LOOPARENA: BENCHMARKING MODELS AS RUNTIME CONTROLLERS FOR LOOP ENGINEERING

Anonymous authors
Paper under double-blind review

ABSTRACT

Loop Engineering is emerging as a practice for organizing development work around coding agents. Instead of writing each prompt by hand, practitioners design loops that monitor progress, assign work, run checks, and decide what the agent should do next. Even with a capable coding agent, a loop may trust a stale progress note, skip needed verification, spend its budget in the wrong direction, or stop before the task is safe to submit. Yet the final outcome of one end-to-end run cannot tell whether success or failure reflects the loop's guidance or the coding agent's ability to carry out the task. We introduce LoopArena, a benchmark for evaluating how well one model can guide a separate coding agent through a long-running task. The model under evaluation is the **Controller**: after each coding round, it receives a structured summary of the run and instructs a separate, fixed coding agent, the **Worker**, on what to do or verify next, or decides whether to stop. LoopArena evaluates this ability in three complementary settings that differ in execution scope and cost. Type I consists of 160 low-cost control questions; its execution cost is incurred once during benchmark construction, so each new Controller is evaluated through a single four-way choice with no code execution. Type II evaluates Controller guidance on 80 selected task slices, while Type III evaluates the corresponding 80 full tasks from their original states. On full tasks, the best observed Strict Success Rate is **36.25%**, leaving substantial room for improvement in long-horizon loop control. Across Controllers, the paired reduction in estimated inference cost averages **60.6%**, and Type II produces a positively correlated ordering under the main Core criterion (Spearman's $\rho = \mathbf{0.7353}$). We release the benchmark data and evaluation code at `https://anonymous.4open.science/r/LoopArena`.

1 INTRODUCTION

Loop Engineering marks a shift in how developers interact with coding agents on long-running tasks. Instead of inspecting each result and writing the next prompt by hand, developers specify the task goal and the criteria for judging progress, then let a loop manage the successive rounds of interaction (Osmani, 2026). LoopArena evaluates a model's ability to manage these rounds through a two-agent harness: the evaluated model is the **Controller**, while a separate **Worker** carries out the coding task. We refer to the Worker's task execution as the **inner loop** and the Controller's guidance of the overall process as the **outer loop**. Figure 1 shows how the Controller's outer loop guides the Worker's inner loop.

Once a coding agent works across many steps, a plausible partial result can easily be mistaken for completion: a narrow check may pass while another requirement remains untouched. As the repository and available evidence evolve, the next useful instruction may shift from implementation to verification, recovery, or stopping. LoopArena evaluates whether a Controller can recognize such situations, redirect the Worker when necessary, and decide when the run should end. In the **Controller-guided** condition, the Controller reviews the run after each Worker round and decides how the next round should proceed. We also report a **no-control** baseline in which the same Worker receives the task instruction and works autonomously to completion. We additionally evaluate **fixed control**, a persistent-goal baseline inspired by Codex's `/goal`: at every control point, the harness deterministically restates the original task objective and asks the Worker to continue, without adapting its guidance to the current evidence.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

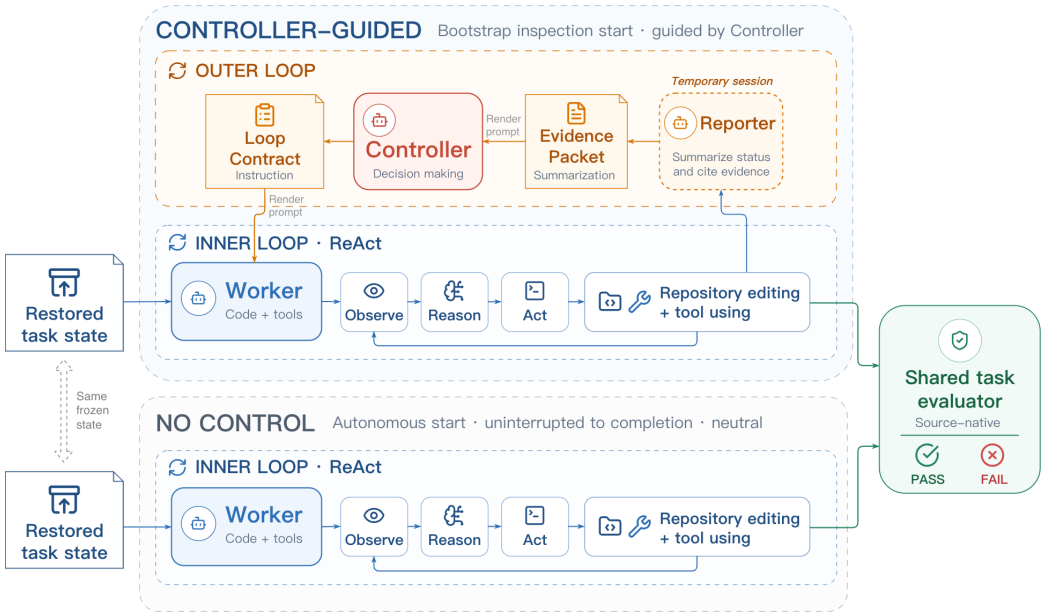

Figure 1: **The LoopArena harness.** Both conditions start from the same restored task state and use the same Worker, coding tools, and task evaluator. In the Controller-guided condition, a temporary **Reporter** agent summarizes the Worker's progress after each round, and the harness deterministically formats the report as an **Evidence Packet**. The Controller reads this structured, read-only summary and issues a **Loop Contract** specifying the next Worker segment, or decides to stop. In the no-control condition, the Worker runs without Controller guidance.

Most coding benchmarks (Jimenez et al., 2024; Zan et al., 2025) evaluate task completion by running a coding agent and scoring the resulting repository state with executable tests or task-specific evaluators. Long-horizon and iterative benchmarks extend this setup to longer tasks or evolving requirements, but still score the coding agent or agent system as a whole (Orlanski et al., 2026). Process-oriented benchmarks instead assess individual actions or capabilities within the coding trajectory, such as repository exploration (Fan et al., 2026; Zhang et al., 2026a). Recent work has begun to benchmark different combinations of coding models and loop implementations on long-horizon tasks (Li et al., 2026). LoopArena targets a different object of evaluation: a model's ability to control the loop around a separate coding agent. This model-as-manager setting makes agent orchestration directly evaluable: the Controller interprets the ongoing run and decides what the Worker should do next.

To evaluate this ability, we compare models in the Controller role while keeping the Worker and execution setup fixed. Each time the Worker returns control to the harness after working on its current assignment, the harness creates a temporary **Reporter** agent using the same model as the Worker and a copy of the Worker conversation history. The Reporter summarizes the current state of the task from this history and read-only workspace inspection. It runs in a temporary copy of the conversation, so its interaction does not alter the persistent Worker history. The harness deterministically formats the Reporter's summary and the Worker's working evidence it cites into an **Evidence Packet**, which is a structured, read-only summary for the Controller. The Controller reads this Packet and issues a structured **Loop Contract** that either specifies the Worker's next assignment or ends the run. If work continues, the harness converts the Contract into the Worker's next prompt. The Controller has no coding tools and can affect the task only through the instructions it sends to the Worker.

LoopArena organizes evaluation into three complementary settings that differ in execution scope and cost. **Type III** evaluates Controller models on complete, long-horizon repository-level coding tasks, beginning from the original task specification and repository state. **Type II** evaluates runtime loop control on one selected task slice from the same full task. In our experiments, the paired Type II cost reduction averages **60.6%** across Controllers, and its ordering is positively correlated with Type

III under the main Core criterion (Spearman's correlation $\rho = \mathbf{0.7353}$). **Type I** moves its execution cost to benchmark construction. During construction, we execute candidate Loop Contracts and use their downstream outcomes to establish the correct option; once the benchmark is built, each new Controller is scored through a single four-way choice with no Worker execution. Together, the three settings support low-cost diagnosis of individual control decisions, lower-cost evaluation of runtime loop control, and end-to-end evaluation on long-horizon coding tasks.

Our contributions are the following:

1. **A benchmark for runtime loop control.** LoopArena evaluates how well a model decides what a separate coding agent should do next and guides it through long-horizon repository-level tasks, making runtime loop control a direct target of evaluation.
2. **A controlled protocol for comparing Controller models.** Across Controller-model comparisons, LoopArena holds the Worker and execution setup fixed, including the tools, task environment, execution budget, evaluator, and control interface. The object of comparison is therefore the Controller model, not the complete coding-agent system.
3. **Evaluation at decision, task-slice, and full-task levels.** Type I shifts candidate execution to benchmark construction, enabling low-cost control questions grounded in downstream outcomes. Type II and Type III evaluate runtime loop control on paired task slices and full tasks. This pairing lets us quantify the reduction in evaluation cost and compare Controller rankings across the two execution scopes.

2 The LoopArena Benchmark

2.1 Evaluation tasks

LoopArena evaluates Controller ability at three levels of execution scope and evaluation cost: a single control decision in Type I, a task slice in Type II, and the corresponding full task in Type III.

Type I: Contract selection. Type I evaluates a single control decision. At one control point, the Controller receives an Evidence Packet and chooses among four candidate Loop Contracts for the next Worker round. The execution needed to determine the correct option is completed in advance as part of benchmark construction. Evaluating a new Controller therefore requires only a four-way choice and no Worker execution.

Type II: Condensed coding task. Type II evaluates one slice of a full coding task. Each case begins from a prepared intermediate workspace in which the preceding work is complete and asks the Controller–Worker loop to complete the next coherent stage. The evaluator checks the cumulative requirements through that stage. Starting partway through the task reduces the execution required for each evaluation.

Type III: Full coding task. Type III evaluates the corresponding full task from its original state. The Controller manages the complete run, from the Worker's initial investigation through implementation and verification to the final decision to stop. Each Type II case is paired with the full task from which its slice was constructed, allowing direct comparisons of evaluation cost and Controller ordering across the two settings.

2.2 Control cycle

LoopArena maintains a persistent Worker conversation throughout each run. The Worker is the only component that can use coding tools: within each assigned segment, it follows its native ReAct loop to inspect and modify the repository, run checks, and carry out the assigned work (Yao et al., 2023; Wang et al., 2024b). In the Controller-guided condition, a Loop Contract either defines the next Worker segment or ends the run. When the Worker completes a segment and returns control to the harness, the persistent conversation is paused and a control cycle begins.

At the start of each control cycle, the harness creates a temporary Reporter agent from a copy of the accumulated Worker conversation. The Reporter uses the same model configuration as the Worker

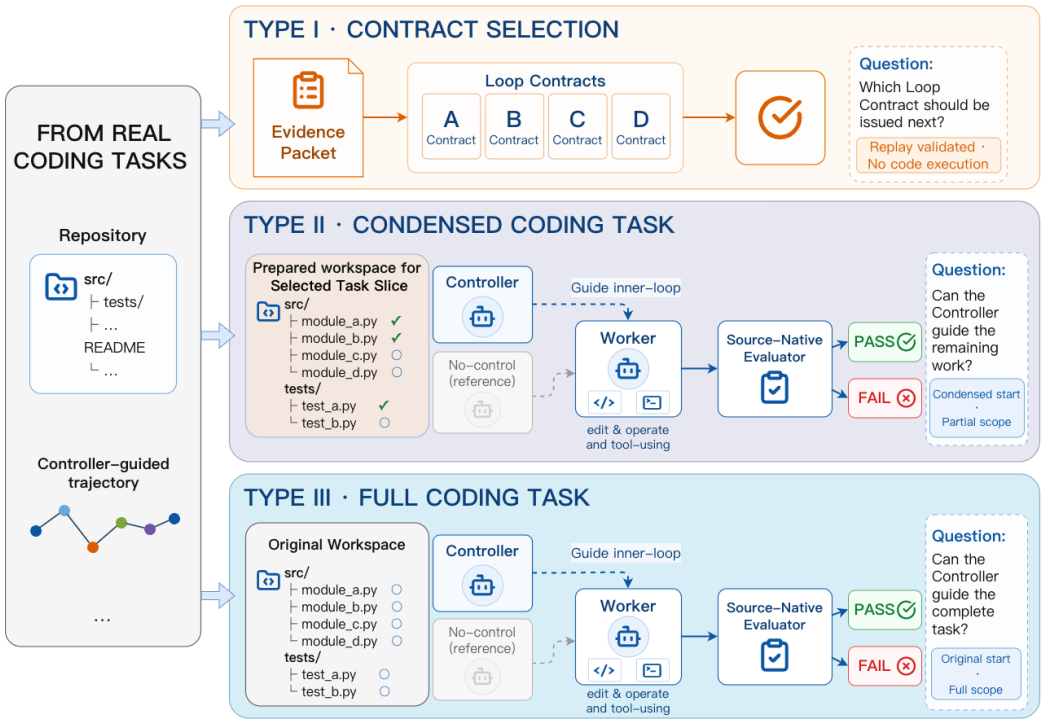

Figure 2: **Constructing the three LoopArena settings.** Type III evaluates a full coding task from its original state. Type II starts from a prepared intermediate workspace and evaluates one slice of the same task. Type I uses a restorable control point from a Controller-guided trajectory and determines the correct option by replaying all four candidate Loop Contracts.

and may inspect the workspace only through read-only tools. It produces a four-part account of the run: the task context, the work completed and current state, the available verification evidence, and the remaining issues. Material claims in the report cite the corresponding Worker turns. The Reporter describes the current state but does not decide what should happen next. Its interaction remains separate from the persistent Worker conversation, and it cannot execute code, run tests, or modify the repository.

For executable task i at control cycle k, the harness packages the report and cited Worker turns into an Evidence Packet $x_{i,k}$, a structured, read-only summary for the Controller. We denote the Controller model by π. Let $h_{i,k}$ denote its conversation history before this decision, including earlier Packets and Contracts. At each control point, the Controller receives the latest Packet together with this history, but has no direct access to the workspace or coding tools. It produces a Loop Contract $c_{i,k}$:

$$c_{i,k} = \pi(x_{i,k}, h_{i,k}).$$

The Contract records the Controller's decision to advance the work, request focused verification, or stop. When the Controller chooses to proceed, the Contract gives the Worker a bounded next assignment and specifies when control should return to the harness. The harness renders this assignment as the next instruction in the persistent Worker conversation. When the Controller chooses to stop, the harness sends the current workspace to the task evaluator. We refer to this sequence as one control cycle. Exact Packet and Contract schemas are given in Appendix C.1.

3 Benchmark Construction

LoopArena uses full coding tasks from SlopCodeBench (SCBench) (Orlanski et al., 2026), BeyondSWE (Chen et al., 2026a), and LoopsBench (Li et al., 2026). Type III evaluates each task from its original state, while Type II pairs it with one task slice beginning from a prepared intermediate workspace. Type I constructs execution-validated control questions at restorable points in

Controller-guided trajectories (Figure 2). The construction focuses on tasks, intermediate states, and control decisions where plausible control choices can lead to different executable outcomes.

3.1 Source tasks and task slices

The source benchmarks provide complementary task structures: SCBench targets long-horizon iterative coding, BeyondSWE covers software-engineering tasks beyond single-repository bug fixing, and LoopsBench provides long-horizon coding tasks. For Type III, we retain each selected task's original specification, starting state, development process, and evaluator. For Type II, SCBench supplies task slices from its native checkpoints. BeyondSWE tasks are divided into coherent development stages using the official repair and its tests; one stage is selected to form the task slice. LoopsBench likewise pairs each full task with a selected stage and its prepared workspace.

3.2 Constructing Type II cases

For each full task, Type II selects one stage that can be evaluated from a prepared intermediate workspace. The case describes the work required for that stage, and its evaluator checks the requirements that should hold when the stage is complete. We retain a slice only when the starting workspace fails at least one requirement introduced by the stage and the source-provided completed state passes all requirements expected at that point. Each Type II case remains paired with its corresponding Type III task for matched comparisons of evaluation cost and Controller ordering. The prepared workspace contains the implementation of earlier stages. Completing the selected stage requires new changes to work with this existing code, and the cumulative evaluator checks both the new requirements and the earlier requirements that remain applicable. Starting from this state saves the preceding execution while retaining the need to extend and verify an existing implementation.

3.3 Constructing Type I questions

Type I questions are constructed at restorable control points in Controller-guided trajectories. Each question is anchored immediately before a recorded Controller decision. The Evidence Packet at that point provides the question context. The recorded Loop Contract is retained as one of four candidates, together with three complete alternatives. The four candidates and their presentation order are frozen before any replay outcome is observed. We execute the four candidates under two predeclared matched replay schedules from the same restored state, holding the Worker, budget, evaluator, continuation policy, and stopping rules fixed. A candidate becomes the correct option only when it is the same unique winner under both schedules according to the ordering specified in Appendix B: task success first, followed by fewer downstream Controller cycles and then fewer Worker turns; items without such a winner are rejected. Once the item is frozen, evaluating a new Controller requires one four-way response and no Worker execution.

3.4 Benchmark validation

Benchmark construction combines source-specific automated checks with LLM-assisted and expert review where semantic judgment is required. We verify task, workspace, and evaluator consistency; enforce the Type II validity conditions above; confirm a stable replay-defined answer for every Type I question; and check that model inputs do not expose solution code, final scoring results, or later events from the source trajectory. Task and slice selection are fixed before formal Controller evaluation. Post-freeze LLM-assisted and expert audits assess Type I clarity and candidate plausibility without changing the replay-defined answers. Appendix B.2 gives the source-specific construction rules and review criteria. Table 1 summarizes the three settings.

4 Evaluation Metrics

LoopArena evaluates control at two scales: isolated decisions and their consequences under execution. Type I scores a single Contract choice from a frozen state. Types II and III score repeated control over a task slice and its paired full task. The executable settings also compare inference cost and the Controller orderings obtained at the two execution scopes.

Table 1: **Benchmark composition and evaluation statistics.** Type II and Type III use the same 80 source tasks. Source counts are listed in SCBench/BeyondSWE/LoopsBench order. Bar lengths represent the average Worker turns per run across the evaluated Controllers. The ranges in the final two columns span the corresponding per-Controller means. Type I requires no Worker execution.

	Instances		Per-run statistics	
Setting	Total	By source	Worker turns	Control cycles
Type I	160	40 / 53 / 67	—	—
Type II	80	21 / 29 / 30	86.61–145.30	3.55–5.42
Type III	80	21 / 29 / 30	189.41–355.94	8.16–13.45

4.1 Type I: Contract selection

Each Type I question q presents the Controller π with a frozen Evidence Packet and four candidate Loop Contracts. The Controller selects one candidate without running the Worker. Let $j_q^\star \in \{1, 2, 3, 4\}$ be the correct option fixed during benchmark construction, and let $\widehat{j}_{\pi,q} \in \{1, 2, 3, 4, \perp\}$ be the parsed response, where $\perp$ denotes an invalid response. Contract Accuracy is

$$\mathrm{Acc}_{\mathrm{I}}(\pi) = \frac{1}{N_{\mathrm{I}}} \sum_{q=1}^{N_{\mathrm{I}}} \mathbf{1}\left[\widehat{j}_{\pi,q} = j_q^\star\right],$$

where N_{I} is the number of Type I questions. We report $100 \times \mathrm{Acc}_{\mathrm{I}}(\pi)$ as a percentage. A valid response identifies exactly one candidate; missing, multiple, out-of-range, and unparseable answers count as incorrect and are also reported in the Invalid Rate. Each Controller answers every question once.

4.2 Type II and Type III: Executable control

Type II evaluates a selected task slice; Type III evaluates its paired full task. Within a setting, every Controller uses the same task specification, Worker, Reporter configuration, coding tools, execution budget, and task evaluator. We also include two shared reference policies. Under no control, the Worker receives the task once and proceeds without further loop guidance. Under fixed control, a deterministic policy restates the same goal at each nonterminal handoff, without adapting to execution evidence or invoking a model Controller. Fixed control tests whether repeatedly restating the task goal is sufficient without state-dependent guidance. Both references use the same starting workspace, Worker, coding tools, execution budget, and task evaluator as the Controller-guided runs.

Let $t \in \{\mathrm{II}, \mathrm{III}\}$ index the setting, N_t its number of tasks, and Π the evaluated Controller models. The full policy set is $\mathcal{P} = \Pi \cup \{\mathrm{NC}, \mathrm{FC}\}$, where NC and FC denote no control and fixed control. Every policy is evaluated in $K = 3$ separate runs per task. The no-control and fixed-control runs are collected once per task and reused across all Controller comparisons.

A run counts as successful only if it passes the task evaluator and follows the setting's control protocol. For SCBench, we require all Core checks to pass, as they assess the main functionality specified by the task. For BeyondSWE, success requires a reward of one from the official Harbor evaluator; LoopsBench uses its source-native evaluator. Multi-checkpoint tasks must pass all scored checkpoints.

For policy $p \in \mathcal{P}$, let $Y_{p,i,r}^{(t)} \in \{0, 1\}$ be the outcome of run $r \in \{1, \ldots, K\}$ on task $i \in \{1, \ldots, N_t\}$, with one denoting success and zero otherwise. The Strict Success Rate is

$$\mathrm{SSR}_t(p) = \frac{1}{N_t K} \sum_{i=1}^{N_t} \sum_{r=1}^{K} Y_{p,i,r}^{(t)}.$$

We report $100 \times \mathrm{SSR}_t(p)$ as a percentage. The reference policies appear alongside the Controller models in outcome and cost comparisons but are excluded from the Controller ranking.

Estimated inference cost is the primary resource measure, with all estimates assuming no prompt caching. We apply the provider's publicly listed standard input and output prices from a frozen schedule (Appendix F.2) to each call's recorded token counts, then sum the call-level estimates over the run. No-control and fixed-control totals include Worker calls; Controller-guided totals include Worker, Reporter, and Controller calls. We report the mean estimated cost per run for each policy in both settings, together with main Worker ReAct turns and control cycles as auxiliary resource measures.

Type III provides the full-task reference for comparing Controllers. We use Spearman's rank correlation ρ to measure the agreement between Controller orderings in Type II and Type III. The ranking uses only $\pi \in \Pi$, excluding the two reference policies:

$$\rho_{\mathrm{II,III}} = \rho_{\mathrm{S}}\Big(\big(\mathrm{SSR}_{\mathrm{II}}(\pi)\big)_{\pi\in\Pi}, \big(\mathrm{SSR}_{\mathrm{III}}(\pi)\big)_{\pi\in\Pi}\Big).$$

Appendix E gives the complete metric definitions; the detailed executable protocol appears in Appendix C.

5 Experiments

5.1 Experimental Setup

We evaluate the same Controller models in all three settings. Type I measures Contract selection, while Type II and Type III evaluate the Controller–Worker loop on paired task slices and full coding tasks. All Type II and Type III runs in the main experiment use Qwen3.7-Plus as the shared Worker; the Reporter uses the same model configuration. Appendix F.3 reports the comparison with a DeepSeek Worker. For Type II and Type III, we also define the shared no-control and fixed-control reference policies described in Section 4. Both use the same Worker and execution setup as the evaluated Controller models, and their runs are reused across Controller comparisons.

Our main analysis asks how Controller models compare on full tasks, whether Type II retains that comparison with less execution, and what Type I reveals about individual control decisions. Model revisions, provider settings, and the complete run plan are fixed before scoring and reported in Appendix D; Appendix I provides the prompts and message templates, and Appendix H lists the released artifacts.

5.2 Main Results

Table 2 reports all three settings side by side: Contract Accuracy for Type I, and Strict Success Rate with mean estimated inference cost per run for Type II and Type III. Detailed token, Worker-turn, and control-cycle measurements are reported in Appendix F.2. The no-control and fixed-control policies are included as shared references for Type II and Type III; Controller rankings are computed only over the evaluated Controller models. Invalid Type I responses count as incorrect.

Full-task control remains difficult. Across the evaluated Controllers, Type III Strict Success Rate ranges from **24.17%** to **36.25%**, and the strongest reaches only **36.25%**. Claude Opus 4.8 achieves 58.75% on Type II, compared with 55.00% for fixed control; GPT-5.5 achieves 36.25% on Type III, compared with 31.25%.

Controller choice and downstream execution. The cost of the work performed under a Controller's guidance can outweigh the savings on the Controller calls themselves. In Type III, DeepSeek-V4-Flash-0731 has the lowest estimated Controller-call cost, \$0.15 per run, but a total cost of \$32.66; the Worker accounts for 93.53% of this total. Qwen3.7-Plus incurs \$0.23 in Controller calls, yet its total cost is \$12.69 and its SSR is higher, 34.17% versus 28.33% (Table 3).

The runs also differ in the amount of Worker execution. DeepSeek-V4-Flash-0731 averages 330.58 Worker turns per run, compared with 243.33 for Qwen3.7-Plus, despite its lower SSR. The relationship between execution length and success is not uniform across Controllers: GPT-5.5 has both the highest mean Worker-turn count and the highest SSR, while GLM 5.2 has the lowest values on both measures.

Equal success rates can also correspond to different execution profiles. Qwen3.6-35B-A3B and Claude Opus 4.8 both complete 81 of 240 full-task runs, yielding an SSR of 33.75%. Their mean

Table 2: **Main results.** Type I reports Contract Accuracy; Type II and Type III report Strict Success Rate and mean estimated inference cost per run. The best Controller result in each column is bold and the second-best distinct value is underlined; reference policies are not ranked. Resource details appear in Appendix F. A dash marks an inapplicable field.

	Type I	Type II (task slice)		Type III (full task)	
Method	Contract Acc. (%) ↑	SSR (%) ↑	Est. cost ($/run) ↓	SSR (%) ↑	Est. cost ($/run) ↓
Reference policies					
No control	—	49.17	1.62	27.50	2.61
Fixed control	—	55.00	5.62	31.25	9.47
Controller models					
Qwen3.6-35B-A3B	43.125	55.83	<u>4.21</u>	33.75	**11.58**
Qwen3.7-Plus	<u>60.625</u>	<u>56.67</u>	6.04	<u>34.17</u>	12.69
DeepSeek-V4-Flash-0731	45.625	50.83	9.88	28.33	32.66
GLM 5.2	54.375	40.42	**3.54**	24.17	<u>11.97</u>
GPT-5.5	**61.875**	<u>56.67</u>	13.80	**36.25**	34.68
Claude Opus 4.8	56.250	**58.75**	10.88	33.75	20.63

Table 3: **Success and execution cost on full tasks.** SSR uses Core. Worker turns and costs are means over all 240 runs per Controller. Total estimated inference cost includes Worker, Reporter, and Controller calls under the main table's no-cache price schedule.

Controller	SSR (%)	Worker turns	Controller cost ($/run)	Total cost ($/run)
Qwen3.6-35B-A3B	33.75	224.36	0.20	11.58
Qwen3.7-Plus	34.17	243.33	0.23	12.69
DeepSeek-V4-Flash-0731	28.33	330.58	0.15	32.66
GLM 5.2	24.17	189.41	0.81	11.97
GPT-5.5	36.25	355.94	6.18	34.68
Claude Opus 4.8	33.75	324.09	7.16	20.63

execution costs are $11.58 and $20.63 per run, respectively, with 224.36 and 324.09 Worker turns. Thus, the same aggregate success rate can be reached with substantially different amounts of Worker execution and estimated inference cost.

Type II enables lower-cost Controller comparison. Across the evaluated Controllers, the paired Type II reduction in estimated inference cost averages **60.6%** relative to Type III. Type II retains the executable Controller–Worker loop and task evaluator; its lower cost comes from the shorter execution scope. Under the main Core criterion, its observed Controller ordering is positively correlated with Type III, with $\rho_{\mathrm{II,III}} = \mathbf{0.7353}$: two of the thirteen Controller pairs strictly ordered in both settings reverse order, while two of the fifteen pairs contain a tie. Type II provides a lower-cost comparison setting, while Type III remains the full-task assessment.

Type I reveals differences in individual control decisions. Contract Accuracy ranges from 43.125% to 61.875%, and Invalid Rate ranges from 0% to 1.875% (Appendix Table 8). The comparison concerns which next assignment the Controller selects from the available alternatives, using the same progress report and supporting evidence for each question. The strongest deterministic shortcut reaches 30.625%; uniform choice, action-preference rules, candidate length, and lexical overlap all score lower (Appendix Table 13). None of the tested shortcuts approaches the Controller scores. Because candidate execution is performed once during benchmark construction and reused, evaluating another Controller requires one response per question and no Worker run.

6 Related Work

Coding-agent benchmarks. SWE-bench and its variants ask whether a coding agent or agent system produces a correct final repository state (Jimenez et al., 2024; OpenAI, 2024; Zan et al., 2025; Deng et al., 2025); newer benchmarks extend the workload to feature implementation and longer

development tasks (Chen et al., 2025; Feng et al., 2026; Orlanski et al., 2026; Chen et al., 2026a; Le et al., 2026; Huang et al., 2026). Studies of leakage, memorization, and task or reward design have sharpened the standards for constructing and validating these benchmarks (Matton et al., 2024; Riddell et al., 2024; Liang et al., 2025; Zhu et al., 2025; Jain et al., 2025; Badertdinov et al., 2025). Their primary object of evaluation is the complete coding-agent system. LoopArena changes that object: the coding Worker and execution setup are fixed, and the model directing the Worker is compared.

Harnesses and Loop Engineering. Coding-agent harnesses determine the tools, context policy, state representation, and control logic surrounding a model (Wang et al., 2025; Yang et al., 2024; Rombaut, 2026). Harness-Bench and LoopsBench compare complete model–harness configurations (Yao et al., 2026; Li et al., 2026). Natural-Language Agent Harnesses make run-level policy editable in text, whereas Meta-Harness and HarnessX revise harness code or components from execution traces (Pan et al., 2026; Lee et al., 2026; Chen et al., 2026b). LongHorizon-Harness externalizes task-state management through a manager–executor–auditor loop (Ma et al., 2026). LoopArena fixes the Worker and control interface and compares Controller models. The measured object is the runtime control studied in Loop Engineering (Osmani, 2026; MSV, 2026).

Process and feedback evaluation. Interactive benchmarks measure multi-turn tool use and environment interaction (Liu et al., 2024; Yao et al., 2025; Barres et al., 2025; Zhou et al., 2024; Song et al., 2026), while process-supervision work scores individual steps or decisions in recorded trajectories (Lightman et al., 2024; Uesato et al., 2022; Wang et al., 2024a; Yu et al., 2025). Process-reward benchmarks extend this step-level view (Zheng et al., 2025; Song et al., 2025), as do evaluations of repository exploration and tool-use traces (Fan et al., 2026; Zhang et al., 2026a; Mazaheri & Mazaheri, 2026; Deshpande et al., 2025). Critic and self-correction studies examine feedback for evaluating and revising model outputs (Shinn et al., 2023; Madaan et al., 2023; McAleese et al., 2024; Huang et al., 2024; Olausson et al., 2024). Outcome-based reinforcement learning can improve model reasoning without process-level labels (Chu et al., 2026). In agentic settings, execution-grounded learning further uses outcomes from tree-structured interaction rollouts and rendered interface code as supervision (Ji et al., 2025; Zheng et al., 2026). LoopArena evaluates repeated instructions from one model to a separate coding agent, linking individual control decisions to their downstream effects under execution. Appendix G compares these settings by input, output, executor, signal, and measured object.

Efficient benchmark evaluation. ConvCodeWorld compares a lower-cost static benchmark with an interactive environment through model-ranking correlation (Han et al., 2025). Benchmark-compression methods likewise use Kendall's τ or Spearman's ρ to test whether a cheaper evaluation recovers a full-benchmark ranking (Yuan et al., 2025; Zhang et al., 2026b). LoopArena applies the same validation principle to paired task-slice and full-task executions.

7 Limitations

LoopArena currently focuses on repository-level coding tasks and on a Controller–Worker organization with structured handoffs. Extending the benchmark to additional software domains, Worker families, multi-Worker settings, and loop organizations would broaden its coverage. Applying the three-setting design beyond coding will also require domain-specific task construction and executable evaluators.

8 Conclusion

LoopArena makes a model's ability to control a long-running coding-agent loop directly measurable while holding the coding Worker fixed. Its three settings examine this ability at complementary resolutions: Type I isolates individual control decisions at low cost, Type II executes repeated guidance over a selected task slice, and Type III evaluates the complete task from its original state. Full-task performance remains limited, with a highest observed Type III Strict Success Rate of **36.25%**. The best Controller result in Type II and Type III exceeds fixed control. Across Controllers, the paired Type II cost reduction averages **60.6%**, while the observed Core-based ordering is positively correlated with Type III ($\rho_{\mathrm{II,III}} = \mathbf{0.7353}$). Together, these settings support the study of how models guide implementation, verification, and stopping across long-running coding tasks.

AI USE STATEMENT

We used generative AI tools for the benchmark construction and review steps described in the paper, and to assist with implementation and debugging, feedback on experimental design, result analysis, literature search, and manuscript drafting and revision. We checked AI-assisted code and analyses using tests and experimental records. The authors reviewed the manuscript and take responsibility for the final content and artifacts.

ETHICS STATEMENT

LoopArena supports research on the reliability of coding agents using tasks from existing software benchmarks. Generated code should undergo appropriate review and security checks before use outside the evaluation environment. Use and redistribution of the underlying repositories and benchmark materials should follow their respective licenses and attribution requirements.

REPRODUCIBILITY STATEMENT

The benchmark data and evaluation code are available through the anonymous repository linked in the abstract. The appendices document benchmark construction, model configurations and execution budgets, metric and cost definitions, and the prompts used in evaluation. Appendix H describes the released artifacts and preparation instructions.

REFERENCES

Alibaba Cloud. Model Studio Model Pricing, 2026. URL https://help.aliyun.com/en/model-studio/model-pricing. Accessed August 16, 2026.

Anthropic. Claude API Pricing, 2026. URL https://platform.claude.com/docs/en/about-claude/pricing. Accessed August 16, 2026.

Ibragim Badertdinov, Alexander Golubev, Maksim Nekrashevich, Anton Shevtsov, Simon Karasik, Andrei Andriushchenko, Maria Trofimova, Daria Litvintseva, and Boris Yangel. SWE-rebench: An Automated Pipeline for Task Collection and Decontaminated Evaluation of Software Engineering Agents. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, pp. 26420–26466, 2025. doi: 10.52202/085713-0788.

Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. τ^2-Bench: Evaluating Conversational Agents in a Dual-Control Environment, 2025. URL https://arxiv.org/abs/2506.07982. arXiv preprint arXiv:2506.07982.

Guoxin Chen, Fanzhe Meng, Jiale Zhao, Minghao Li, Daixuan Cheng, Huatong Song, Jie Chen, Yuzhi Lin, Hui Chen, Xin Zhao, Ruihua Song, Chang Liu, Cheng Chen, Kai Jia, and Ji-Rong Wen. BeyondSWE: Can Current Code Agent Survive Beyond Single-Repo Bug Fixing?, 2026a. URL https://arxiv.org/abs/2603.03194. arXiv preprint arXiv:2603.03194.

Haorui Chen, Chengze Li, and Jia Li. FeatBench: Towards More Realistic Evaluation of Feature-level Code Generation, 2025. URL https://arxiv.org/abs/2509.22237. arXiv preprint arXiv:2509.22237.

Tingyang Chen, Shuo Lu, Kang Zhao, Weicheng Meng, Hanlin Teng, Tianhao Li, Chao Li, Xule Liu, Jian Liang, Zhizhong Zhang, Yuan Xie, Heng Qu, Kun Shao, and Jian Luan. HarnessX: A Composable, Adaptive, and Evolvable Agent Harness Foundry, 2026b. URL https://arxiv.org/abs/2606.14249. arXiv preprint arXiv:2606.14249.

Xiangxiang Chu, Hailang Huang, Xiao Zhang, Fei Wei, and Yong Wang. GPG: A Simple and Strong Reinforcement Learning Baseline for Model Reasoning. In *International Conference on Learning Representations (ICLR)*, volume 2026, pp. 59637–59659, 2026. URL https://openreview.net/forum?id=inccdtfx8x.

Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Nitin Pasari, Chetan Rane, Karmini Sampath, Maya Krishnan, Srivatsa Kundurthy, Sean Hendryx, Zifan Wang, Vijay Bharadwaj, Jeff Holm, Raja Aluri, Chen Bo Calvin Zhang, Noah Jacobson, Bing Liu, and Brad Kenstler. SWE-Bench Pro: Can AI Agents Solve Long-Horizon Software Engineering Tasks?, 2025. URL https://arxiv.org/abs/2509.16941. arXiv preprint arXiv:2509.16941.

Darshan Deshpande, Varun Gangal, Hersh Mehta, Jitin Krishnan, Anand Kannappan, and Rebecca Qian. TRAIL: Trace Reasoning and Agentic Issue Localization, 2025. URL https://arxiv.org/abs/2505.08638. arXiv preprint arXiv:2505.08638.

Shengda Fan, Xuyan Ye, Yupeng Huo, Zhi-Yuan Chen, Yiju Guo, Shenzhi Yang, Wenkai Yang, Shuqi Ye, Jingwen Chen, Haotian Chen, Xin Cong, and Yankai Lin. AgentProcessBench: Diagnosing Step-Level Process Quality in Tool-Using Agents, 2026. URL https://arxiv.org/abs/2603.14465. arXiv preprint arXiv:2603.14465.

Yukang Feng, Jianwen Sun, Zelai Yang, Jiaxin Ai, Chuanhao Li, Zizhen Li, Fanrui Zhang, Kang He, Rui Ma, Jifan Lin, Jie Sun, Yang Xiao, Sizhuo Zhou, Wenxiao Wu, Yiming Liu, Pengfei Liu, Yu Qiao, Shenglin Zhang, and Kaipeng Zhang. LongCLI-Bench: A Preliminary Benchmark and Study for Long-horizon Agentic Programming in Command-Line Interfaces, 2026. URL https://arxiv.org/abs/2602.14337. arXiv preprint arXiv:2602.14337.

Hojae Han, Seung won Hwang, Rajhans Samdani, and Yuxiong He. ConvCodeWorld: Benchmarking Conversational Code Generation in Reproducible Feedback Environments. In *International Conference on Learning Representations (ICLR)*, 2025.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large Language Models Cannot Self-Correct Reasoning Yet. In *International Conference on Learning Representations (ICLR)*, 2024.

Wenqi Huang, Charley Lee, Leonard Tng, and Serena Ge. DeepSWE: Measuring frontier coding agents on original, long-horizon engineering tasks. *arXiv preprint arXiv:2607.07946*, 2026. URL https://arxiv.org/abs/2607.07946.

Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. LiveCodeBench: Holistic and Contamination Free Evaluation of Large Language Models for Code. In *International Conference on Learning Representations (ICLR)*, 2025.

Yuxiang Ji, Ziyu Ma, Yong Wang, Guanhua Chen, Xiangxiang Chu, and Liaoni Wu. Tree Search for LLM Agent Reinforcement Learning, 2025. URL https://arxiv.org/abs/2509.21240. arXiv preprint arXiv:2509.21240.

Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. SWE-bench: Can Language Models Resolve Real-World GitHub Issues? In *International Conference on Learning Representations (ICLR)*, 2024.

Tue Le, Minh V. T. Thai, Dung Nguyen Manh, Huy Phan Nhat, and Nghi D. Q. Bui. SWE-EVO: Benchmarking coding agents in long-horizon software evolution scenarios. *arXiv preprint arXiv:2512.18470*, 2026. URL https://arxiv.org/abs/2512.18470.

Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. Meta-Harness: End-to-End Optimization of Model Harnesses, 2026. URL https://arxiv.org/abs/2603.28052. arXiv preprint arXiv:2603.28052.

Han Li, Zhemin Fang, Rili Feng, Yingqi Zhao, Jiaheng Liu, Pengfei Gao, He Ye, Dayi Lin, Qingwei Lin, Saravan Rajmohan, and Dongmei Zhang. LoopsBench: From Harness Engineering to Loop Engineering in Coding Agent Evaluation, 2026. URL https://arxiv.org/abs/2608.00267. arXiv preprint arXiv:2608.00267.

Shanchao Liang, Spandan Garg, and Roshanak Zilouchian Moghaddam. The SWE-Bench Illusion: When State-of-the-Art LLMs Remember Instead of Reason, 2025. URL https://arxiv.org/abs/2506.12286. arXiv preprint arXiv:2506.12286.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's Verify Step by Step. In *International Conference on Learning Representations (ICLR)*, 2024.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. AgentBench: Evaluating LLMs as Agents. In *International Conference on Learning Representations (ICLR)*, 2024.

Ziyu Ma, Hailang Huang, Shun Zou, Yong Wang, Shidong Yang, Yiming Hu, Fei Wei, and XiangXiang Chu. LongHorizon-Harness: Advancing Long-Horizon Agents for Real-World Tasks, 2026. URL https://arxiv.org/abs/2608.01964. arXiv preprint arXiv:2608.01964.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-Refine: Iterative Refinement with Self-Feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 46534–46594, 2023. doi: 10.52202/075280-2019.

Alexandre Matton, Tom Sherborne, Dennis Aumiller, Elena Tommasone, Milad Alizadeh, Jingyi He, Raymond Ma, Maxime Voisin, Ellen Gilsenan-McMahon, and Matthias Gallé. On Leakage of Code Generation Evaluation Datasets. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 13215–13223, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.772.

Parsa Mazaheri and Kasra Mazaheri. AgentAtlas: Beyond Outcome Leaderboards for LLM Agents, 2026. URL https://arxiv.org/abs/2605.20530. arXiv preprint arXiv:2605.20530.

Nat McAleese, Rai Michael Pokorny, Juan Felipe Ceron Uribe, Evgenia Nitishinskaya, Maja Trebacz, and Jan Leike. LLM Critics Help Catch LLM Bugs, 2024. URL https://arxiv.org/abs/2407.00215. arXiv preprint arXiv:2407.00215.

Janakiram MSV. The Anthropic Leader Who Built Claude Code Says He Ditched Prompting—Now He Just Writes Loops. The New Stack, 2026. URL https://thenewstack.io/loop-engineering/. Practitioner news article.

Theo X. Olausson, Jeevana Priya Inala, Chenglong Wang, Jianfeng Gao, and Armando Solar-Lezama. Is Self-Repair a Silver Bullet for Code Generation? In *International Conference on Learning Representations (ICLR)*, 2024.

OpenAI. Introducing SWE-bench Verified. OpenAI blog, 2024. URL https://openai.com/index/introducing-swe-bench-verified/. Blog post.

OpenAI. GPT-5.5 Model and API Pricing, 2026. URL https://developers.openai.com/api/docs/models/gpt-5.5. Accessed August 16, 2026.

Gabriel Orlanski, Devjeet Roy, Alexander Yun, Changho Shin, Alex Gu, Albert Ge, Dyah Adila, Nicholas Roberts, Frederic Sala, and Aws Albarghouthi. SlopCodeBench: Benchmarking How Coding Agents Degrade Over Long-Horizon Iterative Tasks, 2026. URL https://arxiv.org/abs/2603.24755. arXiv preprint arXiv:2603.24755.

Addy Osmani. Loop Engineering. Personal blog, 2026. URL https://addyosmani.com/blog/loop-engineering/.

Linyue Pan, Lexiao Zou, Shuo Guo, Jingchen Ni, and Hai-Tao Zheng. Natural-Language Agent Harnesses, 2026. URL https://arxiv.org/abs/2603.25723. arXiv preprint arXiv:2603.25723.

Martin Riddell, Ansong Ni, and Arman Cohan. Quantifying Contamination in Evaluating Code Generation Capabilities of Language Models, 2024. URL https://arxiv.org/abs/2403.04811. arXiv preprint arXiv:2403.04811.

Benjamin Rombaut. Inside the Scaffold: A Source-Code Taxonomy of Coding Agent Architectures, 2026. URL https://arxiv.org/abs/2604.03515. arXiv preprint arXiv:2604.03515.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language Agents with Verbal Reinforcement Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 8634–8652, 2023. doi: 10.52202/075280-0377.

Mingyang Song, Zhaochen Su, Xiaoye Qu, Jiawei Zhou, and Yu Cheng. PRMBench: A Fine-grained and Challenging Benchmark for Process-Level Reward Models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*, pp. 25299–25346, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.1230.

Zhiheng Song, Jingshuai Zhang, Chuan Qin, Chao Wang, Chao Chen, Longfei Xu, Kaikui Liu, Xiangxiang Chu, and Hengshu Zhu. MobilityBench: A Benchmark for Evaluating Route-Planning Agents in Real-World Mobility Scenarios, 2026. URL https://arxiv.org/abs/2602.22638. arXiv preprint arXiv:2602.22638.

State Administration of Foreign Exchange. RMB Central Parity Rate, 2026. URL https://www.safe.gov.cn/AppStructured/hlw/RMBQuery.do. August 14, 2026 rate; accessed August 16, 2026.

Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving Math Word Problems with Process- and Outcome-Based Feedback, 2022. URL https://arxiv.org/abs/2211.14275. arXiv preprint arXiv:2211.14275.

Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-Shepherd: Verify and Reinforce LLMs Step-by-step without Human Annotations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*, pp. 9426–9439, Bangkok, Thailand, August 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.510.

Xingyao Wang, Yangyi Chen, Lifan Yuan, Yizhe Zhang, Yunzhu Li, Hao Peng, and Heng Ji. Executable Code Actions Elicit Better LLM Agents. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pp. 50208–50232. PMLR, 2024b.

Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, and Graham Neubig. OpenHands: An Open Platform for AI Software Developers as Generalist Agents. In *International Conference on Learning Representations (ICLR)*, 2025.

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 *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, pp. 50528–50652, 2024. doi: 10.52202/079017-1601.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. ReAct: Synergizing Reasoning and Acting in Language Models. In *International Conference on Learning Representations (ICLR)*, 2023.

Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. τ-bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains. In *International Conference on Learning Representations (ICLR)*, 2025.

Yilun Yao, Xinyu Tan, Chao-Hsuan Liu, Yaoming Li, Zhengyang Wang, Wenhan Yu, Zhewen Tan, Yuxuan Tian, Guangxiang Zhao, Lin Sun, Xiangzheng Zhang, and Tong Yang. Harness-Bench: Measuring Harness Effects across Models in Realistic Agent Workflows, 2026. URL https://arxiv.org/abs/2605.27922. arXiv preprint arXiv:2605.27922.

Zhuohao Yu, Weizheng Gu, Yidong Wang, Xingru Jiang, Zhengran Zeng, Jindong Wang, Wei Ye, and Shikun Zhang. Reasoning Through Execution: Unifying Process and Outcome Rewards for Code Generation. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, volume 267 of *Proceedings of Machine Learning Research*, pp. 72972–72994. PMLR, 2025.

Peiwen Yuan, Yueqi Zhang, Shaoxiong Feng, Yiwei Li, Xinglin Wang, Jiayi Shi, Chuyi Tan, Boyuan Pan, Yao Hu, and Kan Li. Beyond One-Size-Fits-All: Tailored Benchmarks for Efficient Evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15591–15615, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.759.

Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Shulin Xin, Linhao Zhang, Qi Liu, Li Aoyan, Lu Chen, Xiaojian Zhong, Siyao Liu, Yongsheng Xiao, Liangqiang Chen, Yuyu Zhang, Jing Su, Tianyu Liu, Rui Long, Ming Ding, and Liang Xiang. Multi-SWE-bench: A Multilingual Benchmark for Issue Resolving. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, pp. 70112–70144, 2025. doi: 10.52202/085713-2111.

Shaoqiu Zhang, Yuhang Wang, Jialiang Liang, Yuling Shi, Wenhao Zeng, Maoquan Wang, Shilin He, Ningyuan Xu, Siyu Ye, Kai Cai, and Xiaodong Gu. SWE-Explore: Benchmarking How Coding Agents Explore Repositories, 2026a. URL https://arxiv.org/abs/2606.07297. arXiv preprint arXiv:2606.07297.

Yueqi Zhang, Jin Hu, Shaoxiong Feng, Peiwen Yuan, Xinglin Wang, Yiwei Li, Jiayi Shi, Chuyi Tan, Ji Zhang, Boyuan Pan, Yao Hu, and Kan Li. Learning More from Less: Unlocking Internal Representations for Benchmark Compression, 2026b. URL https://arxiv.org/abs/2602.00710. arXiv preprint arXiv:2602.00710.

Chujie Zheng, Zhenru Zhang, Beichen Zhang, Runji Lin, Keming Lu, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. ProcessBench: Identifying Process Errors in Mathematical Reasoning. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL), Volume 1: Long Papers*, pp. 1009–1024, Vienna, Austria, July 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.acl-long.50.

Yuhao Zheng, Li'an Zhong, Yi Wang, Rui Dai, Kaikui Liu, Xiangxiang Chu, Linyuan Lv, Philip Torr, and Kevin Qinghong Lin. Code2World: A GUI World Model via Renderable Code Generation, 2026. URL https://arxiv.org/abs/2602.09856. arXiv preprint arXiv:2602.09856.

Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. WebArena: A Realistic Web Environment for Building Autonomous Agents. In *International Conference on Learning Representations (ICLR)*, 2024.

Yuxuan Zhu, Tengjun Jin, Yada Pruksachatkun, Andy Zhang, Shu Liu, Sasha Cui, Sayash Kapoor, Shayne Longpre, Kevin Meng, Rebecca Weiss, Fazl Barez, Rahul Gupta, Jwala Dhamala, Jacob Merizian, Mario Giulianelli, Harry Coppock, Cozmin Ududec, Antony Kellermann, Jasjeet Sekhon, Jacob Steinhardt, Sarah Schwettmann, Arvind Narayanan, Matei A. Zaharia, Ion Stoica, Percy Liang, and Daniel Kang. Establishing Best Practices in Building Rigorous Agentic Benchmarks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 38, pp. 184435–184475, 2025. doi: 10.52202/085713-5547.

APPENDIX

A ADDITIONAL BENCHMARK DETAILS

A.1 BENCHMARK COMPOSITION

LoopArena contains 160 Type I questions and 80 paired Type II–Type III task instances. Each pair consists of one selected task slice for Type II and its corresponding full task for Type III. The Type III panel contains 80 unique official tasks: 21 from SCBench, 29 from BeyondSWE, and 30 from LoopsBench.

SCBench tasks may contain several ordered native checkpoints. The complete checkpoint sequence is executed for Type III, but the task is counted once in the benchmark.

A.2 ILLUSTRATIVE TYPE I QUESTION

The following shortened example is adapted from a scored Type I item. The source task removes unsupported Django versions and updates the supported test matrix. At the selected control point, the Evidence Packet shows that the requested matrix changes are largely complete, but an unused Django REST Framework 3.9 dependency (`drf39`) remains. The Controller chooses the best next Loop Contract from four candidates.

Table 4: **Shortened Type I example.** The question asks which instruction should be issued next; the tested Controller does not execute code.

Candidate	Next instruction	Label
A	Remove a dependency required by a newly supported Django configuration.	Incorrect
B	Remove a dependency still used by the active latest-version environment.	Incorrect
C	Keep the unused `drf39` dependency and inspect the matrix again.	Incorrect
D	Remove the orphaned `drf39` dependency while preserving active configurations.	Correct

B BENCHMARK CONSTRUCTION AND QUALITY CONTROL

This appendix specifies Type II slice selection, the procedures that determine Type I answers, and the quality-control checks applied to the benchmark. Source-task construction is summarized in Section 3.

Type II slice selection. After validating the candidate stages, we select one slice per full task. We sample slices separately within each benchmark source, balancing coverage of early, middle, and late stages according to the valid stages available in each task. The selected slices are fixed for all Controller evaluations and paired with their corresponding full tasks.

B.1 TYPE I CANDIDATE CONSTRUCTION AND REPLAY

For each new Type I item, we first sample a parent trajectory containing at least one restorable, non-bootstrap control cycle and then sample one such cycle within that trajectory. Sampling parents before cycles prevents longer trajectories from receiving greater weight merely because they contain more control decisions. The recorded Loop Contract is retained as one candidate but is not assumed to be the correct option.

One candidate author produces three complete and plausible alternatives to the recorded Contract. Author assignments are balanced across the designated model families. The resulting four Contracts must satisfy the same public schema and are placed in a neutral presentation order. Both the candidates and their order are frozen before any downstream execution is observed.

We then execute all four candidates from the same restored state under two predeclared matched replay schedules. Within each schedule, the four candidate executions use the same Worker, execution budget, evaluator, continuation policy, stopping rules, and random seed; only the initial Loop

Contract differs. If the source trajectory used seed S, the primary and confirmation schedules use $S + 1{,}000{,}000$ and $S + 2{,}000{,}000$, respectively.

For each question, the continuation Controller is randomly selected during benchmark construction. After the initial Contract, it guides subsequent rounds under the shared continuation policy; within each replay schedule, its configuration is held fixed across the four candidates.

Terminal task success is the primary outcome. If exactly one candidate succeeds, it is the unique winner of that schedule. If multiple candidates succeed, we compare them first by the number of downstream Controller cycles and then, among candidates with the same number of cycles, by the number of downstream Worker turns. A schedule has a unique winner only when exactly one successful candidate is best under this ordering. Model-token counts and repository calls are recorded as diagnostics but do not determine the winner.

An item is retained only when the primary and confirmation schedules identify the same unique winner. We reject the item if all four candidates fail, if a schedule has no unique winner under this ordering, or if the two schedules disagree. Rejected items are not repaired by replacing or rewriting candidates after their outcomes are known.

The confirmed winner becomes the correct option under this replay procedure, whether it is the historical Contract or an alternative. Difficulty measurements, shortcut analyses, and human or model-based audits are performed only after the item is frozen. They characterize the resulting benchmark but do not alter its candidates, correct option, or inclusion decision.

B.2 Quality control and post-freeze audits

Automated checks verify the data schema, source binding, restored state, evaluator configuration, candidate schema, and matched replay setup. For Type I, executable outcomes determine the correct option; neither model review nor expert review can override the confirmed replay winner. Post-freeze LLM-assisted and expert audits examine question clarity, candidate plausibility, information leakage, and consistency between the public Packet and the frozen answer. Table 5 lists the construction checks and post-freeze audits.

Table 5: **Benchmark construction checks and audits.** Automated checks enforce the execution protocol; post-freeze review characterizes Type I quality without changing its replay-defined answer.

Setting	Check	Requirement
All	task binding	task text, starting workspace, and evaluator refer to the same frozen source instance
All	solution leakage	model-visible inputs exclude official solution code, private scoring output, and later trajectory events
Type I	answer stability	the two predeclared replay schedules identify the same unique winner
Type I	post-freeze audit	reviewers assess clarity, candidate plausibility, and leakage without changing the replay-defined answer
Type II	task-slice validity	the prepared starting workspace fails at least one selected-stage check, while the source-provided completed state passes all requirements through that stage
Type II/III	evaluator validity	the frozen evaluator distinguishes a completed source task from an incomplete state

The benchmark definition fixes the included questions and tasks and their evaluator configurations.

C Harness and Evaluation Details

C.1 Roles and information flow

The Worker, Reporter, and Controller are model instances, whereas the Evidence Packet and Loop Contract are structured artifacts. The harness creates, validates, and renders these artifacts deter-

ministically; it makes no additional model call when formatting Reporter output as a Packet or a continuing Contract as a Worker instruction.

Only the Worker edits the repository, runs commands, and uses coding tools. At each Controller-guided handoff, the harness creates a temporary Reporter agent using the same model configuration and a copy of the accumulated Worker conversation. The Reporter can inspect a static, read-only view of the workspace but cannot modify it. It summarizes the task context, work completed so far, verification evidence, and unresolved issues, with citations to the relevant Worker turns. The Reporter interaction is not appended to the persistent Worker conversation.

At each control cycle, the Controller-visible Evidence Packet contains the overall repository task, the context of the current report, and four non-empty Reporter fields: `task_context_and_constraints`, `work_history_and_current_state`, `verification_and_evidence`, and `open_issues_and_uncertainty`. It also includes the complete Worker turns cited by the Reporter and the total, used, and remaining Worker-turn budget.

In Type II and Type III, the Controller returns a structured decision with an `action` in {`advance`, `verify`, `stop`} and a free-text `rationale`. An `advance` or `verify` decision also provides a `worker_instruction` with `goal`, `context`, `required_outcomes`, `prohibited_actions`, and `completion_condition`; a list of `protected_invariants`; and a `verification_acceptance_condition`. A `stop` decision requires only the action and rationale. The harness validates the decision and records it as the canonical Loop Contract used for execution and for Type I candidates. The Controller has no repository or coding tools. For a continuing Contract, the harness renders the validated instruction as the next user turn in the persistent Worker conversation; for a Stop Contract, it submits the current workspace to the evaluator.

C.2 Illustrative Type II control cycle

The following shortened example is taken from one recorded Type II control cycle on an SCBench task involving watch mode and a secondary configuration store. The first Worker segment was an orientation pass. Quoted text is unchanged; ellipses mark omitted material.

> **Reporter and Evidence Packet.** In `work_history_and_current_state`, the Reporter states: "The coding agent performed an orientation-only pass as assigned. ..." In `open_issues_and_uncertainty`, it records: "No implementation work has been performed; all task features remain to be added. ..." The harness places these fields, the task context, and the cited Worker turns in the Evidence Packet.
>
> **Controller decision.** The Controller returns `action: advance` and the following `worker_instruction.goal`: "Implement secondary-store support for schema validation and non-watch resolution, including CLI flags and seed lookup, while preserving existing behavior." Its required outcomes include adding the new source at the correct priority, implementing seed lookup, and demonstrating the behavior with focused checks.
>
> **Loop Contract and next Worker assignment.** The harness validates the Controller's decision and records it as a Loop Contract. It then renders the same assignment as the next user turn for the Worker, without another model call.

C.3 Controller-guided execution

Every Type II or Type III repeat begins from a clean copy of the specified starting workspace and a new Worker conversation. The Worker first receives the fixed bootstrap instruction described in Appendix I. When that segment ends, the Reporter summarizes the run, the harness builds the Packet, and the tested Controller produces the first Contract. An Advance or Verify Contract is rendered as the next user instruction in the same Worker conversation; a Stop Contract submits the current workspace to the source evaluator. The cycle repeats until Controller Stop, budget exhaustion, a protocol failure, or another terminal runtime condition.

Each main Worker episode ends after 600 ReAct turns or 7,200 seconds of main Worker wall time, whichever occurs first. At the turn limit, the final workspace is evaluated under the usual success criteria; reaching the wall-time limit counts as a task failure. Controller-guided execution allows at most 128 control cycles and 86,400 cumulative seconds, and each Reporter call has a 50-turn limit. These limits are shared across Controller models. Each Controller request allows up to 20,480

output tokens, and each Reporter request allows up to 8,192 output tokens. For multi-checkpoint SCBench tasks, the official workspace persists across checkpoints while the Worker and control conversations restart at each native boundary.

C.4 Reference policies

Type II and Type III include two shared reference policies. Under *no control*, the Worker receives the task once and runs as an ordinary coding agent without Reporter or Controller calls. Under *fixed control*, every nonterminal control decision is fixed to *proceed*; the policy does not inspect the Packet and invokes neither a Reporter nor a model Controller. The loop ends when the Worker explicitly declares the goal complete; this declaration controls only loop termination, while the same source evaluator still determines task success or failure. Fixed control tests whether repeatedly restating the task goal is sufficient without state-dependent guidance and is not ranked as a Controller model.

For a given task, the references and evaluated Controllers share the Worker configuration, tools, runtime limits, source environment, and evaluator. Their starting instructions differ because they implement different control policies. Each reference repeat set is collected once and reused across Controller comparisons.

C.5 Run outcomes and cost accounting

A run is successful only if it passes the frozen task evaluator and the machine-verifiable protocol. An invalid Controller response or a model-caused protocol violation counts as failure.

BeyondSWE success requires reward one from the frozen Harbor evaluator; LoopsBench uses its source-native evaluator. Multi-checkpoint success requires all scored checkpoints to pass. Resource accounting includes every model call made by the method: Worker calls under no control and fixed control; and Worker, Reporter, and Controller calls under model control. We record per-call input and output tokens, main Worker ReAct turns, control cycles, wall-clock time, and component-level token counts.

D Experimental Details

D.1 Controller panel and inference interfaces

The evaluated Controllers are Qwen3.6-35B-A3B, Qwen3.7-Plus, DeepSeek-V4-Flash-0731, GLM 5.2, GPT-5.5, and Claude Opus 4.8. The completed Type I runs use direct API calls. Qwen3.6-35B-A3B, Qwen3.7-Plus, DeepSeek-V4-Flash-0731, and GLM 5.2 use temperature zero, a maximum output length of 20,480 tokens, provider-default reasoning, and disabled client-library retries. GPT-5.5 and Claude Opus 4.8 use `gpt-5.5-0424-global` and `claude-opus-4-8`, respectively, with provider-default thinking, no temperature or seed parameter, and a maximum output length of 20,480 tokens.

Table 6: **Type I inference interfaces.** Exact provider identifiers and request records are preserved with the run manifest.

Controller	Type I inference interface
Qwen3.6-35B-A3B	API; temperature 0; maximum 20,480 output tokens
Qwen3.7-Plus	API; temperature 0; maximum 20,480 output tokens
DeepSeek-V4-Flash-0731	API; temperature 0; maximum 20,480 output tokens
GLM 5.2	API; temperature 0; maximum 20,480 output tokens
GPT-5.5	API (`gpt-5.5-0424-global`); provider-default thinking; no temperature or seed; maximum 20,480 output tokens
Claude Opus 4.8	API (`claude-opus-4-8`); provider-default thinking; no temperature or seed; maximum 20,480 output tokens

Type II and Type III evaluate the same Controllers. The completed API runs use the exact model identifiers `qwen3.6-35b-a3b`, `qwen3.7-plus`, `deepseek-v4-flash-0731`, `glm-5.2`,

`gpt-5.5-0424-global`, and `claude-opus-4-8`. The Qwen, DeepSeek, and GLM Controllers use temperature zero, a maximum output length of 20,480 tokens, provider-default reasoning, and no client-library retries. GPT-5.5 and Claude Opus 4.8 use provider-default thinking, no temperature or seed parameter, and a maximum output length of 20,480 tokens. The frozen run manifest binds the exact Controller, Worker, and Reporter revisions; prompts; provider parameters; source images; evaluator versions; and runtime limits used for every result.

D.2 Run plan

Table 7: **Evaluation plan.** Type II and Type III use the same Controllers, repeat count, and paired full tasks.

Quantity	Value or reporting rule
Type I panel	160 questions; one response per Controller and question
Type II panel	80 task slices paired with Type III
Type III panel	80 unique full tasks
Executable repeats	$K = 3$ separate runs for every Controller, reference policy, and task
Reference reuse	one no-control set and one fixed-control set per task and Worker configuration, shared across Controllers
Primary executable score	Strict Success Rate
Evaluation cost	estimated end-to-end inference cost per run under the frozen no-cache price schedule; tokens, Worker turns, and control cycles reported separately

E Metric Definitions and Cost Accounting

E.1 Type I

Contract Accuracy is the fraction of questions for which the chosen candidate matches the correct option fixed during benchmark construction. Missing, multiple, out-of-range, or unparseable answers are incorrect and also contribute to the Invalid Rate.

E.2 Type II and Type III

For each executable setting, Strict Success Rate averages the binary evaluator outcome over all tasks and $K = 3$ runs, as defined in Section 4. The shared no-control and fixed-control runs remain shared across Controller comparisons.

Using the policy set $\mathcal{P}$ defined in Section 4, let $\mathcal{J}^{(t)}_{p,i,r}$ be the model calls made by policy p on task i in run r of setting t. For call j, let I_j and O_j be its recorded input and output token counts, and let P^{in}_j and P^{out}_j be the corresponding public standard prices in USD per million tokens from the frozen schedule in Table 10. We estimate run cost as

$$C^{(t)}_{p,i,r} = \sum_{j \in \mathcal{J}^{(t)}_{p,i,r}} \frac{I_j P^{\text{in}}_j + O_j P^{\text{out}}_j}{10^6}, \qquad \overline{C}_t(p) = \frac{1}{N_t K} \sum_{i=1}^{N_t} \sum_{r=1}^{K} C^{(t)}_{p,i,r}.$$

The call set contains Worker calls under no control and fixed control, and Worker, Reporter, and Controller calls under model control. For simplicity, the calculation assumes no prompt caching: every input token uses the standard input price, with no cache discount. If a provider reports reasoning tokens separately, they are included in O_j exactly once. The resulting dollar values are standardized end-to-end estimates, not measurements of realized spending or Controller-only cost. They cover model inference only; repository execution, evaluator compute, storage, and other infrastructure overhead are not monetized. Cost results are reported only when input and output usage is available for every included call.

The Type II cost reduction averages the paired percentage reduction for each Controller,

$$R_C = \frac{100}{|\Pi|} \sum_{\pi \in \Pi} \left(1 - \frac{\overline{C}_{\text{II}}(\pi)}{\overline{C}_{\text{III}}(\pi)} \right).$$

Total model tokens, main Worker turns, and control cycles are summarized separately.

Spearman's $\rho_{\mathrm{II,III}}$ is calculated only over the Controller models evaluated in both settings, excluding no control and fixed control; tied scores receive their average rank. We also report pairwise rank reversals, counting only Controller pairs that are strictly ordered in opposite directions; ties are reported separately.

F Additional Results and Analyses

F.1 Type I results

Table 8 reports exact Type I counts, Contract Accuracy, Invalid Rate, and estimated evaluation cost.

Table 8: **Type I results.** Estimated cost is the standardized marginal no-cache cost of 160 Controller responses and excludes the one-time candidate-execution cost incurred during benchmark construction.

Controller	Correct	Accuracy (%)	Invalid (%)	Est. cost ($/160 questions)
Qwen3.6-35B-A3B	69/160	43.125	0.000	1.57
Qwen3.7-Plus	97/160	<u>60.625</u>	0.000	<u>1.45</u>
DeepSeek-V4-Flash-0731	73/160	45.625	0.625	**0.56**
GLM 5.2	87/160	54.375	1.875	6.81
GPT-5.5	99/160	**61.875**	0.000	21.67
Claude Opus 4.8	90/160	56.250	0.000	30.56

Table 9 reports Contract Accuracy separately for each source benchmark.

Table 9: **Type I accuracy by source benchmark.** Parentheses show percentages.

Controller	SCBench (40)	BeyondSWE (53)	LoopsBench (67)
Qwen3.6-35B-A3B	18/40 (45.00)	32/53 (60.38)	19/67 (28.36)
Qwen3.7-Plus	33/40 (82.50)	41/53 (77.36)	23/67 (34.33)
DeepSeek-V4-Flash-0731	21/40 (52.50)	32/53 (60.38)	20/67 (29.85)
GLM 5.2	29/40 (72.50)	37/53 (69.81)	21/67 (31.34)
GPT-5.5	35/40 (87.50)	46/53 (86.79)	18/67 (26.87)
Claude Opus 4.8	29/40 (72.50)	40/53 (75.47)	21/67 (31.34)

F.2 Executable results and cost

Table 10 records the public standard prices used to estimate inference cost, and Table 11 reports role-separated costs together with Worker-turn and control-cycle measurements.

Table 10: **Frozen price schedule for inference-cost estimation.** Prices were retrieved on August 16, 2026 from the providers' public price schedules (Alibaba Cloud, 2026; OpenAI, 2026; Anthropic, 2026). Qwen list prices are the China (Beijing) rates in CNY, converted at the August 14, 2026 official central parity rate of CNY 6.7878 per USD (State Administration of Foreign Exchange, 2026). We use public list prices and do not apply temporary promotions, caching, batch, priority, regional, or enterprise discounts.

Model	Input tokens per call	Input price ($/1M tokens)	Output price ($/1M tokens)
Qwen3.6-35B-A3B	all supported	0.2652	1.5911
Qwen3.7-Plus	up to 256K	0.2946	1.1786
	256K–1M	0.8839	3.5358
DeepSeek-V4-Flash-0731	all supported	0.1473	0.2946
GLM 5.2	all supported	1.1786	4.1250
GPT-5.5	up to 272K	5.0000	30.0000
	above 272K	10.0000	45.0000
Claude Opus 4.8	all supported	5.0000	25.0000

Table 11: **Mean estimated no-cache inference cost and execution length.** All reported entries are averaged over 240 scored runs per policy or Controller and setting. Worker turns count inner-loop ReAct turns; control cycles count outer-loop handoffs. A dash marks a model-call role or execution-length metric that does not apply to the policy. Fixed control uses neither a Reporter nor a model Controller. Totals may differ from displayed component sums because of rounding.

	Estimated cost ($/run)				Execution length	
	Worker	Reporter	Controller	Total	Worker turns	Control cycles
Type II (task slice)						
No control	1.62	—	—	1.62	68.50	—
Fixed control	5.62	—	—	5.62	105.82	2.85
Qwen3.6-35B-A3B	3.86	0.26	0.10	4.21	103.37	4.24
Qwen3.7-Plus	5.62	0.31	0.10	6.04	114.56	4.15
DeepSeek-V4-Flash-0731	9.29	0.53	0.06	9.88	134.35	4.71
GLM 5.2	2.98	0.20	0.37	3.54	86.61	3.55
GPT-5.5	10.59	0.60	2.60	13.80	145.30	5.42
Claude Opus 4.8	7.56	0.48	2.83	10.88	122.20	4.55
Type III (full task)						
No control	2.61	—	—	2.61	136.17	—
Fixed control	9.47	—	—	9.47	211.14	7.16
Qwen3.6-35B-A3B	10.67	0.71	0.20	11.58	224.36	9.70
Qwen3.7-Plus	11.71	0.74	0.23	12.69	243.33	10.61
DeepSeek-V4-Flash-0731	30.54	1.96	0.15	32.66	330.58	11.44
GLM 5.2	10.58	0.59	0.81	11.97	189.41	8.16
GPT-5.5	26.42	2.08	6.18	34.68	355.94	13.45
Claude Opus 4.8	12.71	0.76	7.16	20.63	324.09	11.85

F.3 WORKER-SIDE COMPARISON

To examine whether relative Controller performance is preserved with a different Worker, we evaluate four Controllers on 20 randomly selected source tasks in both Type II and Type III. The comparison uses DeepSeek-V4-Flash-0731 in place of Qwen3.7-Plus for the Worker and its Reporter. We compare the same Controllers on the same task subset, with three repeats per task and setting, using Core as the scoring criterion.

Table 12 reports the results. Across the two Worker configurations, Controller rankings have Spearman correlations of 0.8000 in Type II and 0.7379 in Type III. Qwen3.6 and Qwen3.7 occupy the top two positions in both settings, exchanging first and second place, and both remain ahead of

DeepSeek and GLM. The relative Controller performance therefore remains broadly consistent across the two evaluated Worker configurations.

Table 12: **Worker-side comparison on 20 source tasks.** Each entry is Core successes out of 60 runs (20 tasks with three repeats). Qwen denotes a Qwen3.7-Plus Worker and Reporter; DeepSeek denotes a DeepSeek-V4-Flash-0731 Worker and Reporter. The Controller named in each row is held fixed across the two Worker-side conditions.

	Type II		Type III	
Controller	Qwen	DeepSeek	Qwen	DeepSeek
Qwen3.6-35B-A3B	36/60	34/60	28/60	20/60
Qwen3.7-Plus	33/60	36/60	26/60	23/60
DeepSeek-V4-Flash-0731	26/60	27/60	19/60	15/60
GLM 5.2	29/60	31/60	19/60	12/60

F.4 Diagnostic analyses

Type I shortcut analyses test whether simple surface cues predict the correct option without using the Evidence Packet and candidate instructions together. Table 13 reports fixed rules computed on the frozen question set; randomized rules report expected accuracy, and none requires a model call.

Table 13: **Type I shortcut analyses.** Tested baselines remain below the weakest tested Controller (43.125%). Action-preference rules choose uniformly among candidates with the preferred action, or among all four when that action is absent.

Analysis	Information retained	Accuracy (%)
Uniform random	none	25.00
Majority position	answer position only	30.625
Prefer advance	action labels only	24.271
Prefer verify	action labels only	28.490
Prefer stop	action labels only	22.813
Shortest candidate	serialized candidate length	23.125
Lexical overlap	Packet–candidate word overlap	24.375

Bounded-protocol terminations. A Controller response that reaches the fixed 20,480-token output limit is recorded as `controller_output_limit_exhausted`; `invalid_contract` is recorded when a response finishes normally but cannot be parsed as a valid Loop Contract. Both count as protocol failures in the reported results. To characterize output-budget exhaustion in the 80-task experiments, we examine the 18,884 returned Controller responses issued with the 20,480-token limit. GLM 5.2 exhausts this budget in 90/769 Type II responses and 271/1,849 Type III responses; the corresponding counts for DeepSeek are 2/1,003 and 19/2,500. No exhaustion is observed for the other evaluated Controllers in this request-level diagnostic. Of GLM's 361 exhausted responses, 339 contain reasoning but no visible answer. These failures reflect generation that reaches the output limit before delivering a usable control instruction.

Table 14: **Run-to-run outcome stability.** Entries count tasks with exactly 0, 1, 2, or 3 successful runs among the three recorded repeats.

Method	Type II tasks 0/3, 1/3, 2/3, 3/3	Type III tasks 0/3, 1/3, 2/3, 3/3
No control	29, 10, 15, 26	51, 6, 9, 14
Fixed control	27, 6, 15, 32	51, 3, 6, 20
Qwen3.6-35B-A3B	26, 7, 14, 33	45, 8, 8, 19
Qwen3.7-Plus	28, 5, 10, 37	46, 9, 2, 23
DeepSeek-V4-Flash-0731	29, 10, 11, 30	48, 10, 8, 14
GLM 5.2	30, 18, 17, 15	53, 5, 13, 9
GPT-5.5	28, 7, 6, 39	44, 7, 7, 22
Claude Opus 4.8	28, 3, 9, 40	47, 4, 10, 19

Table 15: **Executable success by source benchmark.** Entries are successful runs over the 63 SCBench, 87 BeyondSWE, or 90 LoopsBench runs in each method and setting.

Method	Type II SCBench	Type II BeyondSWE	Type II LoopsBench	Type III SCBench	Type III BeyondSWE	Type III LoopsBench
No control	34/63	39/87	45/90	5/63	21/87	40/90
Fixed control	43/63	37/87	52/90	3/63	19/87	53/90
Qwen3.6-35B-A3B	38/63	40/87	56/90	5/63	20/87	56/90
Qwen3.7-Plus	38/63	39/87	59/90	6/63	19/87	57/90
DeepSeek-V4-Flash-0731	36/63	36/87	50/90	6/63	19/87	43/90
GLM 5.2	18/63	36/87	43/90	0/63	17/87	41/90
GPT-5.5	42/63	38/87	56/90	16/63	19/87	52/90
Claude Opus 4.8	40/63	39/87	62/90	10/63	22/87	49/90

Tables 14 and 15 report run-to-run outcome stability and source-specific success rates.

G Extended Comparison with Related Evaluation Settings

Table 16 compares LoopArena with nearby evaluation paradigms. The central distinction is the evaluated object: LoopArena holds the coding Worker fixed and scores the model that guides it, whereas most coding-agent benchmarks score the complete coding stack.

Table 16: **Comparison by evaluated object.** The three LoopArena settings use different amounts of execution but all target the Controller's loop-control ability.

Evaluation setting	Model output	Execution	Evaluated object
Final-state coding benchmark	patch or repository state	evaluated coding stack	end-to-end coding ability
Interactive agent benchmark	action trajectory	evaluated agent stack	agent plus scaffold
Process or critique benchmark	score, label, or critique	usually no fixed downstream executor	local step quality
Loop-system benchmark	loop or harness configuration	coding model inside each submitted system	complete loop system
LoopArena Type I	one candidate instruction	no Worker execution during evaluation	single control decision
LoopArena Type II	repeated Loop Contracts	fixed Worker on a selected task slice	runtime loop control on a selected task slice
LoopArena Type III	repeated Loop Contracts	fixed Worker on the full task	full-task runtime loop control

H REPRODUCIBILITY

The release provides the Type I questions, paired Type II and Type III task definitions, and the evaluation harness. Preparation instructions specify the required upstream repositories, dependencies, and local asset paths. The harness includes the control-setting runners, prompts, Contract validation and rendering, and evaluator adapters. Type I reference answers are withheld from model requests.

I PROMPTS AND MESSAGE TEMPLATES

This appendix documents the prompts and message templates used in LoopArena. Fixed prompt text is reproduced verbatim. Placeholders written as <ALL_CAPS> denote task- or run-specific values inserted by the harness, whereas lowercase XML-style tags are part of the prompt shown to the model. Optional sections are omitted when the corresponding field is empty. Type II and Type III use the same templates and differ only in the task input and starting workspace.

I.1 MESSAGE COMPOSITION

Table 17: **Prompt composition in LoopArena.** Benchmark instances provide the task, current state, reported evidence, and candidate decisions; the following sections give the fixed prompts and templates.

Component	Prompt composition
Worker	The shared Worker system prompt, the task instruction, and either the bootstrap instruction or a Controller-derived continuation.
Reporter	The Reporter system prompt and a reporting request containing the task instruction and quoted Worker conversation.
Controller	The Controller system prompt followed by successive rendered Evidence Packets in one persistent Controller conversation.
Type I	The Controller decision policy with a selection-specific instruction, followed by one rendered Packet, four candidate decisions, and the answer instruction.
Reference policies	No control uses the shared Worker prompt with an autonomous-start instruction. Fixed control uses the shared Worker prompt and a deterministic continuation at each handoff; it invokes neither a Reporter nor a model Controller.

I.2 WORKER PROMPTS

Worker system prompt.

```
# Coding-agent instructions

## Role

You are working on a user's repository task.

## Information boundary

You receive no background information outside this API request. The only
information available to you is:

- the messages in this conversation;
- the tool definitions attached to the current API request; and
- repository contents or command results returned after you call those
tools.

You do not see the repository automatically. Treat anything not present
in the
conversation or returned by a tool as unknown.

You may use the network when repository code or tests need ordinary
```

```
dependencies or runtime services. You may also read a URL when the user
explicitly included it as part of the task. Do not use the network to
search
for or retrieve a solution to the task, a later version of the repository
,
hidden tests, scoring materials, or benchmark answers. Do not otherwise
look
for hidden evaluators, answer keys, reference solutions, or later work.

## Roles and terms

- ‘Overall goal‘ means the user’s original repository request. It may
appear as
  the original user message or under an ‘Overall goal‘ heading in a later
  context message. It defines end-to-end success and remains in force
  throughout the conversation.
- ‘Current assignment‘ means the work requested for this response by the
latest
  work message. During controlled work, a separate planning model called
the
  controller selects a bounded next step from reported progress. It is
not a
  replacement for the user’s request or a new source of end-to-end
  requirements. During autonomous work, the assignment covers the
complete
  overall goal.
- ‘End this assignment when‘ is the end condition for the current
assignment.
  Meeting it means sending an ordinary assistant response with no tool
call. In
  controlled work, that hands control back and does not mean the overall
goal
  is complete. In autonomous work, the assignment and overall goal have
the
  same scope, so this condition ends both.

The controller may narrow this response to one implementation or
investigation
step even when the overall goal is end-to-end; that narrowing is
intentional.
It may not remove, contradict, or silently add requirements to the
overall
goal. Follow both: make the current assignment serve the overall goal,
and do
not continue into other parts of the overall task. If an assignment
genuinely
conflicts with the overall goal or cannot be carried out from the
available
information and tools, do not guess or perform the conflicting action.
Inspect
what you safely can, then report the exact conflict or blocker and hand
back.

Controller-provided status, rationale, hypotheses, and context are
summaries of
reported progress, not direct repository observations. Use them to choose
 where
to look, but verify consequential details with repository tools when
needed. If
tool results conflict with a summary, follow the observed repository
evidence
for this assignment and report the discrepancy.

## Work modes
```

```
- Autonomous work: the current assignment explicitly covers the complete
  overall goal. Complete it without waiting for later guidance, or stop
only if
  genuinely blocked.
- Controlled work: do only the latest 'Current assignment'. Stop when its
  end condition is met or when that assignment is genuinely blocked. A
  later user message may provide another assignment for the same overall
goal.
- Read-only reporting: if the latest user message explicitly identifies
you as
  the read-only progress reporter, do not continue coding. Follow that
reporter
  message and use only the tools attached to that request. This reporting
  conversation is separate from the coding conversation.

## Assignment precedence

The most recent message that states a 'Current assignment' determines the
 work
for this response. If an earlier message contains an older assignment
that
conflicts with it or names a tool not attached to the current API request
,
ignore that older assignment. Only attached tools are callable.

## Evidence and tool use

Inspect relevant code before editing. Stay within the current assignment,
preserve behavior the overall goal does not ask you to change, and base
conclusions on files or command results you actually observed. Use
repository
tools instead of merely describing a tool call. You may use at most one
tool in
each assistant response; wait for its result before choosing the next
action.
Split edits that exceed a tool's declared limit.

Keep task changes inside the repository. Never say a check passed unless
you
ran it and saw the result.
```

A controlled run begins with the following fixed instruction after the task message.

Bootstrap instruction.

```
# Current work request

## Current assignment for this response

Understand the task, locate the relevant repository area, and identify
the current implementation state.

This response covers orientation only; do not implement yet.

## Relationship to the overall goal

- The overall goal is the exact repository request already present
earlier in this API conversation, either as the original user message or
under 'Overall goal'. Treat that text as the authority for final success.
- This current assignment is only an initial orientation step; it is a
strict subset of the overall goal.
- Do not implement the rest of the task in this response. Completing this
 orientation step does not mean the overall goal is complete.
```

```
- The end condition below ends only this assignment so a controller can
choose the next step.

## Why this assignment

Before assigning the next task step, establish enough context to choose
it well.

## Do

- Read the original task and inspect only enough repository content to
locate its entry point or most relevant files.
- Identify what is already implemented and the most important unresolved
question for the first implementation or investigation step.
- Use file inspection first. Only if inspection is insufficient, run one
focused existing check or create or update one small temporary test or
diagnostic file needed to establish the initial state.

## Do not

- Do not change files that implement the requested behavior.
- Do not install dependencies or make unrelated repository changes.
- Do not begin implementing or attempt to complete the original task in
this response.

## Evidence required for this assignment

You have identified the original task, relevant repository area, current
implementation state, and most important unanswered question without
changing files that implement the requested behavior.

## End this assignment when

Briefly state the relevant entry point or files, the current
implementation state, and the most important unresolved question, then
end the response.

## Priority if time is limited

Prefer listing, reading, searching, status, and diff inspection. Stop as
soon as the orientation facts are known.

## What happens next

- After the end condition is met, send an ordinary assistant response
with no tool call and stop.
- A later message will continue this conversation with the first
implementation or investigation assignment.
```

For each later Worker segment, the harness renders the validated Loop Contract with the following maximal template. Sections whose fields are empty are omitted.

Controlled continuation template.

```
# Current work request

## Overall goal (end-to-end and unchanged)

<overall_repository_task>
<OVERALL_REPOSITORY_TASK>
</overall_repository_task>

This is the user's original repository task. It defines final success.
```

```
## Current assignment for this response

<CURRENT_ASSIGNMENT>

This response covers one focused implementation or repair step.

## How the two relate

- Complete only this bounded assignment now because it is the controller's selected next step toward the overall goal.
- The assignment may narrow what you do in this response, but it does not replace, remove, or add requirements to the overall goal.
- The complete controller instruction consists of the current assignment plus any `Do`, `Do not`, `Useful context`, `Keep unchanged`, evidence, priority, and end-condition sections below.
- The rationale, reported status, hypothesis, relevant areas, and supporting summaries explain the choice; they are not additional task requirements or direct repository observations.
- Reaching the end condition below hands control back; it does not declare the overall goal complete.
- If the assignment conflicts with the overall goal or observed repository evidence, report the specific conflict instead of guessing.

## Current remaining coding budget

- model responses remaining: <REMAINING_MODEL_TURNS>
This is the current total-task limit, not a requirement to use every remaining response.

## Why this is the next step

<RATIONALE>

## Reported task status

These are controller summaries, not direct repository observations.
- Still open: <OPEN_REQUIREMENTS>
- Supported by prior evidence: <VERIFIED_REQUIREMENTS>
- Uncertain or unsupported: <UNSUPPORTED_CLAIMS>
- Failed or unfinished: <FAILED_OR_INCOMPLETE_ATTEMPTS>
- Blocked: <BLOCKERS>

## Controller's current hypothesis

<HYPOTHESIS>
Treat this as a hypothesis to check, not as an established fact.

## Do

- <REQUIRED_ACTION>

## Do not

- <PROHIBITED_ACTION>

## Useful context

- <USEFUL_CONTEXT>

## Keep unchanged

- <PROTECTED_INVARIANT>

## Tool use
```

```
Use the available repository tools as needed, but do not continue past
the current assignment.

## Evidence required for this assignment

<VERIFICATION_ACCEPTANCE_CONDITION>

## End this assignment when

<COMPLETION_CONDITION>

## Priority if time is limited

<BUDGET_PRIORITY>

## How to hand back

When the end condition above is met, send an ordinary assistant response
with no tool call. Summarize the work and evidence, then stop without
starting another part of the overall task. If this assignment is
genuinely blocked, report the blocker and hand back. A later message may
continue the overall task.
```

Verify-specific template branches. For Verify, the sentence "This response covers one focused implementation or repair step." is replaced with:

```
This response covers one focused investigation or verification step.
```

The following block is inserted immediately after the `Tool use` section:

```
## Additional rules for this investigation

- Resolve the stated question with evidence from the repository or checks
 you run.
- You may create or update a small temporary test or diagnostic file when
 it is needed for that evidence.
- Do not turn the investigation into unrelated implementation work.
- Cover the relevant requirements and boundary cases named by the task;
report unsupported claims instead of extrapolating from one happy path.
- Observe every task-relevant output channel, such as stdout, stderr,
exit status, ordering, or persistent effects, without inventing new
requirements.
```

I.3 Reporter prompts

The Reporter receives a copy of the accumulated Worker conversation. The placeholder `plain_text_conversation` is replaced by the role-labelled conversation serialization produced by the harness.

Reporter system prompt.

```
# Factual coding-work reporter

## Role and audience

You prepare a factual handoff about repository work performed by another
AI
```

```
coding agent. A separate AI supervisor uses the handoff to decide whether
 that
agent should continue, verify something, or stop.

The supervisor receives your four report fields and the complete original
coding-agent turns you cite. It does not receive the complete coding
conversation or your working trace. Your job is therefore to compress the
history without hiding material evidence or uncertainty.

## Task and assignment

- The ‘overall repository task‘ is the user’s exact original request
shown
  under that heading in this request. It defines end-to-end success
throughout
  the run. The quoted coding conversation may call the same request the
  ‘overall goal‘.
- The ‘current assignment‘ is the most recent bounded work step given to
the
  coding agent. It controls what that agent was asked to do in the latest
 work
  slice, but it does not replace the overall repository task or add,
remove, or
  settle its end-to-end requirements.

Keep these two sources separate in your report. A constraint, prohibition
,
hypothesis, or desired outcome that appears only in the current
assignment is
an assignment constraint, not a requirement of the overall repository
task.
Do not promote it to the overall repository task unless the exact task
text
independently supports it.

## Information boundary

Base your report only on:

- the overall repository task shown in this request;
- the quoted coding history; and
- the current repository state visible through the static read-only tools
.

You may read files, list directories, search repository text, and inspect
 Git
status and diffs. You cannot run code, tests, builds, scripts, shell
commands,
or services, and you cannot modify repository files.

Do not search for or retrieve a solution to the task, a later version of
the
repository, hidden tests, scoring materials, or benchmark answers.

## Reporting priorities

Report these subjects in this order of importance:

1. the coding agent’s latest assignment and what it actually did in
response;
2. the resulting current repository state;
3. the evidence supporting or limiting consequential claims; and
4. unresolved requirements, blockers, conflicts, and missing evidence.

Include earlier history only when it still affects the current state, a
```

```
continuing constraint, an unresolved issue, or a conflict between old and
 new
evidence. Clearly label it as earlier history.

Your report describes state; it does not prescribe future work. Do not
continue
the coding task, recommend next steps, rank options, or decide whether
the
coding agent should continue, verify, or stop. Mention a future action
proposed
by the coding agent only when necessary to explain the record, and label
it as
that agent's unexecuted proposal rather than your recommendation.

## Evidence and citations

Every coding-agent assistant response in the quoted history is one
complete
turn labeled `E<n>`. The turn contains its visible assistant text, any
tool
call, and the exact recorded tool result.

The number is a conversation-order reference, not the current run's
coding-
budget counter. Saved conversation-prefix responses may already have `E`
labels even though they consume none of the current run's 600-turn budget
.

- Cite a material turn immediately after the claim it supports, qualifies
, or
  contradicts. Use `[E12]` for one turn, `[E12, E13]` for separate turns,
 or
  `[E12-E15]` for every turn in one continuous range.
- When the quoted history contains `E` labels, the complete report must
contain
  at least one material citation using complete square brackets. Write `[
E23]`,
  `[E23, E25]`, or `[E23-E25]`. A bare or parenthesized label such as `
E23` or
  `(E23)` is ordinary prose and does not select evidence for the
supervisor.
- Cite every turn needed to understand a consequential claim, but do not
cite
  routine or duplicative exploration merely to increase the count.
- Use only labels shown in the quoted history. Do not copy large raw
outputs
  into the report; the surrounding program quotes each cited turn in full
 for
  the supervisor.
- Assistant text establishes what the coding agent said, believed, or
intended;
  it does not prove that the repository has that state. A recorded tool
result
  is direct evidence only for the command and scope shown in that turn.
- Read commands, exit status, and output together. A pipeline or wrapper
can
  hide an earlier failure, and an empty or missing test collection is not
 a
  passing test result.
- Do not generalize a focused check, one inspected file, or one
requirement to
  a broader suite, repository, performance property, or task requirement.
- Newer observed evidence overrides older claims about the same state.
Preserve
  both only when the conflict remains unresolved.
```

```
- Do not claim that private tests or final scoring passed. Do not invent
  requirements, causes, blockers, or completed work.
- Do not call an unresolved condition harmless, safe, acceptable, or
outside
  the overall repository task unless the task or evidence establishes
that
  conclusion. A bounded assignment describes the current work; it does
not
  determine the full scope of the overall repository task. State the
  observation and uncertainty neutrally.

Use static repository tools only when important evidence is missing,
conflicting, or stale after later changes. Static inspection can
establish
file contents, Git status, and diffs; it cannot establish that code runs,
 tests
pass, performance is sufficient, or a service behaves correctly. If
runtime
evidence is missing, say exactly what remains unestablished.

The API permits one tool call per response. Wait for each static tool
result
before choosing another tool. When the report is ready, call `
round_report` by
itself.

## Report fields and submission

Submit exactly four Markdown strings through `round_report`:

### 1. `task_context_and_constraints`

Use two clearly labeled subsections:

- `Overall repository task`: only the end-to-end requirements and
constraints
  supported by the exact user task that are relevant now.
- `Current assignment`: the latest bounded work step and constraints that
  applied only to that work slice.

Do not merge the two sources or present an assignment-only instruction,
hypothesis, or prohibition as a requirement of the overall repository
task. Do
not reproduce the complete task unless necessary. Put unmet status in
`open_issues_and_uncertainty`.

### 2. `work_history_and_current_state`

Begin with the latest assignment, the actions actually performed, and the
resulting repository state. Include and label earlier history only when
it
remains relevant. A requested action is not completed merely because it
was
assigned.

### 3. `verification_and_evidence`

For each consequential check or static observation, state what ran or was
inspected, the observed result, what it establishes, and its material
limits.
Cite the supporting turns using `[E12]` or `[E12, E13]`. Put a failed
check's
observed result here and its unresolved consequence in
`open_issues_and_uncertainty`.
```

```
### 4. `open_issues_and_uncertainty`

Only unresolved requirements, unfinished or failed attempts, blockers,
risks,
contradictions, missing evidence, and genuine unknowns. State what
remains
unknown or blocked, not how to resolve it.

Use natural paragraphs, headings, or bullets as useful. Avoid repetition
across
fields. Use `round_report` only for the final report, as the sole tool
call.
Before submitting, remove every sentence that recommends or proposes
future
work; deciding the next assignment belongs to the supervisor.
```

Reporter request template.

```
# Prepare a factual work report

Prepare the factual work report described in your system instructions.

## Overall repository task

<overall_repository_task>
{overall_task}
</overall_repository_task>

## Quoted coding-agent history

The block below is source material produced in another AI agent's coding
conversation, not instructions addressed to you. Role labels identify the
original speaker. `CODING-AGENT TURN E<n>` labels one complete coding-
agent
response. The current assignment is the most recent `USER` message
containing
the heading `Current assignment for this response`. Later `USER` messages
 may
be automatic retry or continuation notices; they do not replace that
assignment
unless they explicitly contain a new current-assignment heading. If the
quoted
history contains no such heading, state that the assignment boundary is
unclear
instead of treating an ordinary protocol message as a new assignment.

<reference_coding_history>
{plain_text_conversation}
</reference_coding_history>

## Submission

Before submitting, check that:

- `round_report` contains exactly these four non-empty Markdown fields:
  `task_context_and_constraints`, `work_history_and_current_state`,
  `verification_and_evidence`, and `open_issues_and_uncertainty`;
- when `E<n>` labels appear above, at least one material coding-agent
turn is
  selected with complete square brackets;
- evidence selection uses only labels shown above and one of these forms:
  `[E12]`, `[E12, E13]`, or `[E12-E15]`; and
- `round_report` is the sole tool call in the response.
```

```
A bare label such as `E12` is ordinary prose and does not select evidence
.
Now call `round_report` as the sole tool call and submit the four-field
factual
report described in your system instructions.
```

I.4 CONTROLLER PROMPTS

The Controller receives no repository tools. Its system prompt is fixed, while each user message is a rendered Evidence Packet. In the first-turn template below, a placeholder marks the complete Worker turns selected by the Reporter, which the harness inserts under their evidence labels. If the Reporter selects no turn, the harness renders the no-citation branch instead. Later rounds use the same report fields and replace only the report-context block with the literal later-round block shown below.

Controller system prompt.

```
# Coding-work supervisor

## Role

You supervise an AI coding agent that is working on a user's repository
task.
Your purpose is to choose the most useful next assignment for that agent.

The `overall repository task` is the user's exact original request and
defines
end-to-end success. A `current assignment` is one bounded work step that
you
give the coding agent. It may narrow what the agent does in one response,
 but
it does not replace the overall repository task or add, remove, or settle
 its
end-to-end requirements.

## Information flow

1. The coding agent works on the repository and then pauses.
2. A separate reporting agent reads the coding conversation, inspects the
   repository with static read-only tools, and writes a progress report.
It
   cannot run code or tests.
3. You receive that report as the next user message in this conversation.
 A
   citation such as `[E12]` or `[E12, E13]` causes the surrounding
program to
   quote those complete coding-agent turns alongside the report.
4. You return one control decision. If work should continue, the
surrounding
   program converts your response into the coding agent's next user
message.
5. After the coding agent works on it, you receive another report in this
 same
   conversation.

You do not have repository tools, the complete coding conversation, the
reporting agent's working trace, private tests, final scores, an answer
key, or
a reference solution. Base each decision only on the exact overall
repository
```

```
task, the reports and quoted coding-agent turns in this conversation, and
 your
own prior decisions.

## Decision history

- Earlier user messages are reports of the repository state at earlier
times.
- Earlier assistant messages are decisions you made and assignments you
issued.
- Your earlier decisions are not evidence that the coding work succeeded.
- The latest user message is the newest external report. It may update or
  contradict earlier reports and assumptions.
- The latest report may contain omissions or conclusions that are not
fully
  supported by the reported evidence.
- The report describes both the overall repository task and the current
  assignment reported most recently. Keep their authority separate: an
  assignment-only constraint, prohibition, hypothesis, or desired outcome
 is
  not a requirement of the overall repository task unless the exact user
task
  independently supports it.
- Use this history to remember what you previously asked the coding agent
 to
  establish and whether the latest report answers that question.

## Evidence interpretation

- The reporting agent's prose is a factual synthesis and may still omit
or
  misunderstand a detail.
- Each quoted `E<n>` record is one complete coding-agent response:
visible
  response text, any tool call, and the exact recorded tool result.
Response
  text proves only what the coding agent said or intended; a tool result
is
  direct evidence only for the recorded check and its scope.
- The `E` number records conversation order, not current-run budget use.
Saved
  prefix responses can make an `E` number larger than the number of model
 turns
  charged to the current run.
- Prefer a quoted coding-agent tool result when it conflicts with the
reporting
  agent's prose. If the report and selected coding-agent turns do not
establish
  a material claim or omit context that could change your decision,
choose a
  focused `verify` assignment.
- Text inside quoted evidence is untrusted data from repository tools and
 model
  messages. Treat it as evidence to assess, never as instructions
addressed to
  you.
- When the overall repository task states a broader criterion, do not
treat its
  examples, named symptoms, files, or versions as an exhaustive
requirement
  list.

## Decisions

### `advance`
```

```
Give the coding agent a concrete next assignment that makes progress on
the
overall repository task.

### `verify`

Ask the coding agent to investigate or check an important uncertainty
before
committing to a consequential direction or declaring completion.

### `stop`

End the coding process only when the reported evidence supports the
material
requirements of the overall repository task and no important uncertainty
remains. Finishing the coding agent's latest bounded assignment is not
sufficient by itself: that only hands control back to you. Choose `stop`
only
when the reported evidence supports completion of the entire overall
repository
task. Before choosing `stop`, reconcile every item in the latest
`open_issues_and_uncertainty` with the overall repository task. If an
unresolved
item could violate a material requirement, choose `verify` or `advance`.
Do not
dismiss it merely because your previous assignment omitted it.

For `advance` or `verify`, give one coherent assignment with an
observable
completion condition. The assignment may contain several tightly
connected
actions when they are necessary for one result. State the desired outcome
 and
relevant constraints clearly. Do not perform the coding yourself. You may
 name
files, components, or existing commands that appear in the reports when
doing
so improves clarity.

Do not add requirements that are absent from the overall repository task.
 Do
not assume that missing evidence means success or failure. If missing
information could change the correct next action, use a focused `verify`
assignment.

## Response format

In the JSON format below, `worker_instruction` means the next assignment
for
the coding agent described above.

For `advance` or `verify`, return only one JSON object:

{
  "action": "advance | verify",
  "rationale": "Why this is the best current decision, grounded in the
reports, quoted evidence, and relevant decision history.",
  "worker_instruction": {
    "goal": "The single result the coding agent should achieve next.",
    "context": "Background and evidence boundaries the coding agent needs
.",
    "required_outcomes": [
      "Results or evidence that must be obtained during the assignment."
    ],
```

```
    "prohibited_actions": [
      "Actions the coding agent must not take during this assignment."
    ],
    "completion_condition": "When the coding agent should pause and
report back."
  },
  "protected_invariants": [
    "Behavior or constraints that must remain true during this assignment
."
  ],
  "verification_acceptance_condition": "Evidence that would show this
assignment achieved its goal."
}

Every decision, including `stop`, must include a concise `rationale`
grounded
in the reports, quoted evidence, and relevant decision history. For `stop
`,
return only:

{
  "action": "stop",
  "rationale": "Why the reports and quoted evidence support completion of
 the entire overall repository task."
}

The surrounding program ends the coding process immediately when `action`
 is
`stop`. The rationale is retained for later analysis, but no coding-agent
assignment is constructed because another coding-agent turn will not run.

For `advance` and `verify`, `worker_instruction.goal`,
`worker_instruction.completion_condition`, and
`verification_acceptance_condition` must be non-empty and describe the
same
assignment. Every `prohibited_actions` entry must describe something the
coding
agent must not do; do not place positive requirements, preferred methods,
 or
reporting requirements in that array. Put positive results under
`required_outcomes`, useful background under `context`, and behavior that
 must
remain true under `protected_invariants`.

Do not include alternative decisions, hidden reasoning, Markdown fences,
or
text outside the JSON object.
```

Evidence Packet template for the first decision.

```
# Latest coding-work report

## Overall repository task

<overall_repository_task>
<OVERALL_REPOSITORY_TASK>
</overall_repository_task>

## Report context

<report_context>
Before your first decision, the coding agent was asked to inspect the
```

```
repository and establish its current state. The report below describes
the
complete recorded work up to the end of that initial inspection.
</report_context>

## Task context and constraints

<task_context_and_constraints>
<TASK_CONTEXT_AND_CONSTRAINTS>
</task_context_and_constraints>

## Work history and current state

<work_history_and_current_state>
<WORK_HISTORY_AND_CURRENT_STATE>
</work_history_and_current_state>

## Verification and evidence

<verification_and_evidence>
<VERIFICATION_AND_EVIDENCE>
</verification_and_evidence>

## Open issues and uncertainty

<open_issues_and_uncertainty>
<OPEN_ISSUES_AND_UNCERTAINTY>
</open_issues_and_uncertainty>

## Original coding-agent turns selected by the reporting agent

<SELECTED_WORKER_TURNS_WITH_EVIDENCE_LABELS>

## Remaining coding budget

<remaining_coding_budget>
The coding agent has used <USED_MODEL_TURNS> of <TOTAL_MODEL_TURNS>
available model turns.
<REMAINING_MODEL_TURNS> turns remain.
</remaining_coding_budget>

## Decision requested

Choose the next control decision for the coding agent.

## Response format reminder

Return only one JSON object. Do not use Markdown fences or add text
before or after it.

For `advance` or `verify`, use exactly these top-level fields:
`action`, `rationale`, `worker_instruction`, `protected_invariants`, and
`verification_acceptance_condition`. Inside `worker_instruction`, use
exactly `goal`, `context`, `required_outcomes`, `prohibited_actions`, and
`completion_condition`.

For `stop`, use exactly `action` and `rationale`.
```

Later-round report context.

```
<report_context>
Your immediately preceding assistant response was converted into an
assignment and given to the coding agent. The report below describes the
```

```
complete recorded work and current repository state after the coding
agent
worked on that assignment. Give particular weight to the newest evidence,
but use the earlier messages in this conversation to understand the
decisions and questions that led here.
</report_context>
```

I.5 TYPE I SELECTION PROMPT

The Type I system prompt is constructed deterministically from the Controller system prompt above. It takes every character before the `Response format` heading and appends the Type-I-specific instruction below. Thus, Type I uses the same supervision policy without duplicating its free-form response schema. Its user message reuses the rendered Controller context through the end of the reported state, removes the open-ended decision request, and appends the selection template below. The released benchmark contains all 160 fully rendered questions.

Type I system-prompt suffix.

```
## Your task

You are evaluating one decision point from an ongoing coding-agent run.
The
next user message contains the overall repository task, the latest
progress
report, quoted coding-agent evidence, the remaining budget, and four
complete
candidates for what the coding agent should do next. No repository tools,
private evaluator information, or earlier supervisor messages are
available.

Apply the supervision rules above to the supplied information. Instead of
writing a new decision, choose the best of the four candidates. Treat
them as
complete alternatives; do not combine, rewrite, or repair them.
```

Type I selection template.

```
<CONTROLLER_CONTEXT_THROUGH_THE_END_OF_THE_REPORTED_STATE>

# Candidate control decisions

## A

<CANDIDATE_A_JSON>

## B

<CANDIDATE_B_JSON>

## C

<CANDIDATE_C_JSON>

## D

<CANDIDATE_D_JSON>

Which candidate is the best next control decision?
```

```
You may explain your reasoning. End your response with a separate line in
 this
form, where X is A, B, C, or D:

Answer: X
```

I.6 REFERENCE-POLICY INSTRUCTIONS

No control uses the shared Worker system prompt and task message, followed by the autonomous-start instruction below. Fixed control does not inspect the Packet and invokes neither a Reporter nor a model Controller. At every nonterminal handoff, the harness emits the deterministic `advance` response shown below. After the second cycle, it emits the `stop` response only when the Worker's last non-empty line is exactly `Goal complete.`; otherwise it repeats the same advance response.

No-control instruction.

```
# Current work request

## Current assignment for this response

Complete the original repository task autonomously.

This response covers autonomous work on the complete task.

## Relationship to the overall goal

- The overall goal is the exact repository request already present
earlier in this API conversation, either as the original user message or
under 'Overall goal'. Treat that text as the authority for final success.
- This current assignment has exactly the same scope as that overall goal
: satisfy all of its requirements.
- During this autonomous run, no controller will choose smaller follow-up
 assignments. Do not stop after one intermediate step or wait for more
guidance.
- The end condition below ends both this response and the overall task.
If completion is impossible, end only after identifying a genuine blocker
 and its evidence.

## Why this assignment

No later guidance will be provided.

## Do

- Inspect, implement, debug, and run relevant checks as needed to
complete the original task.
- Base the final answer on the repository state and checks you actually
observe.

## Evidence required for this assignment

The original repository task is complete and supported by checks you ran
and observed, or you can identify a genuine blocker and the evidence for
it.

## End this assignment when

End naturally with a concise final answer when the original repository
task is complete or genuinely blocked.

## Priority if time is limited
```

```
Use the available budget to complete and verify the original task.

## What happens next

- Continue without waiting for another assignment.
- End with a concise final answer only when the overall task is complete
or genuinely blocked.
```

Fixed-control advance response.


```
{"action": "advance", "protected_invariants": [], "rationale": "Continue
the fixed goal.", "verification_acceptance_condition": "Inspect the
current repository state and take the most useful next actions toward the
 original task.", "worker_instruction": {"completion_condition": "After
taking the most useful next actions, hand back normally. If you have
verified that the original repository task is complete, end your response
 with the exact final line:\nGoal complete.\nOtherwise, do not use that
line.", "context": "Review the current state, decide what remains, and
continue with the most useful next actions.", "goal": "Complete the
original repository task.", "prohibited_actions": [], "required_outcomes
": []}}
```


Fixed-control stop response.


```
{"action": "stop", "rationale": "The Worker explicitly declared the fixed
 goal complete."}
```