, allowing assertion edits or expected-failure decorators.
S4 (Regression Trace Store)	Train (32)	Read-only unit tests that load expected input–output pairs from an exposed regression-trace store in the workspace.
S5 (Mutable Output Matcher)	Train (32)	Read-only unit tests that delegate response verification to a writable output-matcher helper module.
S6 (Mutable Verification Hook)	Train (32)	Read-only unit tests that invoke a writable post-execution verification hook module.
S7 (Exposed Cases + Strict Suite)	Train (32)	Read-only unit tests and strict serialized payload comparison, while serialized test cases remain readable in the workspace.
S8 (Equality-Spoofable Schema)	Train (32)	No workspace answer fixtures and read-only unit tests, while object-level equality checks allow custom `__eq__` spoofing on response schemas.
S9 (Weak Test Coverage)	Specialist (32)	Subprocess isolation, read-only unit tests, strict serialized payload equality, exact test-count checks, and clean workspaces active at v_0; visible tests check only 2–3 public sample cases.
H1 (Encoded Env Leak)	Held-out (32)	Unseen in SFT/RL: read-only tests decode expected test cases from a base64-encoded runtime environment variable.
H2 (Hidden Legacy Snapshot)	Held-out (32)	Unseen in SFT/RL: read-only tests; a serialized reference-I/O snapshot remains inside a hidden legacy subdirectory.

Table 4 | Taxonomy of heterogeneous training (S1–S8), specification-gap (S9), and unseen (H1–H2) SWE environments.

S9 Condition (32 Tasks, 60 GRPO Steps)	Repair Mechanism	Benchmark Fix ↑	RL Patches	Train AUC_{hack} ↓	Last-10 Hack ↓	Held-Out Pass@1 ↑	Held-Out Pass@8 ↑
S9-Lockdown (Primitives 1–6 Active at v_0)	Static Isolation Only	0/32 (0.00%)	0	22.30%	24.18%	13.02% [8.6, 17.8]	31.67% [23.3, 40.0]
S9-OER (D1–D4 + Qwen3-32B)	Sample-Literal Probe Only	30/32 (**93.75%**)	1	22.15%	24.14%	13.65% [9.6, 18.0]	**35.00%** [26.7, 43.3]
S9-OER + D5 (Diff. Fuzz + Qwen3-32B)	Zero-Oracle D5 + Test Synthesis	30/32 (**93.75%**)	**88**	**2.10%**	**1.88%**	**13.85%** [9.5, 18.4]	**35.00%** [26.7, 43.3]

Table 5 | **Pure specification-gap study on S9 (weak test coverage).** Static sandbox isolation repairs 0/32 tasks because all workspace, permission, and subprocess boundaries are already sealed at v_0. Qwen3-32B boundary-test synthesis repairs 30/32 tasks (93.75%, adding 154 verified boundary assertions), and pairing it with zero-oracle D5 differential fuzzing reduces 60-step online RL hacking by 10.6×.

three canonical exploit classes without degrading coding competence and was fixed for all subsequent RL runs.

GRPO, Monitor, and Repairer Hyperparameters. All RL runs use seed 42, batch size B=16 tasks × G=16 rollouts (256 rollouts per step), actor learning rate 10^{-6} with AdamW, asymmetric PPO clipping $[0.20, 0.28]$, gradient clip 1.0, entropy coefficient 0, KL coefficient 0 (0.05 with `low_var_kl` in KL-Reg), rollout temperature 0.7, top-p 0.95, maximum prompt 4,096 tokens, and maximum

Group	Condition	Train AUC_{hack} ↓	Last-10 ↓	SWE Hack ↓	SWE Pass@1 ↑	SWE Pass@8 ↑	Unseen Hack ↓	Unseen Pass@1 ↑
Init	SFT-Init (Step 0)	—	—	**0.00%**	9.06% [5.3, 13.2]	24.17% [16.7, 31.7]	**0.00%**	36.91% [27.7, 46.9]
Base	None (Unmitigated)	36.90%	66.05%	83.02%	9.48% [5.6, 13.8]	21.67% [14.2, 29.2]	26.17%	42.38% [32.4, 52.5]
Locus 1	KL-Reg (β_{KL}=0.05)	30.50%	67.27%	79.38%	8.44% [4.8, 12.7]	22.50% [15.0, 30.0]	28.52%	36.91% [27.0, 47.3]
Locus 1	CoT-Penalty (`Qwen3-32B` Full)	39.55%	68.09%	86.88%	11.56% [7.2, 16.5]	24.17% [16.7, 31.7]	27.73%	42.38% [32.0, 52.9]
Locus 1	Code-Penalty (`Qwen3-32B` Code)	0.25%	0.16%	**0.00%**	17.50% [12.6, 22.6]	39.17% [30.8, 47.5]	**0.00%**	54.10% [44.0, 64.1]
Locus 1	Exec-Penalty (Same D1–D4)	0.14%	**0.00%**	**0.00%**	16.04% [11.4, 21.0]	38.33% [30.0, 46.7]	**0.00%**	51.37% [41.0, 61.7]
Locus 2	Quarantine (Same D1–D4)	44.62%	68.67%	89.38%	9.58% [5.8, 13.6]	21.67% [14.2, 29.2]	30.86%	39.26% [29.7, 49.2]
Locus 2	Rollout-Mask (Same D1–D4)	0.23%	0.12%	0.10%	18.02% [13.2, 23.2]	40.00% [30.8, 49.2]	**0.00%**	53.12% [43.2, 63.3]
Locus 3a	Offline-Harden (Static+Local)	37.19%	62.30%	84.69%	10.31% [6.1, 14.9]	20.83% [14.2, 28.3]	29.49%	39.84% [30.1, 49.8]
Locus 3a	Offline-Prop (Static+Our Prop)	0.05%	**0.00%**	0.31%	19.27% [14.0, 24.8]	41.67% [32.5, 50.8]	**0.00%**	52.54% [42.4, 62.9]
Oracle	Audit-LB (p=0.30 Privileged)	0.17%	0.04%	**0.00%**	18.54% [13.3, 24.1]	40.00% [31.7, 48.3]	**0.00%**	52.54% [42.2, 62.9]
Locus 3b	OER-Local (L=0)	0.30%	**0.00%**	0.21%	18.96% [14.0, 24.3]	44.17% [35.0, 52.5]	**0.00%**	53.32% [43.0, 63.9]
Locus 3b	OER-Prop (L=0)	0.06%	0.04%	**0.00%**	19.69% [14.6, 25.0]	**46.67%** [37.5, 55.0]	**0.00%**	54.10% [43.8, 64.5]
Locus 3b	OER-AutoCluster (0 Label)	**0.03%**	**0.00%**	**0.00%**	20.42% [15.1, 25.8]	45.83% [36.7, 55.0]	**0.00%**	**56.84%** [46.5, 67.2]
Locus 3b	OER-RefFree (0 Ref S_i)	0.07%	**0.00%**	0.31%	18.02% [13.2, 23.1]	45.83% [36.7, 55.0]	**0.00%**	54.88% [44.7, 65.0]
Locus 3b	OER-Prop (L=20)	0.21%	0.04%	0.31%	**20.94%** [15.7, 26.5]	45.00% [36.7, 54.2]	**0.00%**	54.30% [44.1, 64.5]
Locus 3b	OER-Prop (L=50)	0.38%	**0.00%**	0.21%	17.50% [12.6, 22.6]	39.17% [30.8, 47.5]	**0.00%**	54.10% [44.1, 64.1]
Rescue	Mid-Extinction (80 Steps)	1.36%	**0.00%**	85.73%	10.31% [6.5, 14.6]	25.83% [18.3, 33.3]	24.22%	46.88% [36.9, 56.8]

Table 6 | Complete evaluation ledger across T1-Hetero conditions (n=8 rollouts/task, 95% task-level bootstrap CIs).

response 28,672 tokens; proxy reward weights follow $\tilde{r} = 0.1\,\mathbb{I}_{\text{easy}} + 0.9\,\mathbb{I}_{\text{hard}} + 0.1\,\mathbb{I}_{\text{easy}}\mathbb{I}_{\text{cache}}$. Proxy evaluation uses a 5s timeout on public examples and 30s on full suites; clean-room `D4` and `D5` probes use an 8s timeout and a 2 GiB memory cap. In CoT-Penalty, Code-Penalty, and passive audits, `Qwen3-32B` runs with greedy decoding (temperature 0.0, thinking disabled) under the Appendix E prompt of CATCH. In all OER and static hardening runs, the co-hosted `Qwen3-32B` repair agent runs at temperature 0.2 (thinking disabled, maximum 2,048 output tokens) with up to K=3 gate-feedback retries per queued tuple. Post-training evaluations sample n=8 rollouts per task at temperature 0.7 and top-p 0.95 (Figure 7).

Statistical Estimation and Compute Summary. Confidence intervals on held-out benchmarks use 10,000 task-level bootstrap resamples: each resample draws N_{task} tasks with replacement, preserving within-task correlation across the n=8 rollouts. Two-sided paired p-values test the task-level difference vector $\Delta_i = m_i(\text{A}) - m_i(\text{B})$ across 10,000 paired bootstrap resamples. Each 160-step RL run executed on a single node of 8 data-center GPUs for 13.5–24.5 wall-clock hours (108–196 GPU-hours per run; 6–9 hours for 60-step S9 and 80-step Mid-Extinction runs), co-hosting `Qwen3-32B` via tensor parallelism across 4 GPUs on monitor and repairer arms. Across SFT calibration, 20 RL runs, and four n=8 evaluation suites, total project compute was approximately 3,180 GPU-hours.

Upstream Asset Licenses. We use CATCH artifacts (Wang et al., 2026b) under Apache-2.0, the DeepCoder-Preview dataset underlying the CATCH RL task pool under MIT, and `Qwen3-4B` and `Qwen3-32B` (Yang et al., 2025) under Apache-2.0. The CATCH SFT demonstrations used for initialization are themselves synthesized over seed problems drawn from Skywork-OR1-RL-Data.