

LongCLI-Bench: A Preliminary Benchmark and Study for Long-horizon Agentic Programming in Command-Line Interfaces

Yukang Feng^{1,2,3*} Jianwen Sun^{1,2,3*} Zelai Yang^{1*} Jiaxin Ai^{2,4} Chuanhao Li⁴
Zizhen Li^{1,2,3} Fanrui Zhang² Kang He² Rui Ma⁵ Jifan Lin⁵ Jie Sun²
Yang Xiao⁶ Sizhuo Zhou² Wenxiao Wu² Yiming Liu²
Pengfei Liu^{2,5} Yu Qiao^{2,4} Shenglin Zhang¹ Kaipeng Zhang^{2,3†}

¹ NKU ² SII ³ Shanda AI Research Tokyo ⁴ Shanghai AI Laboratory ⁵ SJTU ⁶ PolyU

Project Page: <https://github.com/finyorko/longcli-bench>

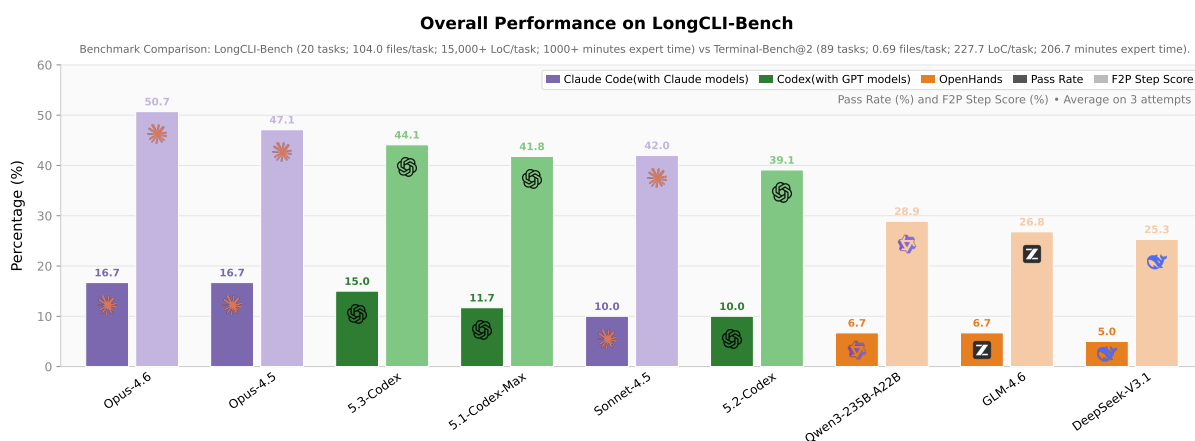


Figure 1: **Overall Performance on LongCLI-Bench.** The comparison reveals a significant capability gap in long-horizon agentic programming, with all evaluated agents achieving less than a 20% pass rate. While agents maintain relatively high P2P step scores, the low F2P fulfillment highlights critical challenges in maintaining execution consistency over extended command-line interactions.

Abstract

Recent advances in AI-assisted programming have empowered agents to execute complex workflows via command-line interfaces, however, existing benchmarks are limited by short task horizons, data contamination from GitHub scraping, and a lack of fine-grained evaluation metrics, fail to rigorously evaluate the long-horizon planning and execution capabilities essential for realistic software engineering. To address these gaps, we introduce **LongCLI-Bench**, a comprehensive benchmark designed to evaluate agentic capabilities across **long-horizon**, realistic tasks. We curated 20 high-quality, long-horizon tasks from over 1,000 computer science assignments and real-world workflows, covering four engineering categories: from scratch, feature addition, bug fixing, and refactoring. We propose a dual-set testing protocol for LongCLI-Bench, which measures requirement fulfillment (*fail*→*pass*) and regression avoidance (*pass*→*pass*), and incorporates step-level scoring to pinpoint execution failures. Extensive experiments reveal

that even state-of-the-art agents achieve pass rates below 20% in LongCLI-Bench. Step-level analysis further indicates that the majority of tasks stall at less than 30% completion, highlighting that critical failures often occur in the early stages. Although self-correction offers marginal gains, human-agent collaboration through plan injection and interactive guidance yields significantly higher improvements. These results highlight that future research must emphasize the development of synergistic human-agent workflows alongside advances in agents' planning and execution capabilities to overcome key challenges in long-horizon task performance.

1 Introduction

The paradigm of AI-assisted programming is undergoing a fundamental transition from code generation to autonomous software engineering. Modern agents such as SWE-agent (Yang et al., 2024a), OpenHands (Wang et al., 2024b), and commercial CLI assistants (OpenAI, 2025; Anthropic, 2025;

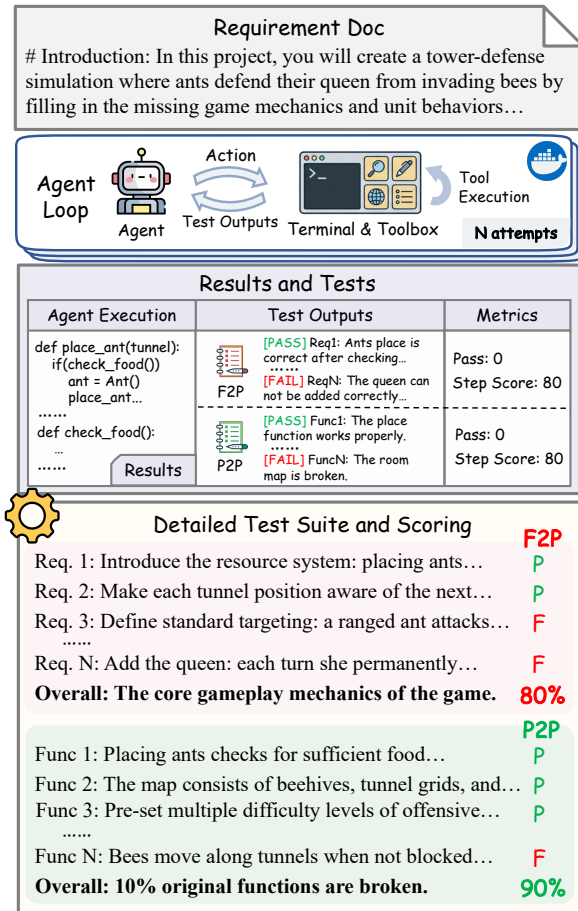


Figure 2: Task sample in LongCLI-Bench.

Google, 2025) can now function like human engineers: they plan architectures, navigate repositories, manage development environments, and execute multi-step workflows via Command Line Interfaces (CLI) based on environment feedback (Yao et al., 2022; Shinn et al., 2023; Madaan et al., 2023).

Despite the advancements, evaluation benchmarks have lagged behind. Early benchmarks, such as HumanEval and MBPP (Chen et al., 2021; Austin et al., 2021), assess only algorithmic logic within isolated code snippets. While subsequent benchmarks like SWE-bench (Jimenez et al., 2024; OpenAI, 2024) extend to repository-level, they remain restricted to short, single-category tasks, and are compromised by data contamination from scraped GitHub repositories. Recent Terminal-Bench (Institute, 2025) introduces interactive sandbox environments with community-contributed tasks, yet its tasks remain simple with short horizons, and lack fine-grained evaluation to deliver effective feedback or diagnose failure modes.

In real-world scenarios, software engineering tasks are inherently long-horizon with continuous

requirements, yet existing benchmarks ignore such sequential dependencies, failing to evaluate the ability to maintain long-term environmental context and navigate complex, interdependent workflows.

To address the above limitations, we introduce **LongCLI-Bench**, a comprehensive benchmark designed to evaluate **long-horizon** agentic capabilities in realistic CLI environments, covering diverse task categories and providing fine-grained evaluation. Unlike benchmarks that focus on a single task type, LongCLI-Bench spans four engineering categories: **From Scratch** ($0 \rightarrow 1$) which involves building projects from the ground up, **Feature Addition** ($N \rightarrow N + 1$) which focuses on enhancing existing codebases, **Bug Fix** ($No \rightarrow Yes$) which addresses diagnosing and resolving bugs, and **Refactor** ($A \rightarrow A'$) which involves optimizing or restructuring code. We curate tasks from two distinct sources: 958 computer science (CS) assignments and 50 real-world workflows. This approach not only mitigates the data contamination risk inherent in GitHub-derived datasets but also ensures that the benchmark encompasses realistic, complex, long-horizon tasks that align with genuine development scenarios. To identify tasks that effectively probe the functional boundaries of agents, we manually crafted test inputs to execute these samples via Codex (OpenAI, 2025), followed by a manual inspection of the execution results. We found that existing agents already exhibit high proficiency in handling the majority of routine assignments, so we eliminated these tasks and focused on the remaining complex tasks. For each task, we manually created the requirement documentation, isolated environments, and test suites and finally filtered the pool into 20 high-quality tasks.

Each task includes an initial repository and a requirement document in an isolated environment. The evaluation utilizes a dual-set protocol: *Fail*→*Pass* (F2P) tests verify that the agent has successfully implemented the new requirements, while *Pass*→*Pass* (P2P) tests ensure that the agent’s modifications do not break existing system functionalities (detecting regressions). Moreover, we provide *step-level scores*, enabling fine-grained measurement of partial progress and pinpointing exactly where long workflows break.

The experiment results show that long-horizon CLI work remains far from solved: all agent systems achieve $< 20\%$ pass rate (e.g., Claude Code (Anthropic, 2025) with Claude-Opus-4.6 and Codex with GPT-5.3-Codex), while step-level

scores reveal that failures are predominantly concentrated in the early stages of tasks. Furthermore, although self-correction improves performance by leveraging error feedback, it still lags behind methods involving human plan injection or human-agent interaction. This finding strongly indicates that strategic planning and execution proficiency remain the key bottlenecks for current autonomous agents. Our contributions are summarized as follows:

LongCLI-Bench: A curated benchmark of 20 long-horizon tasks, filtered from >1,000 samples via manual curation and rigorous validation process to ensure sufficient long-horizon and complex, covering four distinct engineering categories.

Dual-Set Tests with Step-Level Scores: Assesses both the requirement implementation (F2P) and the regressions (P2P). Additionally, we introduce step-level metrics to measure the degree of task completion, enabling a clear distinction between early-stage failures and near-success outcomes.

Experiment Results: All agents yield pass rates below 20%, identifying planning and execution proficiency as the key limitations hindering autonomous software engineering performance.

Rather than focusing solely on boosting agents' standalone planning and execution capacities, future research should prioritize core engineering proficiencies, long-horizon contextual consistency maintenance, and effective human collaboration to navigate complex engineering tasks.

2 Related Work

2.1 Agents and Coding-Oriented LLMs

Agentic coding performance is shaped jointly by the LLMs (He et al., 2026, 2025b; Ding et al., 2025b; Wang et al., 2024a; He et al., 2025a) and the agent scaffold. Firstly, code-oriented LLMs (Rozière et al., 2023; Li et al., 2023; Seed et al., 2025; Huang et al., 2024) such as DeepSeek-Coder (Guo et al., 2024) and Qwen3-Coder (Yang et al., 2025a) have demonstrated considerable capabilities in code understanding and generation. Secondly, SWE-agent (Yang et al., 2024a) can interact with real repositories by interfaces, while OpenHands (Wang et al., 2024b) provides an end-to-end framework; other systems explore different avenues, including structured localization and repair (Zhang et al., 2024), multi-agent optimization (Huang et al., 2023), planning coding (Bairi et al., 2023), baselines (Xia et al., 2024), and retrieval-augmented completion (Zhang et al.,

2023). In practice, commercial CLI assistants (OpenAI, 2025; Anthropic, 2025; Google, 2025), such as Codex, integrate innovative paradigms including general tool-use and self-refinement mechanisms (Yao et al., 2022; Schick et al., 2023; Shinn et al., 2023; Madaan et al., 2023), demonstrating strong intelligent assistance capabilities.

2.2 Coding/SWE Benchmarks

The evaluation of code generation has evolved from simple function-level tasks to increasingly complex, environment-grounded repository-level challenges. Early benchmarks measured function-level (Chen et al., 2021; Austin et al., 2021; Hendrycks et al., 2021) and class-level (Du et al., 2024) tasks, which are useful for assessing local correctness but largely overlook cross-file planning. Consequently, benchmarks expanded to more diverse suites and domains (Lu et al., 2021; Lai et al., 2022; Cassano et al., 2022; Zhuo et al., 2024), freshness-oriented evaluation (Jain et al., 2024), and tasks requiring explicit cross-file reasoning (Liu et al., 2023; Ding et al., 2023).

At the repository level, the SWE-bench series (Jimenez et al., 2024; OpenAI, 2024; Zan et al., 2025; Yang et al., 2024b; Yu et al., 2025; Yang et al., 2025b) evaluates real github issue resolution via patches that satisfy project tests. Complementary benchmarks extend to feature and long-horizon tasks (Li et al., 2025; Ni et al., 2025; Deng et al., 2025), repository generation (Ding et al., 2025a), and test generation (Mündler et al., 2024). Meanwhile, environment-centric benchmarks (Xiao et al., 2025; Fu et al., 2025; Chan et al., 2024) emphasize reproducible execution. But these benchmarks are mined from GitHub and are susceptible to contamination (Xu et al., 2024; Deng et al., 2024). Terminal-Bench (Institute, 2025) provides a standardized sandbox with community-contributed tasks for evaluating agents' coding and terminal capabilities. However, it is limited to short-duration tasks and offers only binary pass/fail feedback, which prevents a detailed analysis of the agent's performance. Furthermore, precisely because these tasks lack the complexity and continuity inherent to real-world workflows, existing benchmarks are increasingly facing saturation. We therefore introduce LongCLI-Bench, a benchmark of multi-category, long-horizon real-world tasks with step-level evaluation.

3 LongCLI-Bench

LongCLI-Bench is designed around five principles: long-horizon, contamination control, clear requirements, solvability, and isolated environments.

3.1 Data Construction

3.1.1 Data Sources and Initial Codebase

Most existing benchmarks scrape tasks from popular GitHub repositories, but this increases contamination risk. We avoid direct mining and instead curate tasks from two primary sources: CS course assignments and real-world workflows.

CS course assignments are rigorously designed by domain experts and typically feature clear specifications alongside realistic codebases. However, establishing a benchmark that genuinely challenges modern agents requires extensive filtration. To ensure high task quality and alignment with our benchmark goals, we undertook a labor-intensive manual curation process. Specifically, we first identified and collected **958** assignments from **108** courses spanning diverse domains such as operating systems, networks, and data processing. For each assignment, we organized the task documentation and executed it using Codex. Since most assignments lack automated evaluation metrics, we then manually assessed the solutions generated by Codex to verify task completion status, which allowed us to evaluate the feasibility and difficulty of each task. Through thorough manual review, we found that Codex excels at most CS homework and many project-level tasks. As a result, we eliminated tasks that were either too easy or hard to evaluate consistently and only a small subset of tasks that required complex, multi-step engineering skills were kept.

Real-world Research & Engineering Workflows. These are long-chain tasks derived from actual research and work scenarios, where all information is manually constructed. Inspired by daily development activities, such as writing code, configuring environments, and building data processing pipelines, these tasks represent a sequence of complex, continuous sub-tasks that mirror realistic workflows. Crucially, they exhibit strong sequential dependency, forming a complete requirement chain where the outcome of previous steps determines the feasibility of subsequent ones.

3.1.2 Requirement Document

Based on the specific task objectives, we craft a requirement document that explicitly defines both

the functional goals and the entry point specifications. This document must clearly articulate the functionality to be implemented. This avoids false negatives where correct logic fails testing due to arbitrary, unlocatable entry points.

For CS course assignments that come with existing descriptions, we partially rewrite the requirements by replacing specific variable/function/file names, and background stories to prevent simple retrieval matching, thereby reducing contamination risks while enhancing readability. For tasks derived from real-world scenarios, the requirements are entirely manually crafted.

3.1.3 Environment and Solution Codebase

We concurrently construct the execution environment and solution codebase within Docker container. This process involves iteratively solving the task from a base docker image while recording necessary dependencies. The final outputs are rigorously separated: dependencies are solidified into a Dockerfile to provide a consistent base environment, while the human solution path form the solution repo to verify solvability.

3.1.4 Test Suite and Scoring Design

Test scripts are written based solely on the **requirement document**, not the solution repo. This prevents “baking in” the implementation into the tests and avoids inheriting potential errors from the human solution. We employ two test sets: 1) *Fail* → *Pass* (F2P): These tests evaluate the completion status of the requirements. 2) *Pass* → *Pass* (P2P): Often ignored in other benchmarks, these tests verify that the agent’s modifications have not broken existing system functionality.

To accurately reflect task completion, we prioritize environment-state verification, such as confirming a service deployment by checking if the port is listening and the log content, rather than just validating that a start command was written. Furthermore, both test sets contribute to a **step-level score**. For CS assignments with built-in grading scripts, we parse the output to calculate this score; for tasks without predefined tests, we manually segment the requirements into sub-tasks and write corresponding evaluation scripts.

3.1.5 Verification and Quality Control

Verification condition. A task is valid if the F2P tests fail on the initial repo and pass on the solution repo, while the P2P tests pass on both.

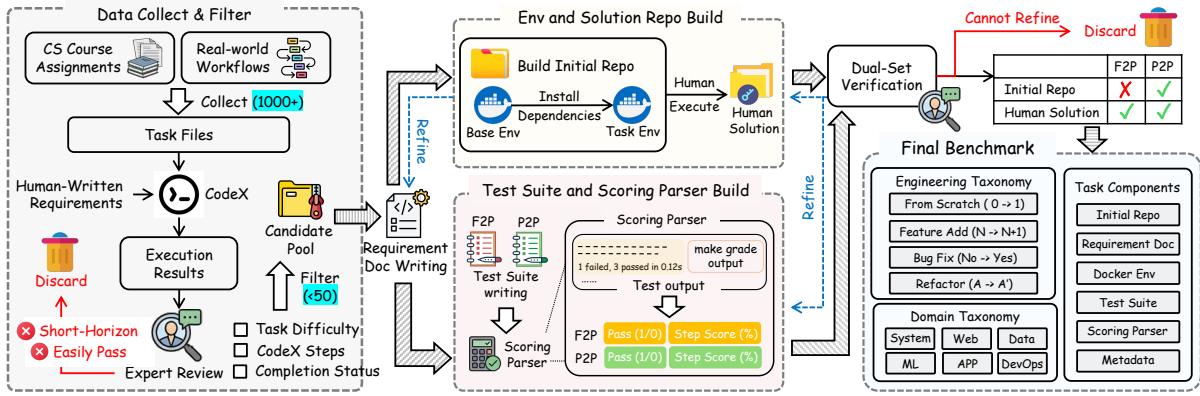


Figure 3: The LongCLI-Bench construction pipeline. We curate tasks from diverse sources and employ a parallel construction method for solutions and tests. The pipeline features a strict Dual-Set Verification mechanism with iterative refinement loops to ensure high-quality, contamination-free benchmarks across various engineering and domain categories.

To ensure benchmark quality, we adopt a rigorous iterative, closed-loop verification procedure: 1) During the environment setup and solution creation, we immediately revise the documentation if any ambiguities or missing content are identified. 2) Test Suite Validation. Execute test scripts against both the initial and solution repos. If the verification condition is not met, it indicates an issue with the test script, the solution, or the environment. Human experts review and fix the issue, triggering a re-validation. If a task cannot meet the verification condition after three iterations, it is discarded. 3) Expert review. Each task undergoes an final expert review to ensure logical correctness and feasibility before inclusion.

3.2 Task Composition and Evaluation

3.2.1 Task Composition

LongCLI-Bench is designed around five principles: long-horizon, isolated environments, clear requirements, solvability, and contamination control. Each task consists of the following components:

Initial Repo: The task starting codebase (e.g., a skeleton or empty directory).

Task Requirement: A task requirement document that explicitly defines functional goals and entry points to ensure clear requirements.

Environment: An isolated environment ensuring consistent execution and reproducibility.

Solution Repo: A human-authored “golden solution” ensuring solvability and validation.

Tests: A dual-set testing suite verifying both requirement fulfillment and regression testing.

Scoring Parser: A parser that analyzes test outputs to calculate a step-level metric.

Metadata: Includes task type, domain, difficulty level, and estimated human completion time.

3.2.2 Evaluation Workflow

Based on above components, the evaluation follows the sequence:

1) **Initialization:** An isolated docker environment is initialized with the initial repo.

2) **Execution:** The agent receives the requirement, plans and interacts with the terminal to execute.

3) **Test:** Upon task completion or timeout, the harness executes the test scripts.

4) **Scoring:** The Scoring Parser aggregates the test results to compute the pass rate and step scores.

LongCLI-Bench also supports optional evaluation approaches:

1) **Multiple Attempts:** The agent executes the task repeatedly for N independent attempts.

2) **Self-Correction:** The agent leverages test feedback from the prior turn and re-executes to refine the solution in a multi-turn process.

3.3 Taxonomy

3.3.1 Engineering Taxonomy

To systematically evaluate agent capabilities across the software development lifecycle, we classify tasks into four distinct engineering categories. This classification enables targeted assessment of agent performance across a full spectrum of core engineering competencies.

From Scratch ($0 \rightarrow 1$): The ability to plan, configure, and build a runnable project from scratch.

Feature Addition ($N \rightarrow N + 1$): The ability to add new modules to an existing codebase.

Bug Fix ($No \rightarrow Yes$): The ability to diagnose, locate, and fix complex bugs.

Refactor ($A \rightarrow A'$): The ability to optimize or restructure code.

3.3.2 Domain Taxonomy

Beyond engineering types, LongCLI-Bench covers six primary domains to assess agent capabilities across diverse technical fields.

System Programming: Operating system, compilers, memory, concurrency, embedded programming, hardware interfacing, and distributed system.

Web Development: Frontend/backend, database, API design, authentication, and service.

Data Engineering: Crawling, filtering, formatting and statistical analysis/visualization.

Machine Learning: training, inference, deployment, evaluation, and signal processing.

Applications: Business or gameplay, simulation engines, physics integration, and interactive tool.

DevOps: CI/CD pipelines, containerization, environment build, system monitoring/diagnostics.

3.4 Statistics of LongCLI-Bench

Table 1: Comparison of LongCLI-Bench and Terminal-Bench@2. Files, LoC and Expert Time are average results; LoC stands for Lines of Code.

Benchmark	Tasks	Files	LoC	Expert Time (min)
Terminal-Bench@2	89	0.69	227.7	206.7
LongCLI-Bench	20	104.0	15,000+	1000+

We present a detailed statistical analysis of the 20 curated tasks in LongCLI-Bench.

LongCLI-Bench averages 15,000+ lines of code (LoC) and 104 source files per task, covering a diverse linguistic landscape, including C, Python, Java, and JavaScript, representing complex, repository-level engineering. In stark contrast, Terminal-Bench@2 averages only 227.7 LoC and 0.69 files. This disparity highlights that LongCLI-Bench evaluates the ability to handle long-horizon complex problems within massive, interdependent systems rather than executing isolated snippets.

This complexity extends to the temporal dimension, imposing a significantly higher cognitive load. Expert completion time for LongCLI-Bench averages 1000+ minutes (vs. 206.7 min for Terminal-Bench@2). These metrics confirm that the benchmark effectively probes long-horizon planning and context maintenance capabilities, differentiating it from short-term execution tasks.

4 Experiments

4.1 Experiment Setup

Agents and Models. We evaluate two distinct categories of agents. The first comprises commercial CLI assistants driven by proprietary models: Codex (gpt-5.x-codex series) and Claude Code (Claude-Sonnet/Opus-4.x series). The second leverages the OpenHands(Wang et al., 2024b) framework to assess the leading open-source models, including DeepSeek-V3.1(Liu et al., 2024), GLM-4.6(Zhipu AI, 2025), and Qwen3-235B-A22B(Yang et al., 2025a).

Evaluation Metrics and Protocol. To ensure fairness, we employed a unified system prompt and averaged results over **three** independent runs. We first measure **step scores** for both requirement fulfillment (F2P) and regression testing (P2P), assessing the percentage of sub-steps completed. A binary “pass” is defined strictly: a test set is considered passed only when its step score reaches 100%. Accordingly, we derive the **F2P pass rate** and **P2P pass rate**. The overall **Pass Rate** is defined as the percentage of tasks where the agent pass on both F2P and P2P tests. Additionally, we report **Pass@3** to assess stability, along with execution time for efficiency analysis.

4.2 Main Results

Table 2 presents the overall performance of various models on LongCLI-Bench in a single-turn setting. The experimental results reveal that LongCLI-Bench poses an extreme challenge to current state-of-the-art agent systems. The pass rate for most models falls below 15%, with even the top-performing Claude-Opus-4.6 achieving only 16.7%.

This result compellingly demonstrates that while existing agent technologies excel in short code generation, a significant capability gap remains when facing real-world long-horizon tasks requiring long-term memory, environmental perception, and complex logical planning. Commercial systems significantly outperformed the open-source framework in requirements completion (F2P step score), indicating that general frameworks struggle to adapt to such complex engineering tasks without targeted optimizations.

While the average P2P step scores are very high (>98%) across all models, the P2P pass rate is notably lower, ranging from 70.0% to 88.3%. The low P2P pass rate indicates that in the tasks where

Table 2: Overall Performance on LongCLI-Bench (Average on 3 attempts). We report the Pass Rate(%), Pass@3(%), and average Step Scores(%) for both F2P and P2P tests. Time is measured in minutes.

Agent	Model	Pass	Pass@3	F2P Pass	F2P Step Score	P2P Pass	P2P Step Score	Time(min)
Codex	GPT-5.1-Codex-Max	11.7	15.0	16.7	41.8	70.0	99.2	12.9
	GPT-5.2-Codex	10.0	15.0	13.3	39.1	73.3	99.4	10.9
	GPT-5.3-Codex	15.0	20.0	18.3	44.1	86.7	99.5	8.8
Claude Code	Claude-Sonnet-4.5	10.0	10.0	13.3	42.0	70.0	98.4	10.6
	Claude-Opus-4.5	16.7	20.0	20.0	47.1	71.7	98.7	11.3
	Claude-Opus-4.6	16.7	25.0	20.0	50.7	78.3	99.1	17.6
OpenHands	DeepSeek-V3.1	5.0	10.0	11.7	25.3	88.3	99.7	17.9
	GLM-4.6	6.7	10.0	11.7	26.8	83.3	99.5	25.3
	Qwen3-235B-A22B	6.7	10.0	13.3	28.9	81.7	99.3	30.2

agents did attempt complex edits, they introduced broken modification in nearly 12%-30% of cases. This corroborates that agents often lose sight of the broader context when focusing on new features, or struggle to adhere to the instructions as task difficulty improves, resulting in unintended breakage of existing functionalities.

Interestingly, OpenHands with DeepSeek-V3.1 achieved the highest P2P pass rate (88.3%), likely because its lower execution capability (F2P Score 25.3) limited the scope of its modifications, reducing the surface area for potential errors.

4.3 Step-level Analysis

To investigate when models fail to complete tasks, we analyzed the distribution of F2P step score.

Table 3 reveals that failures are not evenly distributed and the majority of tasks fail at early stages. Across all evaluated agents, the highest percentage of outcomes is concentrated in the <30% range. This implies that agents encountered severe difficulties at the very onset of the tasks. Notably, commercial agent-model pairs exhibit a lower proportion of scores falling into low range compared to OpenHands, an indicator of their superior early-stage planning capabilities. This strength helps them navigate early task obstacles through more robust upfront preparation and reasoning.

The few cases in the [80, 100) range highlights the sequential dependency of long-horizon tasks, where a failure in a preceding step directly blocks the subsequent steps. This distribution highlights step-level scoring’s core value: unlike binary metrics that treat all failures equally, step scores reveal the exact breakage point in the execution chain. This granularity allows us to discern partial progress and diagnose whether agents fail at fundamental planning or specific logic implementation, confirming that LongCLI-Bench effectively

Table 3: Distribution of F2P step scores. Columns show the percentage of tasks falling into specific score ranges.

Agent	Model	[0,30)	[30,60)	[60,80)	[80,100)	[100]
Codex	GPT-5.1-Codex-Max	45.0	16.7	13.3	8.3	16.7
	GPT-5.2-Codex	51.7	15.0	10.0	10.0	13.3
	GPT-5.3-Codex	41.7	13.3	15.0	11.7	18.3
Claude Code	Claude-Sonnet-4.5	46.7	13.3	11.7	15.0	13.3
	Claude-Opus-4.5	40.0	23.3	6.7	10.0	20.0
	Claude-Opus-4.6	38.3	20.0	10.0	11.7	20.0
OpenHands	DeepSeek-V3.1	65.0	13.3	5.0	5.0	11.7
	GLM-4.6	63.3	15.0	3.3	6.7	11.7
	Qwen3-235B-A22B	58.3	13.3	8.3	6.7	13.3

evaluates the robustness of the entire workflow.

4.4 Self-Correction Capabilities

Under the Self-Correction setting, agents re-executes the task using feedback from previous turn enabling multi-turn self-correction. As illustrated in Figure 4, this setting yields pass rate improvements, with substantial gains from T1 to T2, and continued but comparatively smaller improvements from T2 to T3 for several models (e.g., Claude-Opus-4.6).

F2P keeps improving across rounds, suggesting that agent-model pairs can better leverage the feedback provided by Self-Correction setting to satisfy previously failing requirements and complete the task more effectively.

Meanwhile, P2P stays high overall, but different agent-model pairs diverge at different rounds. Codex with GPT-5.3-Codex pushes P2P pass rate to 95% at T2 and 100% at T3. In contrast, GPT-5.1-Codex-Max reaches 85% at T2 but declines to 80% at T3, and similarly, Claude-Opus-4.6 reaches 95% at T2 but drops to 90% at T3. This suggests that the last self-correction round can widen the change scope. It may complete the remaining hard cases, but it can also introduce new regression testing risks.




Figure 4: Multi-Turn Self-Correction Performance.

Table 4: Results of Human-Agent Collaboration. Pass and F2P Score are reported as percentages (%). Time is measured in minutes. Inter.Avg represents the average number of human interventions per task.

Agent & Model	Mode	Pass	F2P Score	Time	Inter.Avg
Codex GPT-5.3-Codex	Base (Avg)	15.0	44.1	8.8	-
	Self-Correction	40.0	58.4	26.7	-
	Plan	41.7	59.3	10.7	-
	Interactive	45.0	61.3	25.9	2.7
	Plan & Interactive	50.0	62.3	22.4	2.2
Claude Code Claude-Opus-4.6	Base (Avg)	16.7	50.7	17.6	-
	Self-Correction	55.0	63.6	43.8	-
	Plan	58.3	65.4	15.4	-
	Interactive	58.3	67.4	39.6	2.4
	Plan & Interactive	61.7	69.3	35.4	2.1

In terms of time overhead, multi-round self-correction shows mixed effects on runtime across different agents and models, with no consistent trend of increasing time cost in later rounds. The gains are real, but the marginal benefit tends to shrink in later rounds.

4.5 Human-Agent Collaboration

To investigate the performance of human-agent collaboration, we designed two experimental setups: **Static Plan Injection**. Before the execution phase, we inject key plans without specific code details. This setup evaluates whether reducing planning errors leads to improved task accuracy, thereby quantifying the impact of Planning Capability.

Dynamic Interactive Guidance. In this setting, the model autonomously decides whether to request human intervention based on its current state. Upon intervention, humans provide the next steps and future roadmap, offering explanatory guidance rather than direct implementation. If the model explicitly asks for code, the request is denied to force the model to reason independently. The maximum number of interventions is limited to 3.

Table 4 summarizes performance across two agent-model pairs (Codex with GPT-5.3-Codex; Claude Code with Claude-Opus-4.6), showing that current agents face dual limitations in both plan-

ning and execution proficiency, and human collaboration significantly mitigates these bottlenecks.

Plan injection outperformed the self-correction baseline in both pass rate and efficiency. For instance, Claude Code with plan injection achieved a 58.3% pass rate compared to 55.0% with self-correction. This confirms that establishing a correct plan upfront is far more efficient than relying on the agent to iteratively fix errors during the execution process, as it prevents the agent from entering incorrect logical branches initially.

Interactive guidance generally achieved higher performance compared to static plan injection alone. Claude Code improved from 58.3% (Plan) to 58.3% (Interactive) in pass rate, while F2P Score increased from 65.4 to 67.4. While static plans provide a roadmap, they cannot anticipate all runtime anomalies. This dynamic intervention proves more robust, as it effectively guides the agent away from unproductive execution trajectories and logic errors that a static plan cannot anticipate.

The combined Plan & Interactive setting yielded the best results across all metrics. Crucially, the presence of a pre-injected plan reduced the need for human intervention (e.g., Inter.Avg dropped from 2.4 to 2.1 for Claude Code), indicating a synergy where the plan handles the roadmap while human interaction resolves specific execution problems.

Our findings indicate that rather than exclusively chasing full autonomy, future work should focus on developing collaborative systems that leverage the synergy between efficient execution and human strategic guidance.

4.6 Error Analysis

To better understand why current agents achieve low end-to-end success on LongCLI-Bench (most pass rates <20%), we manually inspected 50 failed trajectories across representative agent-model pairs (primarily Codex and Claude Code). Overall, **failures are rarely caused by local syntax-level coding mistakes**. Instead, they are dominated by long-

horizon workflow breakdowns, where the agent must plan, verify state, and preserve consistency across many interdependent steps. This aligns with our step-level findings that the majority of runs remain below 30% completion.

Overall, LongCLI-Bench failures are primarily caused by the following three reasons:

Repetitive loops from weak strategic adaptation.

A common pattern is that the agent encounters an execution failure, proposes a superficial patch, re-runs the same command, observes the same error, and repeats until the step limit is exhausted. Such trajectories reveal that the agent fails to recognize that the current plan is invalid and does not shift focus to addressing the root cause.

Environment grounding and verification gaps.

We observed instances of misdiagnosis in which environment-related issues were incorrectly attributed to code logic flaws, resulting in code edits that failed to address the actual cause of the failure.

Long-term inconsistency and regression. Even when agents make substantial progress on the new requirements, they often break existing functionality: while the average P2P step scores are very high, the P2P pass rates are much lower, indicating that agents introduce regressions risks when they make non-trivial edits. We also saw context drift in long runs (e.g., forgetting earlier constraints), which explains why later self-correction rounds can widen change scope and increase regression risks.

5 Conclusion

LongCLI-Bench introduces a new benchmark for evaluating agentic programming in long-horizon command-line interface (CLI) tasks. By curating 20 complex tasks from over 1,000 real-world workflows and computer science assignments, LongCLI-Bench addresses key gaps in existing benchmarks, such as their focus on short, isolated tasks. Unlike many benchmarks that offer binary pass/fail evaluations, LongCLI-Bench emphasizes long-duration tasks that reflect real-world challenges and evaluates both requirement fulfillment and regression avoidance through step-level scores. The results show that current agent systems struggle with long-horizon tasks, achieving low pass rates. However, human-agent collaboration significantly boosts performance, highlighting the need to prioritize collaborative workflows in future research alongside autonomous agent advancements.

Limitations

Task Creation Requires Significant Manual Effort: Creating tasks for LongCLI-Bench involves extensive manual work, including writing requirement documents, constructing solution paths, setting up test environments, and creating test scripts. This process is time-consuming, with each task taking an average of 40 hours to complete, resulting in a relatively small dataset. Additionally, this makes it challenging to curate high-quality, long-horizon tasks that effectively challenge agents.

Evaluation Metrics: While step-level scores provide more granular insights, they do not fully assess agent performance in areas like code quality or efficiency. Future versions of the benchmark could include additional metrics for a more complete evaluation.

Ethics Statement

Data Privacy and Usage: All tasks in LongCLI-Bench are derived from publicly available resources or manually curated data. No sensitive or proprietary information was used, and all data have been anonymized to ensure privacy.

Human-Agent Collaboration: This study involves human-agent collaboration to explore how agent performance improves with human guidance. We strictly adhere to ethical guidelines to ensure participant rights and safeguard their privacy.

Evaluation Ethics: All data used for evaluation are standardized and anonymized, ensuring compliance with ethical research practices.

Transparency: We will release detailed task creation processes and evaluation results to ensure transparency and encourage open participation in AI research.

References

- Anthropic. 2025. Claude code. <https://github.com/anthropics/claude-code>. GitHub repository.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, and Quoc Le. 2021. *Program synthesis with large language models*. *arXiv preprint arXiv:2108.07732*.
- Ramakrishna Bairi, Atharv Sonwane, Aditya Kanade, Vageesh D C, Arun Iyer, Suresh Parthasarathy, Sriram Rajamani, B. Ashok, and Shashank Shet. 2023. *Codeplan: Repository-level coding using llms and planning*. *arXiv preprint arXiv:2309.12499*.

- Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, and Molly Q Feldman. 2022. **Multipl-e: A scalable and extensible approach to benchmarking neural code generation.** *arXiv preprint arXiv:2208.08227*.
- Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, and Tejal Patwardhan. 2024. **Mle-bench: Evaluating machine learning agents on machine learning engineering.** *arXiv preprint arXiv:2410.07095*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, and Greg Brockman. 2021. **Evaluating large language models trained on code.** *arXiv preprint arXiv:2107.03374*.
- Chunyuan Deng, Yilun Zhao, Xiangru Tang, Mark Gerstein, and Arman Cohan. 2024. **Investigating data contamination in modern benchmarks for large language models.** In *Proceedings of NAACL 2024 (Long Papers)*, pages 8706–8719. Association for Computational Linguistics.
- Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Nitin Pasari, and Chetan Rane. 2025. **Swe-bench pro: Can ai agents solve long-horizon software engineering tasks?** *arXiv preprint arXiv:2509.16941*.
- Jingzhe Ding, Shengda Long, Changxin Pu, Huan Zhou, Hongwan Gao, Xiang Gao, Chao He, Yue Hou, Fei Hu, Zhaojian Li, Weiran Shi, Zaiyuan Wang, Daoguang Zan, Chenchen Zhang, Xiaoxu Zhang, Qizhi Chen, Xianfu Cheng, Bo Deng, Qingshui Gu, and 29 others. 2025a. **NI2repo-bench: Towards long-horizon repository generation evaluation of coding agents.** *Preprint, arXiv:2512.12730*.
- Yangruibo Ding, Zijian Wang, Wasi Uddin Ahmad, Hantian Ding, Ming Tan, Nihal Jain, Murali Krishna Ramanathan, Ramesh Nallapati, Parminder Bhatia, and Dan Roth. 2023. **Crosscodeeval: A diverse and multilingual benchmark for cross-file code completion.** *arXiv preprint arXiv:2310.11248*.
- Yuzhe Ding, Kang He, Bobo Li, Li Zheng, Haijun He, Fei Li, Chong Teng, and Donghong Ji. 2025b. **Zero-shot conversational stance detection: Dataset and approaches.** In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 3221–3235.
- Xueying Du, Mingwei Liu, Kaixin Wang, Hanlin Wang, Junwei Liu, Yixuan Chen, Jiayi Feng, Chaofeng Sha, Xin Peng, and Yiling Lou. 2024. **Evaluating large language models in class-level code generation.** In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering (ICSE 2024)*, pages 81:1–81:13. ACM.
- Kelin Fu, Tianyu Liu, Zeyu Shang, Yingwei Ma, Jian Yang, Jiaheng Liu, and Kaigui Bian. 2025. **Multi-docker-eval: A ‘shovel of the gold rush’ benchmark on automatic environment building for software engineering.** *arXiv preprint arXiv:2512.06915*.
- Google. 2025. **Gemini cli.** <https://github.com/google-gemini/gemini-cli>. GitHub repository.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, and Y. K. Li. 2024. **Deepseek-coder: When the large language model meets programming – the rise of code intelligence.** *arXiv preprint arXiv:2401.14196*.
- Kang He, Boyu Chen, Yuzhe Ding, Fei Li, Chong Teng, and Donghong Ji. 2026. **Pase: Prototype-aligned calibration and shapley-based equilibrium for multimodal sentiment analysis.** In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 30960–30968.
- Kang He, Yuzhe Ding, Bobo Li, Haining Wang, Fei Li, Chong Teng, and Donghong Ji. 2025a. **Harnessing dimensional contrast and information compensation for sentence embedding enhancement.** In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Kang He, Yuzhe Ding, Haining Wang, Fei Li, Chong Teng, and Donghong Ji. 2025b. **DALR: Dual-level alignment learning for multimodal sentence representation learning.** In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 3586–3601.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, and Dawn Song. 2021. **Measuring coding challenge competence with apps.** *arXiv preprint arXiv:2105.09938*.
- Dong Huang, Jie M. Zhang, Michael Luck, Qingwen Bu, Yuhao Qing, and Heming Cui. 2023. **Agent-coder: Multi-agent-based code generation with iterative testing and optimisation.** *arXiv preprint arXiv:2312.13010*.
- Siming Huang, Tianhao Cheng, J. K. Liu, Jiaran Hao, Liuyihan Song, Yang Xu, J. Yang, Jiaheng Liu, Chenchen Zhang, and Linzheng Chai. 2024. **Open-coder: The open cookbook for top-tier code large language models.** *arXiv preprint arXiv:2411.04905*.
- Laude Institute. 2025. **Terminal-bench.** <https://github.com/laude-institute/terminal-bench>. GitHub repository.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. **Live-codebench: Holistic and contamination free evaluation of large language models for code.** *arXiv preprint arXiv:2403.07974*.

- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2024. [Swe-bench: Can language models resolve real-world github issues?](#) *arXiv preprint arXiv:2310.06770*.
- Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Scott Wen tau Yih, Daniel Fried, Sida Wang, and Tao Yu. 2022. [Ds-1000: A natural and reliable benchmark for data science code generation.](#) *arXiv preprint arXiv:2211.11501*.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, and Jenny Chim. 2023. [Starcoder: may the source be with you!](#) *arXiv preprint arXiv:2305.06161*.
- Wei Li, Xin Zhang, Zhongxin Guo, Shaoguang Mao, Wen Luo, Guangyue Peng, Yangyu Huang, Houfeng Wang, and Scarlett Li. 2025. [Fea-bench: A benchmark for evaluating repository-level code generation for feature implementation.](#) *arXiv preprint arXiv:2503.06680*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024. [Deepseek-v3 technical report.](#) *arXiv preprint arXiv:2412.19437*.
- Tianyang Liu, Canwen Xu, and Julian McAuley. 2023. [Repubench: Benchmarking repository-level code auto-completion systems.](#) *arXiv preprint arXiv:2306.03091*.
- Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, and Duyu Tang. 2021. [Codexglue: A machine learning benchmark dataset for code understanding and generation.](#) *arXiv preprint arXiv:2102.04664*.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, and Yiming Yang. 2023. [Self-refine: Iterative refinement with self-feedback.](#) *arXiv preprint arXiv:2303.17651*.
- Niels Mündler, Mark Niklas Müller, Jingxuan He, and Martin Vechev. 2024. [Swb-bench: Testing and validating real-world bug-fixes with code agents.](#) *arXiv preprint arXiv:2406.12952*.
- Ziyi Ni, Huacan Wang, Shuo Zhang, Shuo Lu, Ziyang He, Wang You, Zhenheng Tang, Yuntao Du, Bill Sun, and Hongzhang Liu. 2025. [Gittaskbench: A benchmark for code agents solving real-world tasks through code repository leveraging.](#) *arXiv preprint arXiv:2508.18993*.
- OpenAI. 2024. [Introducing swe-bench verified.](#) <https://openai.com/index/introducing-swe-bench-verified/>. Blog post.
- OpenAI. 2025. [Codex.](#) <https://github.com/openai/codex>. GitHub repository.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, and Tal Remez. 2023. [Code llama: Open foundation models for code.](#) *arXiv preprint arXiv:2308.12950*.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools.](#) *arXiv preprint arXiv:2302.04761*.
- ByteDance Seed, Yuyu Zhang, Jing Su, Yifan Sun, Chenguang Xi, Xia Xiao, Shen Zheng, Anxiang Zhang, Kaibo Liu, and Daoguang Zan. 2025. [Seed-coder: Let the code model curate data for itself.](#) *arXiv preprint arXiv:2506.03524*.
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning.](#) *arXiv preprint arXiv:2303.11366*.
- Haining Wang, Kang He, Bobo Li, Lei Chen, Fei Li, Xu Han, Chong Teng, and Donghong Ji. 2024a. [Refining and synthesis: A simple yet effective data augmentation framework for cross-domain aspect-based sentiment analysis.](#) In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10318–10329.
- Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, and Jaskirat Singh. 2024b. [Openhands: An open platform for ai software developers as generalist agents.](#) *arXiv preprint arXiv:2407.16741*.
- Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. 2024. [Agentless: Demystifying llm-based software engineering agents.](#) *arXiv preprint arXiv:2407.01489*.
- Yijia Xiao, Runhui Wang, Luyang Kong, Davor Golac, and Wei Wang. 2025. [Csr-bench: Benchmarking llm agents in deployment of computer science research repositories.](#) *arXiv preprint arXiv:2502.06111*.
- Cheng Xu, Shuhao Guan, Derek Greene, and M-Tahar Kechadi. 2024. [Benchmark data contamination of large language models: A survey.](#) *arXiv preprint arXiv:2406.04244*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, and Chenxu Lv. 2025a. [Qwen3 technical report.](#) *arXiv preprint arXiv:2505.09388*.
- John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024a. [Swe-agent: Agent-computer interfaces enable automated software engineering.](#) *arXiv preprint arXiv:2405.15793*.

- John Yang, Carlos E. Jimenez, Alex L. Zhang, Kilian Lieret, Joyce Yang, Xindi Wu, Ori Press, Niklas Muennighoff, Gabriel Synnaeve, and Karthik R. Narasimhan. 2024b. [Swe-bench multimodal: Do ai systems generalize to visual software domains?](#) *arXiv preprint arXiv:2410.03859*.
- John Yang, Kilian Lieret, Carlos E. Jimenez, Alexander Wettig, Kabir Khandpur, Yanzhe Zhang, Binyuan Hui, Ofir Press, Ludwig Schmidt, and Diyi Yang. 2025b. [Swe-smith: Scaling data for software engineering agents](#). *arXiv preprint arXiv:2504.21798*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. [React: Synergizing reasoning and acting in language models](#). *arXiv preprint arXiv:2210.03629*.
- Boxi Yu, Yuxuan Zhu, Pinjia He, and Daniel Kang. 2025. [Utboost: Rigorous evaluation of coding agents on swe-bench](#). *arXiv preprint arXiv:2506.09289*.
- Daoguang Zan, Zhirong Huang, Wei Liu, Hanwu Chen, Linhao Zhang, Shulin Xin, Lu Chen, Qi Liu, Xiaojian Zhong, and Aoyan Li. 2025. [Multi-swe-bench: A multilingual benchmark for issue resolving](#). *arXiv preprint arXiv:2504.02605*.
- Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang Lou, and Weizhu Chen. 2023. [Repcoder: Repository-level code completion through iterative retrieval and generation](#). *arXiv preprint arXiv:2303.12570*.
- Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. 2024. [Autocoderover: Autonomous program improvement](#). *arXiv preprint arXiv:2404.05427*.
- Zhipu AI. 2025. [Glm-4.6: Advanced agentic, reasoning and coding capabilities](#). Accessed: 2026-01-06.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, and Indraneil Paul. 2024. [Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions](#). *arXiv preprint arXiv:2406.15877*.