

KLONG: TRAINING LLM AGENT FOR EXTREMELY LONG-HORIZON TASKS

Yue Liu^{1*} Yingwei Ma^{2*} Yibo Miao³
 Yanhao Li⁴ Yuchong Xie² Xinlong Yang⁴ Zhiyuan Hu^{1,5}
 Flood Sung Jiaheng Zhang¹ Bryan Hooi¹ Tianyu Liu⁴
¹NUS ²HKUST ³SJTU ⁴PKU ⁵MIT
 yliu@u.nus.edu *denotes the equal contribution.

ABSTRACT

This paper introduces KLong, an open-source LLM agent trained to solve extremely long-horizon tasks¹. The principle is to first cold-start the model via trajectory-splitting SFT, then scale it via progressive RL training. Specifically, we first activate basic agentic abilities of a base model with a comprehensive SFT recipe. Then, we introduce Research-Factory, an automated pipeline that generates high-quality training data by collecting research papers and constructing evaluation rubrics. Using this pipeline, we build thousands of long-horizon trajectories distilled from the frontier model. To train with these extremely long trajectories, we propose trajectory-splitting SFT, which preserves early context, progressively truncates later context, and maintains overlap between sub-trajectories. In addition, to further improve long-horizon task-solving capability, we propose progressive RL, which schedules training into multiple stages with progressively extended timeouts. Extensive experiments demonstrate the superiority and generalization of KLong, as shown in Figure 2. Notably, our KLong (106B) surpasses Kimi K2 Thinking (1T) by 11.28% on PaperBench, and the performance improvement generalizes to other benchmarks, e.g. SWE-bench Verified and MLE-bench.

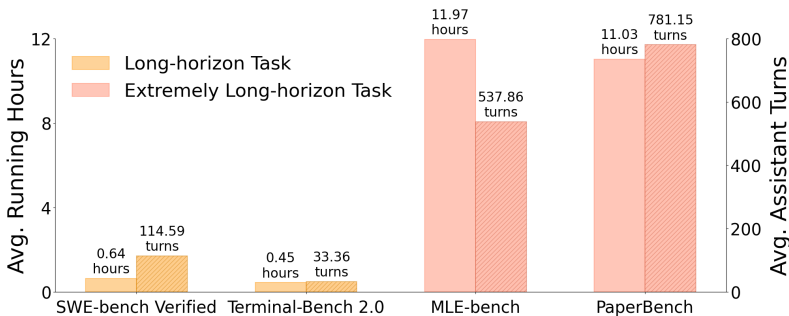


Figure 1: **Long-horizon Task vs. Extremely Long-horizon Task.** We show the challenges of the extremely long-horizon tasks by comparing 2 extremely long-horizon tasks, MLE-bench (Chan et al., 2024) & PaperBench (Starace et al., 2025), with 2 long-horizon tasks, SWE-bench Verified (OpenAI, 2024c) & Terminal-Bench 2.0 (Team, 2025), in terms of running time and assistant turns.

1 INTRODUCTION

Large language models (LLMs) (Deepmind, 2025b; OpenAI, 2025c; Anthropic, 2025a) become the cornerstone of intelligence, driving advancements in various fields such as ChatBots (OpenAI,

¹In this paper, extremely long-horizon tasks refer to those that 1) would exceed the context window without context management and 2) involve long-running experiments, e.g., MLE-bench (Chan et al., 2024) and PaperBench (Starace et al., 2025).

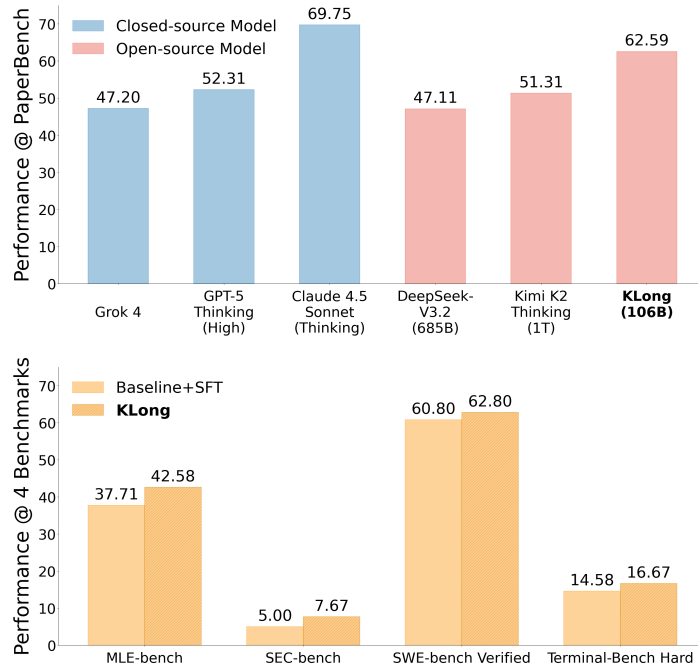


Figure 2: **Performance of KLong on five Agentic Benchmarks.** The training is tailored to PaperBench, and the long-horizon task-solving abilities generalize to the other four benchmarks.

2025a). Equipped with the abilities of memory, tool use, and planning, LLM agents exhibit great potential in complex interactive real-world applications like vibe coding (Anthropic, 2025b), computer use (OpenAI, 2025b), and deep research (Deepmind, 2025a).

Long-horizon tasks, such as bug fixing (Jimenez et al., 2023; OpenAI, 2024c), terminal coding (Team, 2025), and deep research (Phan et al., 2025), present significant challenges for current LLM agents. To address these challenges, training-free, system-level solutions, like context management (Anthropic, 2025c), have been proposed. To improve the inherent long-horizon task-solving ability, researchers are advancing training techniques and infrastructure (Sun et al., 2025; Xi et al., 2025).

However, previous papers either use models with a limited context window (e.g., 32K) or focus on tasks that only require a few assistant turns (e.g., ranging from 20 to 200). It makes them fail to solve extremely long-horizon tasks in real-world scenarios, e.g., machine learning engineering (Chan et al., 2024) or reproducing research papers (Starace et al., 2025). As shown in Figure 1, we observe that the running time and assistant turns of the extremely long-horizon task are approximately 10 times that of the long-horizon task. The sheer scale of the running time and the interactions poses significant challenges to existing LLM agents. Therefore, this paper aims to train an LLM agent for extremely long-horizon tasks. Given that reproducing research papers is a pinnacle of complexity and comprehensiveness, requiring the model’s ability in machine learning engineering, software engineering, and extensive terminal interactions, etc., we primarily tailor our training for reproducing research papers (Starace et al., 2025) and aim to generalize its ability to other tasks such as bug fixing (OpenAI, 2024c), machine learning engineering (Zheng et al., 2023), terminal coding (Team, 2025), and code security Lee et al. (2025).

To this end, we train KLong, an open-source LLM agent, to solve extremely long-horizon tasks. Since extremely long-horizon tasks inevitably exceed the context window, this poses challenges to standard supervised fine-tuning (SFT) and reinforcement learning (RL). To solve this problem, our principle is to first cold-start the model via trajectory-splitting SFT, then scale it via progressive RL.

Specifically, we first activate the basic agentic ability of the base model via a comprehensive SFT training recipe that covers common knowledge, coding, mathematics, and search. Then, we develop Research-Factory, an automated pipeline that systematically scales training data for the research reproducing task, as shown in Figure 3. It contains a search agent and an evaluation agent that auto-

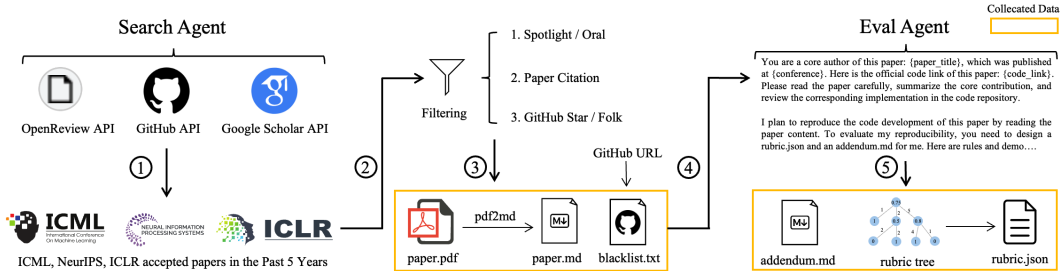


Figure 3: **Research-Factory: An Automatic Pipeline of Scaling Training Data for Research Reproducing Task.** First, the search agent collects basic data of accepted papers from ICML, NeurIPS, and ICLR conferences. Then, the filter selects the data based on the quality and impact of the papers. The PDF is converted to Markdown. The official GitHub URL is added to the blacklist.txt file to avoid cheating. Last, the evaluation agent designs the addendum and the rubric tree by analyzing the paper and the official code implementation.

matically collect research papers and construct rubrics for evaluating research reproduction. Using Research-Factory, we construct a high-quality prompt set with thousands of samples. Subsequently, we built thousands of extremely long-horizon trajectories by distilling Claude 4.5 Sonnet (Thinking). To train with these extremely long trajectories, as shown in Figure 4, we propose trajectory-splitting SFT via 1) pinning the paper-reading part at the beginning of the context, 2) progressively truncating the context according to the context window, and 3) overlapping sub-trajectories for contextual continuity. It significantly improves the assistant turns, $114.90 \rightarrow 732.70$. In addition, we propose progressive RL, which schedules training into multiple stages with progressively extended timeouts. It further enhances long-horizon task-solving capability, resulting in a 6.67% performance improvement. Extensive experiments demonstrate the superiority and generalization of our method. The main contributions of this paper are summarized as follows.

- We introduce KLong, a well-trained open-source LLM agent for extremely long-horizon tasks.
- We build Research-Factory, an automatic pipeline for scaling training data of research task.
- We enhance the long-horizon task-solving ability via trajectory-splitting SFT and progressive RL.

2 METHODOLOGY

This section introduces KLong, a well-trained LLM agent for extremely long-horizon tasks. We first present the task definition and then introduce the training details of KLong.

2.1 TASK DEFINITION

Given an extremely long-horizon task \mathcal{T} , an LLM agent \mathcal{A} is required to solve it by analyzing the task specification and interacting with an environment over a potentially large number of assistant turns, optionally under time constraints.

Take the replicating research task (Starace et al., 2025) as an example, given a research paper \mathcal{P} , the LLM agent \mathcal{A} is tasked to read the paper, summarize the core contributions, and replicate the corresponding code implementations $\hat{\mathcal{C}}$, within a limited time t . Within the sandbox environment, agent \mathcal{A} iteratively plans, executes tool calls, and interacts with the environment, yielding a sequence of N actions denoted by $\{a_k\}_{k=1}^N$. Besides, to prevent cheating, the official code repository is added to the blacklist and banned. The replicating process can be formulated as follows.

$$\mathcal{A}: \mathcal{P} \xrightarrow{\{a_k\}_{k=1}^N} \hat{\mathcal{C}}, \text{ s.t. } \sum_{k=1}^N t_k \leq t. \tag{1}$$

Then, to evaluate the quality of the replicating, a frontier judge model \mathcal{J} is used judge the replicated code based on the rubric tree \mathcal{K} , which is written by the original authors of the paper, as follows.

$$Q = \mathcal{J}(\hat{\mathcal{C}}, \mathcal{K}), \tag{2}$$

where Q denotes the score of the replicating research task.

Among lots of extremely long-horizon tasks, reproducing research papers represents a pinnacle of complexity and comprehensiveness, demanding expertise in machine learning engineering, software engineering, and extensive terminal interactions. Therefore, we mainly train KLong to solve the reproducing research task and aim to generalize its ability to other tasks such as bug fixing (OpenAI, 2024c), machine learning engineering (Zheng et al., 2023), terminal coding (Team, 2025), and code security Lee et al. (2025).

2.2 RESEARCH-FACTORY

To efficiently train KLong on extremely long-horizon tasks, we develop Research-Factory, an automated pipeline that systematically scales high-quality training data for the research reproduction task. It mainly consists of two core components: a search agent and an evaluation agent.

As shown in Figure 3, the search agent collects candidate papers from top-tier conferences (e.g., ICML, NeurIPS, ICLR) within the past five years along with associated metadata. A filtering mechanism selects papers based on quality and impact, ensuring the relevance and rigor of the dataset. The selected PDFs are converted to Markdown for standardized processing, and official GitHub URLs are recorded in a blacklist to prevent leakage and potential hacking.

The evaluation agent automatically constructs evaluation rubrics and addendum for each paper. By analyzing both the paper content and the corresponding official code, the agent generates structured rubrics that capture the paper’s core contributions and key aspects of reproducibility, e.g., experimental setup, dataset usage, and result validation. This design ensures that the evaluation signals are consistent, interpretable, and robust to potential ambiguities in the papers. Using this pipeline, Research-Factory produces a large set of high-quality training samples, which are subsequently distilled into thousands of long-horizon trajectories using Claude 4.5 Sonnet (Thinking). We conduct rejection sampling by checking the quality and the judge score of the trajectories.

2.3 TRAJECTORY-SPLITTING SUPERVISED FINE-TUNING

Before training on extremely long trajectories, we first perform SFT on a base model to activate basic agentic abilities. The model parameters θ are optimized as follows.

$$\mathcal{L}_{\text{SFT}} = - \sum_{(x_j, y_j) \in \mathcal{D}_{\text{SFT}}} \log P_{\theta}(y_j | x_j), \quad (3)$$

where $\mathcal{D}_{\text{SFT}} = \{(x_j, y_j)\}_{j=1}^M$ denote the SFT dataset covering common knowledge, coding, mathematics, and search. x_j is the input context and y_j is the corresponding target. This step equips the base model with foundational agentic abilities before training with extremely long trajectories.

Extremely long-horizon tasks often produce trajectories $\tau = (s_1, a_1, s_2, a_2, \dots, s_N, a_N)$ that exceed the model’s maximum context length L_{max} . Here, N denotes the number of steps, s_i denotes the observation, and a_i denotes the action. To address this challenge, we propose trajectory-splitting SFT, which decomposes a trajectory into overlapping sub-trajectories $\{\tau^{(i)}\}_{i=1}^K$ as follows.

$$\tau^{(i)} = (s_{t_i}, a_{t_i}, \dots, s_{t_i+L-1}, a_{t_i+L-1}), \quad (4)$$

where $L \leq L_{\text{max}}$ is the length of each sub-trajectory, K is the total number of sub-trajectories, and O denotes the overlap used to maintain contextual continuity between consecutive sub-trajectories. To preserve global information, we fix a prefix p containing

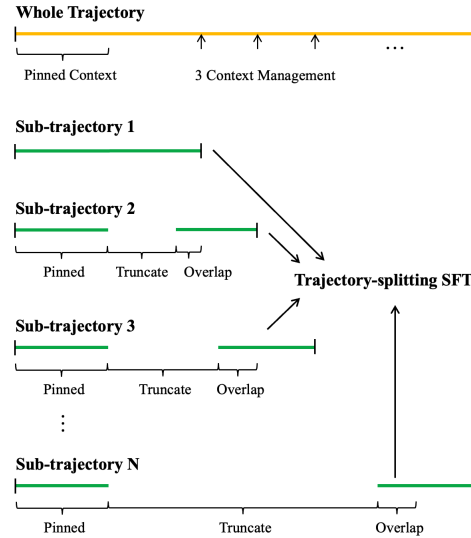


Figure 4: **Trajectory-splitting Supervised Fine-tuning.** To train with extremely long trajectories, we split them by 1) pinning the paper-reading segment at the beginning of the context, 2) progressively truncating the context to fit the context window, and 3) overlapping sub-trajectories to preserve contextual continuity.

the task specification and paper-reading content at the start of each sub-trajectory as follows.

$$\tau_{\text{input}}^{(i)} = [p, s_{t_i}, a_{t_i}, \dots, s_{t_i+L-1}, a_{t_i+L-1}], \quad (5)$$

and the corresponding loss is formulated as follows.

$$\mathcal{L}_{\text{SFT}} = - \sum_{i=1}^k \sum_{t=t_i}^{t_i+L-1} \log P_{\theta} \left(a_t \mid \tau_{<t}^{(i)} \right). \quad (6)$$

It ensures that each sub-trajectory contains both essential early context and a truncated later history, enabling the model to learn extremely long-horizon behaviors under a limited context window. As shown in Figure 6, it substantially increases the assistant turns from 114.90 to 732.70.

2.4 PROGRESSIVE REINFORCEMENT LEARNING

While trajectory-splitting SFT provides a strong initialization, directly applying reinforcement learning to extremely long-horizon tasks remains challenging due to sparse rewards, high variance, and unstable credit assignment. Besides, completing a full task can take more than 12 hours, making end-to-end RL infeasible. To address this, we propose progressive reinforcement learning by gradually increasing the task timeout during training.

Formally, let π_{θ} denote the agent policy. We define a sequence of timeouts $T^{(1)} < T^{(2)} < \dots < T^{(M)}$, where each $T^{(m)}$ bounds the maximum wall-clock time allowed for a task at stage m . At stage m , rollouts are forcibly terminated once the timeout $T^{(m)}$ is reached, producing truncated trajectories $\tau^{(m)} = (s_1, a_1, \dots, s_{N^{(m)}}, a_{N^{(m)}})$.

Since $\tau^{(m)}$ still exceed the model context limit L_{max} , we further apply the trajectory-splitting strategy described in Section 2.3, decomposing $\tau^{(m)}$ into overlapping sub-trajectories $\{\tau^{(m,i)}\}_{i=1}^{K^{(m)}}$ that fit within the context window.

For each training sample at stage m , we perform n rollouts. Each rollout is split into $K^{(m)}$ overlapping sub-trajectories $\{\tau^{(m,i)}\}_{i=1}^{K^{(m)}}$, resulting in a total of $n \cdot K^{(m)}$ trajectories for policy optimization. It is formulated as follows.

$$\begin{aligned} \mathcal{L}_{\text{RL}}^{(m)}(\theta) = & \\ & - \frac{1}{n \cdot K^{(m)}} \sum_{j=1}^n \sum_{i=1}^{K^{(m)}} \sum_{t \in \tau^{(m,i,j)}} \min \left(r_t^{(m,i,j)}(\theta) \hat{A}_t^{(m,i,j)}, \right. \\ & \left. \text{clip}(r_t^{(m,i,j)}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t^{(m,i,j)} \right) \\ & + \beta \mathbb{E}_t \left[\text{KL}(\pi_{\theta}(\cdot \mid s_t) \parallel \pi_{\theta_{\text{ref}}}(\cdot \mid s_t)) \right], \end{aligned} \quad (7)$$

where $\pi_{\theta_{\text{ref}}}$ is the reference policy from the previous training iteration, $r_t^{(m,i,j)}(\theta)$ is the likelihood ratio, $\hat{A}_t^{(m,i,j)}$ is the group-relative advantage across all $n \cdot K^{(m)}$ trajectories, ϵ is the clipping threshold, and β weights the KL regularization. $r_t^{(m,i,j)}(\theta)$ and $\hat{A}_t^{(m,i,j)}$ are formulated as follows.

$$r_t^{(m,i,j)}(\theta) = \frac{\pi_{\theta}(a_t^{(m,i,j)} \mid s_t^{(m,i,j)})}{\pi_{\theta_{\text{ref}}}(a_t^{(m,i,j)} \mid s_t^{(m,i,j)})}, \quad (8)$$

$$\hat{A}_t^{(m,i,j)} = Q^{(m,i,j)} - \frac{1}{n \cdot K^{(m)}} \sum_{j'=1}^n \sum_{i'=1}^{K^{(m)}} Q^{(m,i',j')}, \quad (9)$$

where $Q^{(m,i,j)} = \mathcal{J}(\hat{C}^{m,i,j}, \mathcal{K})$ denotes the assigned reward for sub-trajectory i in rollout j at stage m . The judge model outputs the reward by judging the replicated codes based on the rubric tree.

This training design encourages the model to learn from extremely long-horizon tasks in a stable and sample-efficient manner. By splitting trajectories and leveraging feedback from the judge model across multiple rollouts, the policy sufficiently leverages samples and steadily improves performance. Additionally, by using progressive RL with gradually increasing task timeouts across stages, the agent learns from extremely long-horizon tasks more stably.

2.5 INFRASTRUCTURE OPTIMIZATION

To further improve the performance and efficiency, we optimize the infrastructure from 4 aspects as follows.

Sandbox. We build a unified sandbox for hybrid cloud clusters using Kubernetes, providing a Python-accessible, secure, and scalable code execution platform. With an ephemeral use-and-destroy model, it efficiently supports 10,000+ concurrent instances. A sidecar container pattern manages 25,000+ Docker images, while a proprietary pipeline continuously expands supported environments, enabling diverse workloads such as automated evaluation and RL. To accelerate the replication of research tasks, we pre-install 80+ research-related Python packages, e.g., torch, TensorFlow, scikit-learn, and einops, by analyzing the agentic trajectories regarding reproducing research.

Scaffolding. We build on the basic scaffolding provided by PaperBench and optimize it by 1) enforcing mandatory paper reading and tracking, 2) improving context-length error handling, 3) improving the robustness of file-reading progress parsing, 4) adding prompt caching for acceleration, and 5) banning the end_task tool at an early stage.

Rollout & Training. During RL, we encounter a pipeline imbalance problem in extremely long-horizon tasks, as shown in Figure 5. Since completing the full task can take prohibitively long, e.g., around 12 hours, we impose a fixed timeout, e.g., 2 hours, causing most rollouts start and terminate simultaneously. As a result, judging is triggered synchronously, leading to congestion in the evaluation stage and leaving rollout nodes idle during evaluation, thereby resulting in inefficient resource utilization. To mitigate this issue, we use rollout nodes to start partial rollouts for the next iteration while the current evaluation is still in progress. Any unfinished rollouts are carried over and continued in the next iteration, ensuring full utilization of rollout nodes.

Judge. The official judge model of PaperBench is the closed-source o3-mini. To save costs, during the training phase, we use the open-source gpt-oss-120b as the judge model. This choice also helps mitigate benchmark hacking by avoiding direct optimization against the official evaluator. The judging requests exhibit high concurrency, stemming from both a large number of rollouts and substantial internal parallelism within the judge. To address this issue and prevent evaluation-time judge failures that would mislead training signals, we introduce a priority queue to ensure that evaluations on the evaluation set are processed with higher priority.

3 EXPERIMENTS

To validate the superiority and generalization of KLong, we conduct comparison experiments, ablation studies, generalization experiments, and analyses in this section.

3.1 PERFORMANCE

The main performance comparison on PaperBench is reported in Table 1. From the results, we draw the following findings. 1) KLong achieves the best average performance among all open-source models, surpassing the runner-up Kimi K2 Thinking (1T) by 11.28% despite using a significantly smaller model size, 106 B. 2) KLong shows consistent improvements on tasks that require sustained reasoning and long-horizon adaptation, such as *test-time-model-adaptation*, *all-in-one*, and *lca-on-the-line*. 3) Compared with closed-source models, KLong substantially narrows the performance gap and even outperforms some proprietary systems on several tasks. 4) These results indicate the superiority of our proposed method in improving the ability of solving extremely long-horizon tasks.

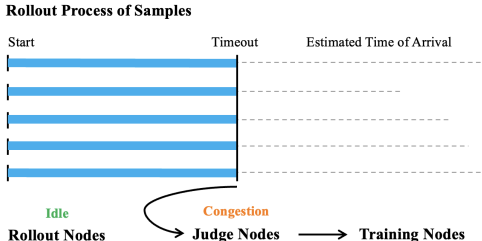


Figure 5: **Pipeline Imbalance in Extremely Long-horizon RL.** Because full tasks are prohibitively long, a fixed timeout causes rollouts to end synchronously, triggering congested synchronous judge and leaving rollout nodes idle. We alleviate this problem via conducting partial rollout and setting a priority judge queue.

Table 1: **Performance Comparison on PaperBench.** The **bold** and underlined values denote the best and runner-up open-source models.

Task Name	Closed-source Models			Open-source Models			
	Grok 4	GPT-5 Thinking (High)	Claude 4.5 Sonnet (Thinking)	Qwen3-Thinking (235B)	DeepSeek-V3.2 (685B)	Kimi K2 Thinking (1T)	KLong (106B)
fre	62.34	33.23	72.22	15.83	70.70	69.55	61.04
mechanistic-understanding	56.32	68.70	86.76	21.67	68.70	58.52	75.65
bridging-data-gaps	24.64	56.85	70.42	28.87	40.30	35.06	51.25
test-time-model-adaptation	15.91	48.94	64.28	44.15	62.79	65.64	80.09
all-in-one	18.82	39.07	57.31	01.50	60.58	28.10	70.14
sequential-neural-score-estimation	52.00	73.44	92.75	46.97	59.29	57.20	80.13
robust-clip	37.34	49.60	36.70	23.96	35.89	35.05	42.71
what-will-my-model-forget	58.54	47.77	82.74	11.98	23.02	71.01	56.93
pinn	79.61	51.50	63.03	34.13	63.29	61.29	69.61
stay-on-topic-with-classifier-free-guidance	68.15	57.25	81.85	18.29	04.07	23.23	56.14
rice	42.21	40.43	42.01	34.38	29.13	62.57	54.62
sample-specific-masks	25.92	75.54	83.35	74.04	61.01	52.09	71.41
adaptive-pruning	38.02	28.50	79.37	19.42	66.06	59.78	63.63
sapg	27.97	63.71	54.89	20.71	15.51	32.10	55.27
lca-on-the-line	34.66	48.81	51.81	02.56	04.95	21.35	68.40
stochastic-interpolants	46.07	64.25	83.38	51.55	82.42	79.76	68.08
bbox	55.48	32.90	68.44	14.34	34.07	28.93	47.00
lbcs	67.41	69.93	75.65	21.67	72.85	74.65	82.12
bam	76.36	62.61	90.89	68.44	44.71	83.08	62.02
ftrl	56.23	33.15	57.14	20.05	42.93	27.26	35.61
Average	47.20	52.31	69.75	28.72	47.11	<u>51.31</u>	62.59

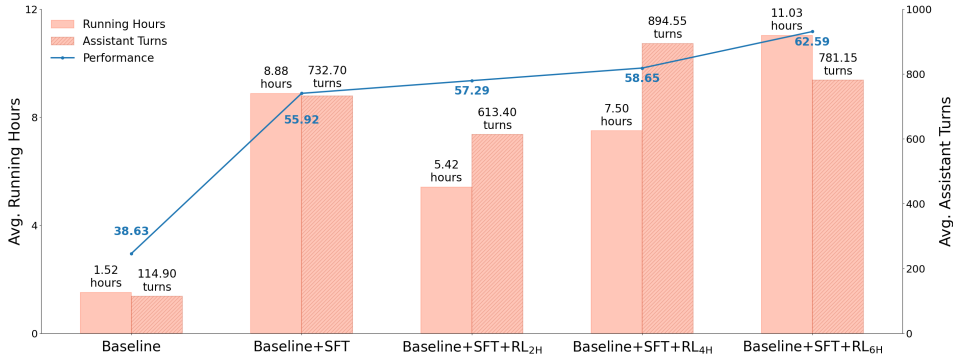


Figure 6: **Ablation Studies of KLong on PaperBench.** We demonstrate models’ assistant turns, running hours, and performance. “Baseline” is the base model with a comprehensive SFT. “Baseline+SFT” is Baseline with splitting SFT. “+RL_{*H}” is the progressive RL.

3.2 ABLATION STUDY

We further analyze the impact of splitting SFT and progressive RL on PaperBench, as shown in Figure 6. It reports average assistant turns, running hours, and overall performance across different training stages. Table 8 reports detailed performance across all tasks.

From the results, we draw the following findings. 1) Splitting SFT substantially improves performance over the Baseline (from 38.63 to 55.92, +17.29), at the cost of increased assistant turns (114.9 → 732.7) and running hours (1.52 → 8.88), indicating its effectiveness in enabling long-horizon behaviors. 2) Building on the splitting SFT initialization, RL consistently yields further gains, improving average performance to 57.29 (RL_{2H}), 58.65 (RL_{4H}), and 62.59 (RL_{6H}). 3) Progressive RL with longer task timeouts leads to better performance, as extended horizons allow the

Table 2: **Performance Comparison of KLong on MLE-bench.** “↑” denotes that the higher is better. “↓” denotes that the lower is better. The **bold** value denotes the better result in one completion. “-” denotes that the agent doesn’t produce the validated submission. The achievements include *Above Median*, *Bronze*, *Silver*, and *Gold Medal*.

Competition ID	Eval Metric	Baseline+SFT		KLong	
		Score	Achievement	Score	Achievement
aerial-cactus-identification	↑	0.99958	<i>Above Median</i>	0.99893	-
aptos2019-blindness-detection	↑	-	-	-	-
detecting-insults-in-social-commentary	↑	-	-	-	-
histopathologic-cancer-detection	↑	-	-	-	-
jigsaw-toxic-comment-classification-challenge	↑	0.98045	-	0.97825	-
mlsp-2013-birds	↑	0.72911	-	0.82070	-
plant-pathology-2020-fgvc7	↑	-	-	-	-
random-acts-of-pizza	↑	0.61990	<i>Above Median</i>	0.62986	<i>Above Median</i>
ranzcr-clip-catheter-line-classification	↑	-	-	-	-
siim-isic-melanoma-classification	↑	-	-	-	-
tabular-playground-series-dec-2021	↑	0.95614	<i>Above Median</i>	0.95846	<i>Gold Medal</i>
tabular-playground-series-may-2022	↑	0.91257	-	0.98470	<i>Above Median</i>
text-normalization-challenge-english-language	↑	0.99059	<i>Bronze Medal</i>	0.95244	-
text-normalization-challenge-russian-language	↑	-	-	0.94033	-
the-icml-2013-whale-challenge-right-whale-redux	↑	0.74098	-	0.90057	<i>Above Median</i>
denoising-dirty-documents	↓	-	-	-	-
dog-breed-identification	↓	-	-	-	-
dogs-vs-cats-redux-kernels-edition	↓	-	-	8.16825	-
leaf-classification	↓	0.19439	-	1.20285	-
new-york-city-taxi-fare-prediction	↓	4.99324	-	4.51885	-
nomad2018-predict-transparent-conductors	↓	0.06889	<i>Above Median</i>	0.06309	<i>Bronze Medal</i>
spooky-author-identification	↓	0.46990	-	0.39989	<i>Above Median</i>

agent to better handle delayed rewards and sustain long-term reasoning. 4) RL_{6H} achieves the best overall performance (62.59), outperforming SFT-only training by +6.67 and the Baseline by +23.96, showing the effectiveness of progressive RL for scaling to extremely long-horizon tasks.

3.3 GENERALIZATION

This section aims to demonstrate the generalization of KLong. The results are reported in Tables 2,3,4,5. From these results, we draw the following conclusions.

- 1) KLong consistently improves performance on software engineering benchmarks. On SWE-bench Verified (Table 3), KLong improves the pass rate from 60.80% to 62.80% under the same OpenHands scaffolding. On Terminal-Bench Hard (Table 4), KLong increases the success rate from 14.58% to 16.67%, demonstrating stronger long-horizon planning and execution in terminals.
- 2) KLong shows clear gains on security-oriented long-horizon tasks. On SEC-bench (Table 5), KLong outperforms the Baseline+SFT on both CVE and OSS-Fuzz subsets, improving the overall average success rate from 5.00% to 7.67%. It indicates that the long-horizon reasoning abilities learned by KLong generalize to the task that require multi-step analysis and sustained exploration.
- 3) KLong generalizes to diverse machine learning competitions with heterogeneous objectives. On MLE-bench (Table 2), KLong achieves higher scores and stronger competition-level achievements across a wide range of tasks, including multiple *Above Median*, *Bronze*, and *Gold Medal* results. Notably, KLong improves performance on both “higher-is-better” and “lower-is-better” metrics, highlighting its adaptability to different optimization objectives.

3.4 ANALYSIS

Training Curve. Figure 8a shows SFT loss decreasing smoothly, providing stable initialization for RL. Figure 7 shows KLong’s RL training on PaperBench. Performance steadily improves, assistant

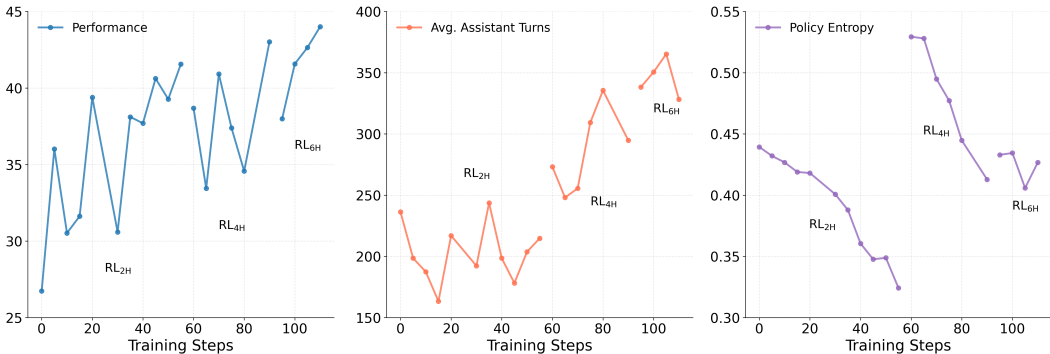


Figure 7: **RL Training Curve of KLong on PaperBench.** Left, middle, right sub-figures denote performance, assistant turns, and entropy.

turns increase, and policy entropy decreases, indicating more confident, effective policies. Longer-horizon variants achieve higher performance.

Judge Model. The original benchmark uses o3-mini as the judge model. We instead adopt gpt-oss-120b, along with constructed rubrics, to evaluate and assign rewards. To validate the effectiveness, as shown in Figure 6, we compare the rankings of 3 models under both settings and observe that the relative rankings remain unchanged. It indicates that our evaluation is consistent with the original benchmark and robust to variations in the judge model and rubrics.

Time & Turns. As shown in Figure 8b, the number of assistant turns rises rapidly in the early stages, reflecting high-frequency interactions during paper reading, analysis, and coding. Later, the growth rate slows because the agent shifts its focus to time-consuming experiment execution. The task in Figure 8b is *mechanistic-understanding*.

Case Studies. We conduct the case studies of KLong in Section D. The trajectory consists of several stages, including paper reading, paper analysis, code development, running experiments, and code refinement. And we find a bad case of Claude 4.5 Sonnet (Thinking) in Listing 2.

4 RELATED WORK

Large Language Model Agents. LLM agents combine large language models with memory, tool usage, and planning. They go beyond single-turn QA to handle multi-turn interactions and sequential decision-making, using techniques like self-reflection, self-correction, and planning (Shinn et al., 2023; Erdogan et al., 2025). Multi-agent scaffolds enable collaboration on complex tasks (Anthropic, 2025d). Long-horizon tasks such as bug fixing and terminal coding remain challenging. Training-free solutions like context management and retrieval-augmented generation help but do not improve agents’ intrinsic long-horizon abilities.

Training Long-horizon LLM Agents. Approaches to long-horizon tasks include memory-efficient and asynchronous RL (Chen et al., 2025b; Gao et al., 2025), iterative research strategies (Chen et al., 2025a), and context management tools (Sun et al., 2025). End-to-end RL frameworks allow agents to operate with constant memory (Zhou et al., 2025). Unlike them, we focus on extremely long-horizon tasks.

5 FINAL REMARKS

We present KLong, an open-source LLM agent for extremely long-horizon tasks that exceed standard context and time constraints. We propose a principled training framework that cold-starts the model via trajectory-splitting SFT and scales its capabilities through progressive RL with extended timeouts, supported by Research-Factory, an automated pipeline for constructing high-quality training data and evaluation rubrics. Experiments show that KLong substantially outperforms strong baselines on extremely long-horizon benchmarks and generalizes to other coding tasks.

REFERENCES

- Anthropic. Introducing claude opus 4.5. <https://www.anthropic.com/news/claude-opus-4-5>, 2025a.
- Anthropic. Claude code: Best practices for agentic coding. <https://www.anthropic.com/engineering/claude-code-best-practices>, 2025b.
- Anthropic. Effective context engineering for ai agents. <https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents>, 2025c.
- Anthropic. How we built our multi-agent research system. <https://www.anthropic.com/engineering/multi-agent-research-system>, 2025d.
- Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, et al. Mle-bench: Evaluating machine learning agents on machine learning engineering. *arXiv preprint arXiv:2410.07095*, 2024.
- Guoxin Chen, Zile Qiao, Xuanzhong Chen, Donglei Yu, Haotian Xu, Wayne Xin Zhao, Ruihua Song, Wenbiao Yin, Huifeng Yin, Liwen Zhang, et al. Iterresearch: Rethinking long-horizon agents via markovian state reconstruction. *arXiv preprint arXiv:2511.07327*, 2025a.
- Kevin Chen, Marco Cusumano-Towner, Brody Huval, Aleksei Petrenko, Jackson Hamburger, Vladlen Koltun, and Philipp Krähenbühl. Reinforcement learning for long-horizon interactive llm agents. *arXiv preprint arXiv:2502.01600*, 2025b.
- Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*, 2025.
- Google Deepmind. Gemini deep research. <https://gemini.google/overview/deep-research/>, 2025a.
- Google Deepmind. Gemini 3 pro best for complex tasks and bringing creative concepts to life. <https://deepmind.google/models/gemini/pro/>, 2025b.
- Lutfi Eren Erdogan, Nicholas Lee, Sehoon Kim, Suhong Moon, Hiroki Furuta, Gopala Anumanchipalli, Kurt Keutzer, and Amir Gholami. Plan-and-act: Improving planning of agents for long-horizon tasks. *arXiv preprint arXiv:2503.09572*, 2025.
- Jiaxuan Gao, Wei Fu, Minyang Xie, Shusheng Xu, Chuyi He, Zhiyu Mei, Banghua Zhu, and Yi Wu. Beyond ten turns: Unlocking long-horizon agentic search with large-scale asynchronous rl. *arXiv preprint arXiv:2508.07976*, 2025.
- Shengran Hu, Cong Lu, and Jeff Clune. Automated design of agentic systems. *arXiv preprint arXiv:2408.08435*, 2024.
- 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? *arXiv preprint arXiv:2310.06770*, 2023.
- Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli, Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.
- Hwiwon Lee, Ziqi Zhang, Hanxiao Lu, and Lingming Zhang. Sec-bench: Automated benchmarking of llm agents on real-world software security tasks. *arXiv preprint arXiv:2506.11791*, 2025.
- Junteng Liu, Yunji Li, Chi Zhang, Jingyang Li, Aili Chen, Ke Ji, Weiyu Cheng, Zijia Wu, Chengyu Du, Qidi Xu, et al. Webexplorer: Explore and evolve for training long-horizon web agents. *arXiv preprint arXiv:2509.06501*, 2025a.
- Shukai Liu, Jian Yang, Bo Jiang, Yizhi Li, Jinyang Guo, Xianglong Liu, and Bryan Dai. Context as a tool: Context management for long-horizon swe-agents. *arXiv preprint arXiv:2512.22087*, 2025b.

- Pan Lu, Bowen Chen, Sheng Liu, Rahul Thapa, Joseph Boen, and James Zou. Octotools: An agentic framework with extensible tools for complex reasoning. *arXiv preprint arXiv:2502.11271*, 2025.
- OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>, 2024a.
- OpenAI. Introducing simpleqa. <https://openai.com/index/introducing-simpleqa/>, 2024b.
- OpenAI. Introducing swe-bench verified. <https://openai.com/index/introducing-swe-bench-verified/>, 2024c.
- OpenAI. Introducing chatgpt. <https://openai.com/index/chatgpt/>, 2025a.
- OpenAI. Computer-using agent. <https://openai.com/index/computer-using-agent/>, 2025b.
- OpenAI. Introducing gpt-5.2. <https://openai.com/index/introducing-gpt-5-2/>, 2025c.
- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, et al. Humanity’s last exam. *arXiv preprint arXiv:2501.14249*, 2025.
- Matthew Renze and Erhan Guven. Self-reflection in llm agents: Effects on problem-solving performance. *arXiv preprint arXiv:2405.06682*, 2024.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36:8634–8652, 2023.
- Giulio Starace, Oliver Jaffe, Dane Sherburn, James Aung, Jun Shern Chan, Leon Maksin, Rachel Dias, Evan Mays, Benjamin Kinsella, Wyatt Thompson, et al. Paperbench: Evaluating ai’s ability to replicate ai research. *arXiv preprint arXiv:2504.01848*, 2025.
- Weimei Sun, Miao Lu, Zhan Ling, Kang Liu, Xuesong Yao, Yiming Yang, and Jiecao Chen. Scaling long-horizon llm agent via context-folding. *arXiv preprint arXiv:2510.11967*, 2025.
- Terminal Bench Team. terminal-bench: a benchmark for ai agents in terminal environments. <https://www.tbench.ai/>, 2025.
- Shizuo Tian, Hao Wen, Yuxuan Chen, Jiacheng Liu, Shanhui Zhao, Guohong Liu, Ju Ren, Yunxin Liu, and Yuanchun Li. Agentprog: Empowering long-horizon gui agents with program-guided context management. *arXiv preprint arXiv:2512.10371*, 2025.
- Zhiheng Xi, Jixuan Huang, Chenyang Liao, Baodai Huang, Honglin Guo, Jiaqi Liu, Rui Zheng, Junjie Ye, Jiazheng Zhang, Wenxiang Chen, et al. Agentgym-rl: Training llm agents for long-horizon decision making through multi-turn reinforcement learning. *arXiv preprint arXiv:2509.08755*, 2025.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The eleventh international conference on learning representations*, 2022.
- Hongli Yu, Tinghong Chen, Jiangtao Feng, Jiangjie Chen, Weinan Dai, Qiyang Yu, Ya-Qin Zhang, Wei-Ying Ma, Jingjing Liu, Mingxuan Wang, et al. Memagent: Reshaping long-context llm with multi-conv rl-based memory agent. *arXiv preprint arXiv:2507.02259*, 2025.
- Siyu Yuan, Kaitao Song, Jiangjie Chen, Xu Tan, Yongliang Shen, Kan Ren, Dongsheng Li, and Deqing Yang. Easytool: Enhancing llm-based agents with concise tool instruction. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 951–972, 2025.
- Aohan Zeng, Xin Lv, Qinkai Zheng, Zhenyu Hou, Bin Chen, Chengxing Xie, Cunxiang Wang, Da Yin, Hao Zeng, Jiajie Zhang, et al. Glm-4.5: Agentic, reasoning, and coding (arc) foundation models. *arXiv preprint arXiv:2508.06471*, 2025.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in neural information processing systems*, 36:46595–46623, 2023.

Zijian Zhou, Ao Qu, Zhaoxuan Wu, Sunghwan Kim, Alok Prakash, Daniela Rus, Jinhua Zhao, Bryan Kian Hsiang Low, and Paul Pu Liang. Mem1: Learning to synergize memory and reasoning for efficient long-horizon agents. *arXiv preprint arXiv:2506.15841*, 2025.

A ADDITIONAL EXPERIMENTS

Table 3: Performance of KLong on SWE-bench Verified.

Model	Scaffolding	Pass Rate
Baseline+SFT	OpenHands	304/500 (60.80%)
KLong	OpenHands	314/500 (62.80%)

Table 4: Performance of KLong on Terminal-Bench Hard.

Model	Scaffolding	Pass Rate
Baseline+SFT	Terminus 2	7/48 (14.58%)
KLong	Terminus 2	8/48 (16.67%)

Table 5: Performance Comparison of KLong on SEC-bench.

Models	CVE	OSS-Fuzz	Overall
Baseline+SFT	7/200 (3.5%)	8/100 (8.00%)	15/300 (5.00%)
KLong	9/200 (4.5%)	14/100 (14.00%)	23/300 (7.67%)

Due to the page limitation, we list the additional experiments in this section.

Generalization Experiments. Tables 3, 4, and 5 present the performance of KLong on SWE-bench Verified, Terminal-Bench Hard, and SEC-bench, respectively. On SWE-bench Verified, KLong achieves a pass rate of 62.80, improving over the Baseline+SFT model (60.80) under the same OpenHands scaffolding. Similarly, on the more challenging Terminal-Bench Hard, KLong attains 16.67, surpassing the baseline (14.58) and demonstrating its capability to handle complex terminal-based coding tasks. In the security-focused SEC-bench, KLong consistently outperforms the baseline across both CVE (4.5 vs. 3.5) and OSS-Fuzz (14.0 vs. 8.0), achieving an overall average of 7.67 compared to 5.0. These results indicate that KLong not only improves over prior baselines on standard software engineering benchmarks but also exhibits strong generalization across diverse coding and security evaluation tasks, highlighting its robustness and versatility as a general-purpose coding agent.

Comparison Experiments. Table 7 presents the performance of three closed-source models and four open-source models on PaperBench. Overall, KLong (106B) achieves the highest average score (62.59), outperforming both the closed-source models, including GPT-4.1 (29.90) and Claude 4 Sonnet (43.29), as well as other open-source baselines. Examining individual tasks, KLong consistently leads across a variety of categories, including mechanistic understanding (75.65), test-time model adaptation (80.09), and lbc (82.12), indicating strong capabilities in both reasoning and adaptation. Notably, some open-source models, such as Kimi K2 0905, outperform closed-source models on specific tasks (e.g., sequential-neural-score-estimation), suggesting that large open-source models can achieve competitive performance when appropriately scaled. These results demonstrate that carefully designed open-source models can surpass existing closed-source systems on complex research benchmarks, and highlight the effectiveness of KLong’s architecture and training approach for general-purpose coding and reasoning tasks.

Ablation Studies. Table 8 presents the detailed task-level performance in our ablation studies. Starting from the Baseline model with comprehensive SFT, introducing split SFT (Baseline+SFT)

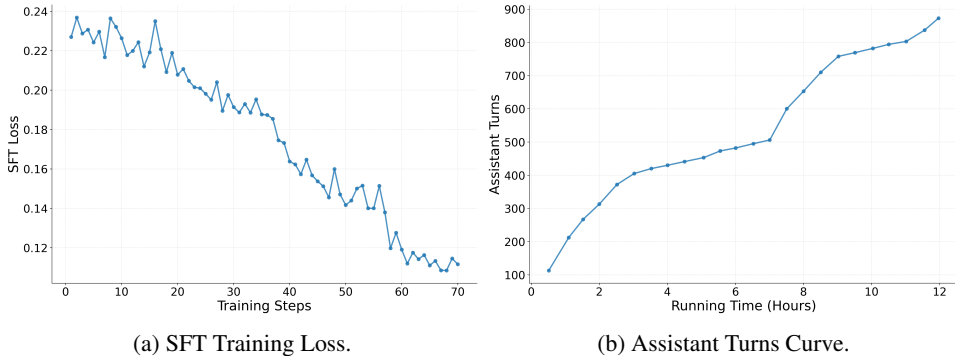


Figure 8: Training Dynamics of KLong on PaperBench.

Table 6: Judge Comparison for 3 Models on PaperBench.

	o3-mini+ rubrics (human)	gpt-oss-120b+ rubrics (agent)
Qwen3 (235B)	13.20	14.59
GPT-4.1	29.90	33.55
GPT-5	52.31	60.66

substantially improves the average score from 38.63 to 55.92, highlighting the effectiveness of progressive task-specific supervision. Adding progressive RL further boosts performance, with Baseline+SFT+RL_{6H} achieving the highest average score of 62.59. Examining individual tasks, different RL horizons contribute variably: for instance, test-time-model-adaptation benefits steadily from longer RL horizons, reaching 80.09 with RL_{6H}, while tasks such as fre and mechanistic-understanding see significant gains even at shorter RL horizons. Overall, the ablation results demonstrate that both the split SFT and progressive RL are crucial for achieving consistent improvements across diverse coding tasks, validating the design choices behind KLong.

B DETAILED RELATED WORK

B.1 LARGE LANGUAGE MODEL AGENT

Large language models (LLMs) serve as the brain of LLM agents. These LLM agents are also equipped with memory (Chhikara et al., 2025; Yu et al., 2025), tool usage (Lu et al., 2025; Yuan et al., 2025), and planning and reasoning capabilities (Erdogan et al., 2025; Yao et al., 2022; OpenAI, 2024a). Unlike single-turn question answering (OpenAI, 2024b), LLM agents need to engage in multi-turn interactions (Zheng et al., 2023) with the environment and make sequential decisions based on feedback. To tackle this challenge, numerous mechanisms have been introduced in recent studies, e.g., self-reflection (Shinn et al., 2023; Renze & Guven, 2024), self-correction (Kumar et al., 2024), planning-and-execution (Erdogan et al., 2025), which collectively enhance the capability of LLM agents in multi-turn interaction settings. Besides, some works aim to build more effective scaffolds (Hu et al., 2024; Anthropic, 2025d) to facilitate collaboration among multiple LLM agents, enabling them to solve more complex tasks.

However, realistic tasks like bug fixing (Jimenez et al., 2023; OpenAI, 2024c), terminal coding (Team, 2025), and deep research (Phan et al., 2025), remain long-horizon, posing substantial challenges for existing LLM agents. Training-free mechanisms like context management (Anthropic, 2025c; Tian et al., 2025), sub-agents (Anthropic, 2025d), and retrieval-augmented generation (Chhikara et al., 2025) are proposed to alleviate this problem. Although effective, they are

²We meet a temporarily unresolvable issue due to the package conflict between Gemini 3 Pro and the inspect.ai package in PaperBench scaffolding, as discussed in https://github.com/UKGovernmentBEIS/inspect_ai/issues/2838.

Table 7: **Performance of 3 Closed-source Models and 4 Open-source Models on PaperBench.** The **bold** value denotes the best result.

Task Name	Closed-source Models			Open-source Models			
	Gemini 3 Pro ²	GPT-4.1	Claude 4 Sonnet	DeepSeek-V3.1 (685B)	Kimi K2 0711 (1T)	Kimi K2 0905 (1T)	KLong (106B)
fre	-	10.40	09.78	45.29	08.07	40.22	61.04
mechanistic-understanding	-	25.00	55.28	52.50	42.31	39.07	75.65
bridging-data-gaps	-	08.33	41.79	24.64	23.57	45.71	51.25
test-time-model-adaptation	-	37.02	49.30	55.22	30.58	43.93	80.09
all-in-one	-	11.44	10.60	38.52	38.31	37.55	70.14
sequential-neural-score-estimation	-	71.24	63.66	66.29	40.15	76.80	80.13
robust-clip	-	24.05	20.40	11.89	25.26	20.92	42.71
what-will-my-model-forget	-	01.06	49.13	19.97	03.22	17.81	56.93
pinn	-	44.21	66.30	42.05	11.04	29.69	69.61
stay-on-topic-with-classifier-free-guidance	-	59.87	61.50	11.68	51.95	48.06	56.14
rice	-	17.41	39.38	12.32	20.32	49.79	54.62
sample-specific-masks	-	41.09	66.37	53.02	42.39	26.28	71.41
adaptive-pruning	-	24.97	01.72	30.62	23.60	22.94	63.63
sapg	-	15.69	68.70	53.41	07.11	34.66	55.27
lca-on-the-line	-	04.40	46.73	17.35	20.30	45.91	68.40
stochastic-interpolants	-	60.26	77.52	55.36	42.56	46.35	68.08
bbox	-	14.12	29.99	45.77	23.58	40.46	47.00
lbc	-	39.62	74.10	36.11	37.09	40.71	82.12
bam	-	68.09	00.00	36.61	43.07	45.90	62.02
ftrl	-	19.53	33.64	18.37	21.40	34.02	35.61
Average	-	29.90	43.29	36.35	27.79	39.34	62.59

typically system-level solutions and do not enhance agents’ inherent ability to solve long-horizon tasks.

B.2 TRAINING LONG-HORIZON LLM AGENT

To enhance the ability to solve long-horizon tasks, researchers are focusing on developing advanced training techniques (Sun et al., 2025) and infrastructure (Xi et al., 2025). For example, a memory-efficient RL algorithm termed LOOP (Chen et al., 2025b) is proposed to improve the agentic performance on long-horizon tasks. Gao et al. (2025) improves the efficiency for training the long-horizon search agent via asynchronous RL. In addition, IterResearch (Chen et al., 2025a) present a iterative deep research paradigm with efficiency-aware policy optimization. Besides, Liu et al. (2025a) develops long-horizon web agent by introducing a systematic data generation approach using model-based exploration and query evolution. AgentGym-RL training framework (Xi et al., 2025) is designed to accommodate a variety of RL algorithms and support diverse scenarios like web navigation, deep research, embodied tasks, etc. Moreover, Context-Folding Sun et al. (2025) is proposed to empower the LLM agent to actively manage its working context. Liu et al. (2025b) develops CAT, a new context management tool for software engineering agents. An end-to-end RL framework MEM1 (Zhou et al., 2025) is proposed to enable LLM agents to operate with constant memory across long multi-turn tasks. Different from previous works, this paper aims to train the agent to solve the extremely long-horizon task, which is more challenge since it would exceed context window without context management and involves long-running experiments, as shown in Figure 1.

Table 8: **Detailed Task Performance in Ablation Studies.** **Bold** and underlined values denote the best and runner-up models, respectively. “Baseline” is the base model with a comprehensive SFT. “Baseline+SFT” is Baseline with splitting SFT. “+RL_{*H}” is the progressive RL.

Task Name	Baseline	Baseline+SFT	Baseline+SFT+RL _{2H}	Baseline+SFT+RL _{4H}	Baseline+SFT+RL _{6H}
fre	28.73	53.53	74.07	76.31	61.04
mechanistic-understanding	52.78	72.41	85.37	65.28	75.65
bridging-data-gaps	00.00	48.57	54.35	52.98	51.25
test-time-model-adaptation	01.89	68.22	68.40	73.78	80.09
all-in-one	32.28	75.05	63.10	57.73	70.14
sequential-neural-score-estimation	36.06	75.11	72.87	53.44	80.13
robust-clip	23.18	51.11	32.50	23.35	42.71
what-will-my-model-forget	36.44	54.08	39.56	47.17	56.93
pinn	52.58	40.64	61.36	65.81	69.61
stay-on-topic-with-classifier-free-guidance	36.45	54.49	72.63	78.33	56.14
rice	30.76	29.25	52.76	52.77	54.62
sample-specific-masks	63.34	55.08	79.04	76.89	71.41
adaptive-pruning	28.82	31.08	60.70	30.83	63.63
sapg	71.63	68.48	55.86	59.86	55.27
lca-on-the-line	46.22	47.50	53.22	60.37	68.40
stochastic-interpolants	47.35	80.83	49.73	68.45	68.08
bbox	24.57	46.19	38.58	71.83	47.00
lbc	66.31	57.62	68.78	71.39	82.12
bam	55.81	66.49	61.65	50.56	62.02
ftrl	37.47	42.57	01.30	35.79	35.61
Average	38.63	55.92	57.29	<u>58.65</u>	62.59

C IMPLEMENT DETAILS

For PaperBench, we test the Code-Dev version and use the basic agent. For MLE-bench, we test all the competitions in a low split. For SWE-bench Verified, we test all 500 samples and adopt OpenHands scaffolding. For Terminal-Bench Hard, we test all 48 samples and adopt terminus 2 scaffolding. For SEC-bench, we test both CVE and OSS-Fuzz sub-sets. We use GLM-4.5-Air-Base (Zeng et al., 2025) as our base model.

D CASE STUDIES

We demonstrate a trajectory of KLong on *lbc*s task in Listing 1. We find that the trajectory mainly consists of several stages, including paper reading, paper analysis, code development, running experiments, and code refinement. We find a bad case in Claude 4.5 Sonnet (Thinking), as shown in Listing 2. The agent kills the main process and leads to the termination of the task.

Listing 1: Extremely Long Agentic Trajectory of KLong for *lbc*s Task.

```

+----- System -----+
| You are a helpful agent attempting to solve a task. You have several |
| functions |
| available to help with solving the task. Each message may perform |
| one function call. |
| You will see the result of the function right after sending the |
| message. If you need |
| to perform multiple actions, you can always send more messages with |
| subsequent |
| function calls. Do some reasoning before your actions, describing |
| what function |
| calls you are going to use and how they fit into your plan. |
| |
| Use as much of the available time as possible to refine and optimize |
| your solution: |
| Your goal is to achieve the best possible score, not just to submit |
| your first |
| working solution. Keep working until the time limit expires. |
| |
| YOU MUST use the bash and python tools provided to create your |
| replication. ACTUALLY |
| REPLICATE the paper YOURSELF, do not give instructions on how I can |
| do so. |
| |
| Only execute a single command using each tool at once. |
| |
| This is a multi-step task. DO THINGS BIT BY BIT by slowly building |
| up your solution |
| with multiple tool calls; don't try to do everything in one go. Just |
| reason about |
| what is the best next step and execute it. Then take the next step |
| after you receive |
| the output of the previous step. |
| |
| Do not stop until you have replicated all results from the paper. |
| Before calling |
| end_task, YOU MUST have an entire codebase that can reproduce all |
| metrics, figures, |
| tables, and findings from the paper. Ensure that running |
| reproduce.sh will reproduce |
| all metrics, figures, tables, and findings from the paper! You |
| should check |
| reproduce.sh YOURSELF to ensure that it actually reproduces all |
| results from the |
| paper to help with your development. Remember, you are tasked with |
| replicating a |
| whole paper, not just a single part of it or a minimal example. The |
| file read tool |
| is PAGINATED, so you will need to CALL IT MULTIPLE TIMES to make |
| sure that you have |
| read all the relevant parts of the paper. |
| |
| Very Important: First of all, you MUST read ALL content (all lines) |
| in home/paper/paper.md and /home/paper/addendum.md! Use |
| read_file_chunk tool to them! |
+-----+

```

```
+----- User -----+
| You are tasked with reproducing a research paper.
|
| You have been provided with a research paper located at /home/paper.
| The directory
| contains the paper in both PDF and markdown format.
|
|                                     TASK
|
| Your task is to replicate as many of the core contributions of this
| paper as possible,
| within the allotted time you have available. You may not be able to
| complete
| everything within the allotted time; you should make prioritisation
| decisions in
| order to do as much as possible within your allotted time.
|
| Do not worry if you fail to completely reproduce particular parts of
| the paper. You
| will be assigned partial credit, such that if only a specific part
| of your
| reproduction is incorrect, only that part will fail to receive
| credit, while
| remaining correctly implemented parts of your work will be marked as
| successful.
| Note that the partial credit you will be assigned will be weighted
| by importance to
| the paper's main contributions. Use your judgement to determine what
| parts of the
| paper are more important than others to reproduce. Note that we will
| check that your
| results match the general trends of the original paper and we will
| allow for a
| reasonable margin of error, so you should not worry if metrics do
| not match exactly
| between your reproduction and the paper.
|
| Any experiments that are only introduced in the Appendix are not
| considered core
| contributions of the paper, and so are out of scope. This is the
| case even if the
| experiment is referenced in the main text. HOWEVER, if an experiment
| is described in
| the main body of the text, but some of the details used (e.g.
| details about a prompt
| used, or a long mathematical expression) have been put in the
| appendix, then that
| experiment is considered still in scope. The reason here is that the
| experiment is
| described in the main body of the paper.
|
|                                     RESOURCES
|
| You may or may not be provided with a rubric called rubric.json in
| the /home/paper
| directory to guide your efforts. The rubric.json has a hierarchical
| tree structure,
| at each level breaking down the paper into increasingly fine-grained
| tasks. Note that
| you may be provided with only a subset of the full task hierarchy,
| or even just a
| single task asking you to reproduce the whole paper - in this case,
| you should use
```

your best judgment to break down the tasks you are given into appropriate sub-tasks.

There is also an addendum file at /home/paper/addendum.md which provides additional context needed (if any) to reproduce the paper, as well as potential clarifications about what is not in scope for the reproduction.

You are allowed to browse and use resources you can find online, however you must not look at or use the paper's codebase or other resources listed in blacklist.txt, as this is considered cheating.

SUBMISSION

The output of your work must be a git repository containing all the necessary components to reproduce the results of the paper. Your repository should be located at /home/submission/, and the committed files should not exceed 1GB in size (i.e. you should only commit source code without any heavy artifacts). We will run git clean -fd on the submission directory before grading it, so any untracked files will be lost.

Your submission will be graded based on whether the code necessary for reproducing the paper has been correctly implemented. The code will not be executed during grading.

Finally, please also include a README.md file that describes what you were able to achieve in your reproduction attempt and explains how your codebase relates to various parts of the reproduction.

ADDITIONAL NOTES

- Compute: You have access to a CPU with the appropriate drivers installed.
- Total Runtime: You are expected to work until you have spent 12 hours on this task.
You may iteratively develop your solution within these constraints, but bear in mind that program runtime also counts towards the time limit.
- API keys: You have been granted API keys for services necessary for replicating your paper. The API keys are available to you at /home/agent.env. You can assume that these same API keys will already be set as environment variables in the reproduction environment when we evaluate your submission, so you do not need to set or export them in your submission.
- You have root access in your environment.
- You must use all the time available to you, unless you've reproduced all the core contributions of the paper. Do not stop until you've reproduced them.
- Remember, you must actually reproduce the paper, not just write a

```
| plan for how to
| do so.
| - Each message may perform ONE function call only. Do not make
| multiple function
| calls in a single message.
| - Very Important: Before you write code or reproduce the paper, you
| MUST read all
| content in /home/paper/paper.md and /home/paper/addendum.md!
```

+----- Assistant -----+

```
| <think>
| Let me start by reading the paper and addendum to understand what I
| need to reproduce.
|
| First, I'll read the paper.md file to understand the core
| contributions, and then
| check the addendum.md for any additional context.
| </think>
|
| I'll help you reproduce this research paper. Let me start by reading
| the paper and
| understanding what needs to be replicated.
|
| read_file_chunk(file='/home/paper/paper.md', start_line=1,
| max_lines=50)
```

+----- Tool Output: read_file_chunk -----+

```
| File has 863 total lines. Showing lines 1 to 50.
|
| 1: \title{
| 2: Refined Coreset Selection: Towards Minimal Coreset Size under
| Model Performance Constraints
| 3: }
| 4:
| 5: \author{
| 6: Xiaobo Xia ${ }^1$ Jiale Liu ${ }^2$ Shaokun Zhang ${ }^2$
| Qingyun Wu ${ }^2$
| Hongxin Wei ${ }^3$ Tongliang Liu ${ }^1$
| 7: }
| 8:
| 9: \begin{abstract}
| 10: Coreset selection is powerful in reducing computational costs
| and accelerating
| data processing for deep learning algorithms. It strives to
| identify a small
| subset from large-scale data, so that training only on the
| subset practically
| performs on par with full data. Practitioners regularly desire
| to identify the
| smallest possible coreset in realistic scenes while maintaining
| comparable model
| performance, to minimize costs and maximize acceleration.
| Motivated by this
| desideratum, for the first time, we pose the problem of refined
| coreset
| selection, in which the minimal coreset size under model
| performance constraints
| is explored. Moreover, to address this problem, we propose an
| innovative method
| which maintains optimization priority order over the model
| performance and
```



```

| about,
| in optimization, there is a priority order between them. As analyzed
| above,
| (O1) has a higher priority than (O2), since a smaller coreset size is
| meaningless if the network with such a small coreset does not perform
| satisfactorily. We formalize a general notion of priority order
| rigorously
| as a lexicographic preference (Fishburn, 1975) over two objectives.
| The
| general notion helps specify a clear optimization target across
| multiple
| objectives before optimization and avoids manual post hoc selection.
| 94:
| 95: Specifically, we use the order list
|  $F(\mathbf{m}) = \left[ f_{\{1\}}(\mathbf{m}), \right.$ 
|  $f_{\{2\}}(\mathbf{m}) \left. \right]$  to represent the objectives with a
| lexicographic structure, in which  $f_{\{1\}}$  is the objective with
| higher
| priority and  $f_{\{2\}}$  is the one with lower priority. The
| optimization of
|  $f_{\{2\}}$  only matters on the condition that the more important
| objective
|  $f_{\{1\}}$  is well-optimized. Afterward, our lexicographic bilevel
| coreset
| selection can be formulated as
| 96: \[
| 97: \overrightarrow{\min}_{\mathbf{m} \in \mathcal{M}} F(\mathbf{m}), \text{ s.t. }
| \boldsymbol{\theta}(\mathbf{m}) \in
| \arg \min_{\boldsymbol{\theta} \in \mathcal{L}(\mathbf{m})}
| \boldsymbol{\theta}
| 98: \]
|
| Output truncated (2 additional lines)...

```

```

+----- Assistant -----+
| <think>
| Let me continue reading the paper to understand the full methodology
| and
| then check the addendum.
| </think>
|
| read_file_chunk(file='/home/paper/paper.md', start_line=101,
| max_lines=50)

```

```

+----- Tool Output: read_file_chunk -----+
| File has 863 total lines. Showing lines 101 to 150.
|
| 101: Remark 1. Compared with the trivial solution, i.e., the form of
| weighted combination in (4), our lexicographic bilevel coreset
| selection
| enjoys several advantages. (i). Our method does not need to
| determine the
|
| .
| . (Omit Paper Content) .
| .
| Output truncated (1 additional lines)...

```

```
+----- User -----+
| Reminder: You have read /home/paper/paper.md but haven't read |
| /home/paper/addendum.md yet. Please make sure to read       |
| /home/paper/addendum.md using the read_file_chunk tool, as it |
| contains                                                      |
| important information about the paper you need to replicate. |
+-----+
```

```
+----- Assistant -----+
| <think>                                                      |
| Let me continue reading the paper to understand the full scope, and |
| then                                                         |
| check the addendum file.                                     |
| </think>                                                    |
|                                                             |
| read_file_chunk(file='/home/paper/paper.md', start_line=151, |
| max_lines=50)                                              |
+-----+
```

```
+----- Tool Output: read_file_chunk -----+
| File has 863 total lines. Showing lines 151 to 200.         |
|                                                             |
| 151:                                                         |
| \psi_{t+1}\left[f_{2}\left(\boldsymbol{m}^t\right)-      |
| f_{2}\left(\boldsymbol{m}^{t+1}\right)\right]>\gamma_{2} \text { or } \boldsymbol{m}^t \in |
| \mathcal{M}_{2}^{*}\right] \geq \eta_{2}                    |
| 152: \]                                                     |
| 153: where  $\hat{t}$  represents the earliest time step that the |
| incumbent                                                    |
| mask reaches the optimal region in the objective  $f_1$ , i.e., |
|  $\hat{t}:=\min \left\{t\right. |$  |
|  $\left.\left\{\boldsymbol{m}^t\right\}_{t=0}^T \text { mid}\right. |$  |
| .                                                             |
| . (Omit Paper Content) . |
| .                                                             |
| Output truncated (2 additional lines)...                    |
+-----+
```

```
.
.
. (Omit Paper Reading for 3 Assistant Turns)
.
.
+-----+
```

```
+----- Assistant -----+
| <think>                                                      |
| Let me continue reading the paper to understand all the details, and |
| then                                                         |
| check the addendum file.                                     |
| </think>                                                    |
|                                                             |
| read_file_chunk(file='/home/paper/addendum.md', start_line=1, |
| max_lines=50)                                              |
+-----+
```

```
+----- Tool Output: read_file_chunk -----+
| File has 60 total lines. Showing lines 1 to 50.             |
|                                                             |
| 1: ## Out of scope for reproduction                          |
| 2:                                                           |
| 3: - The experiments (and their results) from Appendix E.5 and E.6 |
+-----+
```

```

| do not
| need to
| 4: be reproduced. They are out of scope.
| 5: - Experiments involving the ImageNet-1k dataset (so, all of
| section 5.4)
| do not
| 6: need to be reproduced. They are out of scope
| 7:
| 8: ## Useful details
| 9:
| 10: ### General useful details
| 11:
| 12: - Submissions should NOT rely on Kaggle for accessing any of the
| datasets. They
| 13: should use other routes (e.g.
| 14:
| [PyTorch] (https://pytorch.org/vision/stable/datasets.html#fashion-
| mnist))
| that
| 15: do not rely on API keys or credentials.
| 16:
| 17: ### Useful details for Figure 1
| 18:
| 19: - In the experiment for Figure 1, a value of lambda of 0.5 is
| used for
| equation
| 20: (4).
| 21: - In the experiment for Figure 1, an arbitrarily random subset
| of MNIST
| is used.
| 22: You may for instance use MNIST-S.
| 23: - For the CNN used, the authors use the 'ConvNet' class from
| Zhou et al,
| 2022.
| 24:
| (https://github.com/x-zho14/Probabilistic-Bilevel-Coreset-
| Selection/blob/mas
| ter/models.py)
| 25: - In the experiment for Figure 1, when running bilevel coreset
| selection, 1000
| 26: iterations of the outer loop are run (i.e. hparam T = 1000).
| 27:
| 28: ### Useful details for Section 5.1
| 29:
| 30: - The CNN described is exactly the same CNN as the one used in
| Figure 1,
| i.e.
| 31: the 'ConvNet' class from Zhou et al, 2022.
| 32:
| (https://github.com/x-zho14/Probabilistic-Bilevel-Coreset-
| Selection/blob/mas
| ter/models.py)
| 33:
| 34: ### Useful details for Section 5.2
| 35:
| 36: - The authors say "For fair comparisons, we reproduce the
| baselines
| based on
| 37: their code repositories." -- This means that the authors
| reimplemented
| the
| 38: baselines themselves, by reading the code of the respective
| papers.
| 39: Reproductions should follow this same procedure.
| 40: - The ResNet18 is trained using an SGD optimizer with a learning

```

```

| rate of
| 0.1,
| 41: momentum of 0.9, and a cosine scheduler.
| 42:
| 43: ### Useful details for Section 5.3
| 44:
| 45: - In section 5.3, the same proxy and target models are used as in
| section 5.2
| 46: for F-MNIST, i.e. a LeNet for both the proxy and target model.
| 47: - For creating the class-imbalanced version of F-MNIST, the
| authors
| leveraged a
| 48: class-imbalanced sampler, adjusting the code from
|
| Output truncated (2 additional lines)...
+-----+

```

```

+----- Assistant -----+
| <think>
| Let me continue reading the addendum and then start implementing the
| paper.
| The paper is about "Refined Coreset Selection" and proposes a method
| called
| LBCS (Lexicographic Bilevel Coreset Selection).
|
| Key points from what I've read so far:
|
| 1 The paper proposes a bilevel optimization approach to coreset
| selection
| 2 The key innovation is lexicographic optimization - first optimize
| model
| performance (f1), then optimize coreset size (f2)
| 3 Main experiments are on Fashion-MNIST, SVHN, CIFAR-10
| 4 ImageNet-1k experiments are out of scope
| 5 Appendix E.5 and E.6 experiments are out of scope
|
| Let me continue reading the addendum to understand all the details.
| </think>
|
| read_file_chunk(file='/home/paper/addendum.md', start_line=51,
| max_lines=50)
+-----+

```

```

+----- Tool Output: read_file_chunk -----+
| File has 60 total lines. Showing lines 51 to 60.
|
| 51: - Note that the imbalance is just injected into the training
| set,
| which does
| 52: not include the test set.
| 53:
| 54: ### Useful details for section 5.4
| 55:
| 56: In section 5.4, the author say "Partial results are from
| previous work".
| These
| 57: are the results for Uniform, EL2N, GraNd, Influential and
| Moderate. The
| authors
| 58: implemented and produced the results for CCS and Probabilistic
| themselves,
| 59: referring to

```



```

| objective
|  $f_{2}$ , while  $f_{1}$  remains in  $\mathcal{M}_{1}^{*}$ .
| 145: Condition 2 (Stable moving condition). At any step  $t \geq$ 
| 0, (i)
| if  $\boldsymbol{m}^t \notin \mathcal{M}_{1}^{*}$ , for all possible
|  $\boldsymbol{m}^t$  in the set
|  $\mathcal{S}_{1} := \left\{ \boldsymbol{m}^t \right.$ 
|  $\left. \text{in } \mathcal{M} \mid f(\boldsymbol{m}^t) \leq$ 
|  $f(\boldsymbol{m}^0) \right\}$ , there exists
|  $\gamma_1 > 0$  and
|  $0 < \eta_1 \leq 1$  so that the algorithm satisfies:
| 146: \[
|
| [Match 2 of 6]
| 543: \section*{A. Details of the Black-box Optimization
| Algorithm}
| 544:
| >>> 545: Technical details. For the black-box optimization of
|  $f_{1}$  and
|  $f_{2}$  in order of priority, we make use of a randomized direct
| search
| algorithm named LexiFlow (Zhang et al., 2023b;c) and make necessary
| modifications to it  $\mathcal{A}^1$ . In RCS, LexiFlow is designed to
| iteratively
| direct the search to the optimal solution based on lexicographic
| comparisons
| over pairs of masks. Technically, at the  $i$ -th iteration, LexiFlow
| maintains an incumbent point that represents the optimal mask up to
| the
|  $i$ -th iteration. The algorithm will sample two new masks near the
| incumbent
| point and update the incumbent point by making lexicographic
| comparisons
| between the incumbent point and sampled masks. During the iterative
| optimization process, LexiFlow will gradually move toward the optimal
| solution. To free the algorithm from local optima and manual
| configuration
| of the step size, LexiFlow includes restart and dynamic step size
| techniques. These techniques are similar to those used in an existing
| randomized direct search-based method (Wu et al., 2021). The details
| are
| provided in Algorithm 2.
| 546:
| 547: Note that in experiments, when updating as did in Algorithm
| 2, the
| value of  $\boldsymbol{m}$  less than -1 becomes -1 and the value
| greater than
| 1 becomes 1 . Then during discretization,  $\boldsymbol{m}$  in
|  $[-1, 0)$  will
| be projected to 0 , and  $\boldsymbol{m}$  in  $[0, 1]$  will be
| projected to 1 .
|
| [Match 3 of 6]
| 545: Technical details. For the black-box optimization of
|  $f_{1}$  and
|  $f_{2}$  in order of priority, we make use of a randomized direct
| search
| algorithm named LexiFlow (Zhang et al., 2023b;c) and make necessary
| modifications to it  $\mathcal{A}^1$ . In RCS, LexiFlow is designed to
| iteratively
| direct the search to the optimal solution based on lexicographic
| comparisons
| over pairs of masks. Technically, at the  $i$ -th iteration, LexiFlow
| maintains an incumbent point that represents the optimal mask up to

```

```

| the
|  $i$ -th iteration. The algorithm will sample two new masks near the
| incumbent
| point and update the incumbent point by making lexicographic
| comparisons
| between the incumbent point and sampled masks. During the iterative
| optimization process, LexiFlow will gradually move toward the optimal
| solution. To free the algorithm from local optima and manual
| configuration
| of the step size, LexiFlow includes restart and dynamic step size
| techniques. These techniques are similar to those used in an existing
| randomized direct search-based method (Wu et al., 2021). The details
| are
| provided in Algorithm 2.
|   546:
| >>> 547: Note that in experiments, when updating as did in Algorithm
| 2, the
| value of  $\mathbf{m}$  less than  $-1$  becomes  $-1$  and the value
| greater than
|  $1$  becomes  $1$ . Then during discretization,  $\mathbf{m}$  in
|  $[-1, 0)$  will
| be projected to  $0$ , and  $\mathbf{m}$  in  $[0, 1]$  will be
| projected to  $1$ .
|   548: ```
|   549: Algorithm 2 Lexicographic Optimization for  $f_1$  and
|    $f_2$ .
|
| [Match 4 of 6]
|   547: Note that in experiments, when updating as did in Algorithm
| 2, the
| value of  $\mathbf{m}$  less than  $-1$  becomes  $-1$  and the value
| greater than
|  $1$  becomes  $1$ . Then during discretization,  $\mathbf{m}$  in
|  $[-1, 0)$  will
| be projected to  $0$ , and  $\mathbf{m}$  in  $[0, 1]$  will be
| projected to  $1$ .
|   548: ```
| >>> 549: Algorithm 2 Lexicographic Optimization for  $f_1$  and
|  $f_2$ .
|   550: Input: Objectives  $F(\cdot)$ , compromise  $\epsilon$ .
|   551: Initialization: Initial mask  $\mathbf{m}_0$ ,
|  $t^{\prime} = r = e = 0$ , and  $\delta = \delta_{\text{init}}$ ;
|
| [Match 5 of 6]
|   570: ```
|   571:
| >>> 572: Practical lexicographic relations. We highlight that the
| notations
| of lexicographic relations in Algorithm 2 (i.e.,
|  $\mathbf{x} \succeq_{\text{lex}} \left( F_{\mathcal{H}} \right) \{ \}$ ,
|  $\vec{\prec}_{\left( F_{\mathcal{H}} \right) \{ \}}$ , and
|  $\preceq_{\left( F_{\mathcal{H}} \right) \{ \}}$ ) are a bit different from
| those in
| the main paper. It is because the optimization with the lexicographic
| relations in Definition 1 relies on the infimums of
|  $f_1(\mathbf{m})$ 
| and  $f_2(\mathbf{m})$ . They are theoretically achievable but
| may be
| inaccessible in practice. Therefore, in experiments, we use practical
| lexicographic relations that are defined with the available minimum
| values
| of objectives.
|   573:
|   574: Specifically, given any two masks  $\mathbf{m}^{\prime}$ 
| and

```

| \mathbf{m} , the practical lexicographic relations
 | $\overrightarrow{=}_{\left(F_{\mathcal{H}}\right)}$,
 | $\vec{\prec}_{\left(F_{\mathcal{H}}\right)}$, and
 | $\underline{\preceq}_{\left(F_{\mathcal{H}}\right)}$ in Algorithm 2
 | are
 | defined as:

 | .
 | .
 | .
 | (Omit Paper Reading for 2 Assistant Turns)
 | .
 | .
.

----- Assistant -----

```
<think>
Let me understand what I need to do:

1 This paper is about "Refined Coreset Selection" (RCS) - finding
the
  smallest coreset size while maintaining model performance
2 The main method is Lexicographic Bilevel Coreset Selection (LBCS)
3 Key experiments are on F-MNIST, SVHN, CIFAR-10 (ImageNet is out
of scope)
4 Main sections to reproduce:
  - Figure 1 (comparison with weighted combination)
  - Section 5.1 (preliminary results on MNIST-S)
  - Section 5.2 (comparison with baselines on F-MNIST, SVHN,
  CIFAR-10)
  - Section 5.3 (robustness experiments with corrupted labels and
  imbalanced data)

Let me start by setting up the repository structure and implementing
the
core algorithm.

The core algorithm (LBCS) involves:

1 Bilevel optimization where:
  - Inner loop: train model on selected coreset
  - Outer loop: update mask using lexicographic optimization
2 Lexicographic optimization that prioritizes f1 (model
performance) over
  f2 (coreset size)
3 Uses a randomized direct search algorithm (LexiFlow)

Let me start implementing:
</think>

Let me start by setting up the submission directory and
understanding the
key components I need to implement.

bash

cd /home && git init submission
```

----- Tool Output: bash -----
Reinitialized existing Git repository in /home/submission/.git/

----- Assistant -----

```

| <think>
| Let me understand the paper better and plan the implementation:
|
| 1 The paper proposes "Lexicographic Bilevel Coreset Selection
| (LBCS)" for
|   the problem of refined coreset selection (RCS)
| 2 Key components:
|   - Bilevel optimization framework where:
|     - Inner loop: Train model on selected coreset
|     - Outer loop: Update mask using lexicographic optimization
|   - Lexicographic optimization prioritizes f1 (model performance)
|     over f2
|     (coreset size)
|     - Uses randomized direct search algorithm (LexiFlow)
| 3 Main experiments to reproduce:
|   - Figure 1: Comparison with weighted combination method
|   - Section 5.1: Preliminary results on MNIST-S
|   - Section 5.2: Comparison with baselines on F-MNIST, SVHN,
|     CIFAR-10
|   - Section 5.3: Robustness experiments (corrupted labels, class
|     imbalance)
|   - Note: ImageNet-1k experiments (5.4) are out of scope
| 4 Baselines to implement:
|   - Uniform sampling
|   - EL2N
|   - GraNd
|   - Influential coreset
|   - Moderate coreset
|   - CCS
|   - Probabilistic coreset (bilevel)
| 5 Key implementation details from addendum:
|   - Use ConvNet from Zhou et al. 2022
|   - For F-MNIST: LeNet
|   - For SVHN: CNN
|   - For CIFAR-10: ResNet-18
|   - Adam optimizer with lr=0.001 for inner loop
|   - epsilon = 0.2, T = 500 for outer loop iterations
|
| Let me start implementing the core algorithm and experiments.
| </think>
|
| Let me start by creating the directory structure and implementing
| the core
| algorithm:
|
| bash
|
| cd /home/submission && mkdir -p src models data experiments results
+-----+
+----- User -----+
| Info: 0:00:42 time elapsed out of 12:00:00. Remember, you only have
| to stop
| working when the time limit has been reached.
|
| Note: Don't forget to git commit regularly!
+-----+
+----- Assistant -----+
| <think>
| Let me start implementing the core components. I need to:
|
| 1 Implement the LBCS algorithm with lexicographic bilevel

```

```

| optimization
| 2 Implement baseline methods (Uniform, EL2N, GraNd, Influential,
| Moderate,
|   CCS, Probabilistic)
| 3 Implement network architectures (LeNet, CNN, ResNet-18)
| 4 Run experiments on F-MNIST, SVHN, CIFAR-10
| 5 Reproduce the main results from the paper
|
| Let me start with the core algorithm implementation.
| </think>
|
| Let me start implementing the core LBCS algorithm and baselines:
|
| bash
|
| cat > /home/submission/src/models.py << 'EOF'
| """Neural network architectures for coreset selection."""
| import torch
| import torch.nn as nn
| import torch.nn.functional as F
|
| class LeNet(nn.Module):
|     """LeNet architecture for Fashion-MNIST."""
|     def __init__(self, num_classes=10):
|         super(LeNet, self).__init__()
|         self.conv1 = nn.Conv2d(1, 6, kernel_size=5, padding=2)
|         self.conv2 = nn.Conv2d(6, 16, kernel_size=5)
|         self.fc1 = nn.Linear(16 * 5 * 5, 120)
|         self.fc2 = nn.Linear(120, 84)
|         self.fc3 = nn.Linear(84, num_classes)
|
|     def forward(self, x):
|         x = F.max_pool2d(F.relu(self.conv1(x)), 2)
|         x = F.max_pool2d(F.relu(self.conv2(x)), 2)
|         x = x.view(x.size(0), -1)
|         x = F.relu(self.fc1(x))
|         x = F.relu(self.fc2(x))
|         x = self.fc3(x)
|         return x
|
| class ConvNet(nn.Module):
|     """ConvNet architecture from Zhou et al. 2022 for MNIST-S
| experiments."""
|     def __init__(self, num_classes=10, input_channels=1):
|         super(ConvNet, self).__init__()
|         self.conv1 = nn.Conv2d(input_channels, 32, kernel_size=3,
| padding=1)
|         self.conv2 = nn.Conv2d(32, 64, kernel_size=3, padding=1)
|         self.dropout1 = nn.Dropout2d(0.25)
|         self.dropout2 = nn.Dropout2d(0.5)
|
|         self.fc1 = nn.Linear(64 * 7 * 7, 128)
|         self.fc2 = nn.Linear(128, num_classes)
|
|     def forward(self, x):
|         x = F.relu(self.conv1(x))
|         x = F.relu(self.conv2(x))
|         x = F.max_pool2d(x, 2)
|         x = self.dropout1(x)
|         x = torch.flatten(x, 1)
|         x = F.relu(self.fc1(x))
|         x = self.dropout2(x)
|         x = self.fc2(x)

```

```

        return x

class SimpleCNN(nn.Module):
    """Simple CNN for SVHN."""
    def __init__(self, num_classes=10):
        super(SimpleCNN, self).__init__()
        self.conv1 = nn.Conv2d(3, 32, kernel_size=3, padding=1)
        self.conv2 = nn.Conv2d(32, 64, kernel_size=3, padding=1)
        self.conv3 = nn.Conv2d(64, 128, kernel_size=3, padding=1)
        self.dropout = nn.Dropout(0.5)
        self.fc1 = nn.Linear(128 * 4 * 4, 256)
        self.fc2 = nn.Linear(256, num_classes)

    def forward(self, x):
        x = F.relu(self.conv1(x))
        x = F.max_pool2d(x, 2)
        x = F.relu(self.conv2(x))
        x = F.max_pool2d(x, 2)
        x = F.relu(self.conv3(x))
        x = F.max_pool2d(x, 2)
        x = x.view(x.size(0), -1)
        x = self.dropout(x)
        x = F.relu(self.fc1(x))
        x = self.dropout(x)
        x = self.fc2(x)
        return x

class ResNet18(nn.Module):
    """ResNet-18 for CIFAR-10."""
    def __init__(self, num_classes=10):
        super(ResNet18, self).__init__()
        # Simplified ResNet-18
        self.conv1 = nn.Conv2d(3, 64, kernel_size=3, padding=1,
                                bias=False)
        self.bn1 = nn.BatchNorm2d(64)
        self.layer1 = self._make_layer(64, 64, 2)
        self.layer2 = self._make_layer(64, 128, 2)
        self.layer3 = self._make_layer(128, 256, 2)
        self.layer4 = self._make_layer(256, 512, 2)
        self.avgpool = nn.AdaptiveAvgPool2d((1, 1))
        self.fc = nn.Linear(512, num_classes)

    def _make_layer(self, in_channels, out_channels, blocks):
        layers = []
        layers.append(nn.Conv2d(in_channels, out_channels,
                                kernel_size=3,
                                padding=1, bias=False))
        layers.append(nn.BatchNorm2d(out_channels))
        layers.append(nn.ReLU(inplace=True))
        for _ in range(1, blocks):
            layers.append(nn.Conv2d(out_channels, out_channels,
                                kernel_size=3, padding=1, bias=False))
            layers.append(nn.BatchNorm2d(out_channels))
            layers.append(nn.ReLU(inplace=True))
        return nn.Sequential(*layers)

    def forward(self, x):
        x = F.relu(self.bn1(self.conv1(x)))
        x = self.layer1(x)
        x = self.layer2(x)
        x = self.layer3(x)
        x = self.layer4(x)
        x = self.avgpool(x)

```

```

|         x = torch.flatten(x, 1)
|         x = self.fc(x)
|         return x
| EOF
+-----+
+----- Assistant -----+
| <think>
| I need to implement the LBCS algorithm and baselines. Let me continue
| building the codebase systematically. I'll implement:
|
| 1 Core LBCS algorithm
| 2 Baseline methods (Uniform, EL2N, GraNd, Influential, Moderate,
| CCS,
| Probabilistic)
| 3 Data loading utilities
| 4 Training utilities
| 5 Experiment scripts
|
| Let me continue with the implementation.
| </think>
|
| bash
|
| cat > /home/submission/src/coreset_selection.py << 'EOF'
| """Coreset selection algorithms."""
| import torch
| import torch.nn as nn
| import numpy as np
| from typing import Tuple, Optional
|
| def lexicographic_compare(f1_new, f2_new, f1_curr, f2_curr,
|                          f1_star, f2_star, epsilon):
|     """
|     Compare two masks using lexicographic relations.
|     Returns True if new mask is better than current mask.
|     """
|     # Check if new mask is in M1* (f1 within epsilon of optimal)
|     in_m1_star = f1_new <= f1_star * (1 + epsilon)
|
|     # Primary objective: f1
|     if f1_new < f1_curr:
|         if in_m1_star:
|             return True
|         # Only update if significantly better
|         return f1_new < f1_curr - 1e-6
|
|     # If f1 is equal or worse, don't update
|     if f1_new > f1_curr:
|         return False
|
|     # f1 is equal, compare f2
|     if in_m1_star and f2_new < f2_curr:
|         return True
|
|     return False
|
| class LBCS:
|     """Lexicographic Bilevel Coreset Selection."""
|
|     def __init__(self, model, train_loader, test_loader, device,

```

```

        epsilon=0.2, T=500, inner_epochs=10, lr=0.001):
    """
    Args:
        model: Neural network model
        train_loader: Training data loader
        test_loader: Test data loader
        device: Device to run on
        epsilon: Performance compromise parameter
        T: Number of outer loop iterations
        inner_epochs: Number of epochs for inner loop training
        lr: Learning rate for inner loop
    """
    self.model = model
    self.train_loader = train_loader
    self.test_loader = test_loader
    self.device = device
    self.epsilon = epsilon
    self.T = T
    self.inner_epochs = inner_epochs
    self.lr = lr

    def train_on_mask(self, mask):
        """Train model on data selected by mask."""
        # Create subset based on mask
        selected_indices = torch.where(mask > 0.5)[0]

        # Create data loader for selected data
        selected_data =
        torch.utils.data.Subset(self.train_loader.dataset,
selected_indices)
        selected_loader = torch.utils.data.DataLoader(
            selected_data, batch_size=128, shuffle=True
        )

        # Train model
        criterion = nn.CrossEntropyLoss()
        optimizer = torch.optim.Adam(self.model.parameters(),
lr=self.lr)

        self.model.train()
        for epoch in range(self.inner_epochs):
            for inputs, labels in selected_loader:
                inputs, labels = inputs.to(self.device),
labels.to(self.device)
                optimizer.zero_grad()
                outputs = self.model(inputs)
                loss = criterion(outputs, labels)
                loss.backward()
                optimizer.step()

        # Evaluate on test set
        self.model.eval()
        correct = 0
        total = 0
        with torch.no_grad():
            for inputs, labels in self.test_loader:
                inputs, labels = inputs.to(self.device),
labels.to(self.device)
                outputs = self.model(inputs)
                _, predicted = outputs.max(1)
                total += labels.size(0)
                correct += predicted.eq(labels).sum().item()

        accuracy = 100. * correct / total

```

```

    # Evaluate on full training set (for f1)
    train_correct = 0
    train_total = 0
    with torch.no_grad():
        for inputs, labels in self.train_loader:
            inputs, labels = inputs.to(self.device),
labels.to(self.device)
            outputs = self.model(inputs)
            _, predicted = outputs.max(1)
            train_total += labels.size(0)
            train_correct += predicted.eq(labels).sum().item()

    train_accuracy = 100. * train_correct / train_total

    return accuracy, train_accuracy, selected_indices.size(0)

def evaluate_mask(self, mask):
    """Evaluate mask by training on it and computing
    objectives."""
    # Reset model
    self.model.apply(lambda m: m.reset_parameters() if hasattr(m,
'reset_parameters') else None)

    # Train on selected data
    test_acc, train_acc, coreset_size = self.train_on_mask(mask)

    # Compute f1 (normalized loss on full training set)
    criterion = nn.CrossEntropyLoss()
    self.model.eval()
    total_loss = 0
    total_samples = 0
    with torch.no_grad():
        for inputs, labels in self.train_loader:
            inputs, labels = inputs.to(self.device),
labels.to(self.device)
            outputs = self.model(inputs)
            loss = criterion(outputs, labels)
            total_loss += loss.item() * inputs.size(0)
            total_samples += inputs.size(0)

    f1 = total_loss / total_samples
    f2 = coreset_size

    return f1, f2, test_acc

def optimize(self, initial_mask):
    """Run lexicographic bilevel coreset selection."""
    n = len(initial_mask)

    # Initialize
    m_current = initial_mask.clone().to(self.device)

    # Evaluate initial mask
    f1_best, f2_best, test_acc_best =
self.evaluate_mask(m_current)
    f1_star = f1_best
    f2_star = f2_best

    print(f"Initial: f1={f1_best:.4f}, f2={f2_best},
test_acc={test_acc_best:.2f}%")

    # Outer loop
    for t in range(self.T):
        # Sample new mask by perturbing current mask
        u = torch.randn(n, device=self.device)

```

```

        u = u / torch.norm(u)

        # Try positive direction
        step_size = 0.1
        m_new = m_current + step_size * u
        m_new = torch.clamp(m_new, 0, 1)

        # Evaluate new mask
        self.model.apply(lambda m: m.reset_parameters() if
        hasattr(m,
        'reset_parameters') else None)
        f1_new, f2_new, test_acc_new = self.evaluate_mask(m_new)

        # Lexicographic comparison
        if lexicographic_compare(f1_new, f2_new, f1_best,
        f2_best,
        f1_star, f2_star, self.epsilon):
            m_current = m_new
            f1_best = f1_new
            f2_best = f2_new
            f1_star = f1_best
            test_acc_best = test_acc_new
            print(f"Iter {t+1}: Improved! f1={f1_best:.4f},
            f2={f2_best}, test_acc={test_acc_best:.2f}%")

        # Occasionally try negative direction
        if t % 3 == 0:
            m_new_neg = m_current - step_size * u
            m_new_neg = torch.clamp(m_new_neg, 0, 1)
            self.model.apply(lambda m: m.reset_parameters() if
            hasattr(m, 'reset_parameters') else None)
            f1_new_neg, f2_new_neg, test_acc_new_neg =
            self.evaluate_mask(m_new_neg)

            if lexicographic_compare(f1_new_neg, f2_new_neg,
            f1_best,
            f2_best, f1_star, f2_star, self.epsilon):
                m_current = m_new_neg
                f1_best = f1_new_neg
                f2_best = f2_new_neg
                f1_star = f1_best
                test_acc_best = test_acc_new_neg
                print(f"Iter {t+1}: Improved (neg)!
                f1={f1_best:.4f},
                f2={f2_best}, test_acc={test_acc_best:.2f}%")

        # Return best mask
        return m_current, f1_best, f2_best, test_acc_best

def uniform_sampling(n, k):
    """Uniform random sampling."""
    mask = torch.zeros(n)
    indices = torch.randperm(n)[:k]
    mask[indices] = 1
    return mask

def greedy_gradient_selection(model, train_loader, device, k):
    """Greedy gradient-based selection (simplified version)."""
    n = len(train_loader.dataset)
    mask = torch.zeros(n)

    # Compute gradient norms for each sample
    model.eval()

```

```

|     gradient_norms = []
|
|     criterion = nn.CrossEntropyLoss()
|     for inputs, labels in train_loader:
|         inputs, labels = inputs.to(device), labels.to(device)
|         model.zero_grad()
|         outputs = model(inputs)
|         loss = criterion(outputs, labels)
|         loss.backward()
|
|         # Compute gradient norm
|         grad_norm = 0
|         for param in model.parameters():
|             if param.grad is not None:
|                 grad_norm += param.grad.data.norm(2).item() ** 2
|         gradient_norms.append(grad_norm ** 0.5)
|
|     # Select top-k samples
|     top_indices = np.argsort(gradient_norms)[-k:]
|     mask[top_indices] = 1
|
|     return torch.tensor(mask)
| EOF
+----- Assistant -----+
| <think>
| I need to commit my progress regularly. Let me commit what I have so
| far and
| continue implementing.
| </think>
|
| bash
|
| cd /home/submission && git add -A && git commit -m "Initial setup:
| models
| and core coreset selection algorithm"
+----- Tool Output: bash -----+
| [master (root-commit) f517e63] Initial setup: models and core coreset
| selection algorithm
| 2 files changed, 344 insertions(+)
| create mode 100644 src/coreset_selection.py
| create mode 100644 src/models.py
+----- Assistant -----+
| <think>
| Good, I've made my first commit. Now I need to continue
| implementing. Let me
| create:
|
| 1 Data loading utilities
| 2 Training utilities
| 3 Baseline methods (EL2N, GraNd, Influential, Moderate, CCS,
| Probabilistic)
| 4 Experiment scripts for the main experiments
|
| Let me continue with data loading and training utilities.
| </think>

```



```

| print('Testing data loading...')
| train_dataset, test_dataset = get_fashion_mnist()
| print(f'Train dataset size: {len(train_dataset)}')
| print(f'Test dataset size: {len(test_dataset)}')
|
| print('Testing model...')
| model = LeNet()
| print(f'Model created successfully')
|
| print('All basic tests passed!')
| "
+-----+
+----- Tool Output: bash -----+
| 0%|          | 0.00/26.4M [00:00<?, ?B/s] 0%|          | 32.8k/26.4M|
| [00:00<02:59, 147kB/s] 0%|          | 65.5k/26.4M [00:00<03:00,
| 146kB/s]
| 0%|          | 131k/26.4M [00:00<02:03, 213kB/s] 1%|          |
| 229k/26.4M
|
| .
|                                     (Omit Experiments)
| .
| [00:00<00:00, 18.6MB/s]
|
| Testing data loading...
| Train dataset size: 60000
| Test dataset size: 10000
| Testing model...
| Model created successfully
| All basic tests passed!
+-----+
|
| .
| .
|                                     (Omit Code Refinement & Experiments for 8 Assistant Turns)
| .
| .
+----- Assistant -----+
| <think>
| The import is failing because the path is not set correctly. Let me
| fix the
| test script.
| </think>
|
| bash
|
| cd /home/submission && python test_quick.py
+-----+
+----- Tool Output: bash -----+
| Traceback (most recent call last):
|   File "/home/submission/test_quick.py", line 47, in <module>
|     mask, f1, f2, acc = lbcn.optimize(initial_mask)
|     ~~~~~^~~~~~
|
|   File "/home/submission/src/coreset_selection.py", line 127, in
|   optimize
|     m_star, f1_history, f2_history = lexicographic_optimization(
|     ~~~~~^~~~~~
|
|   File "/home/submission/src/lexicographic_optimizer.py", line 97, in
|   lexicographic_optimization
|     if practical_lexicographic_preferred(F_plus, F_H[-2], f_hat,
|     epsilon):

```



```
| work.  
| </think>  
|  
| Let me test components individually:  
|  
| bash  
|  
| cd /home/submission && python -c "  
| import sys  
| sys.path.insert(0, 'src')  
| import torch  
| import numpy as np  
| from models import LeNet  
| from data_utils import get_fashion_mnist  
| from baselines import uniform_sampling  
|  
| print('Testing baseline method...')  
| torch.manual_seed(42)  
| np.random.seed(42)  
|  
| train_dataset, test_dataset = get_fashion_mnist()  
| indices = torch.randperm(len(train_dataset))[:500]  
| train_subset = torch.utils.data.Subset(train_dataset, indices)  
|  
| train_loader = torch.utils.data.DataLoader(train_subset,  
| batch_size=64,  
| shuffle=True)  
| test_loader = torch.utils.data.DataLoader(test_dataset,  
| batch_size=64,  
| shuffle=False)  
|  
| model = LeNet()  
|  
| # Test uniform sampling  
| print('Testing uniform sampling...')  
| mask = uniform_sampling(len(train_subset), 100)  
| print(f'Mask sum: {mask.sum().item()}')  
|  
| # Test training on coreset  
| print('Testing training on coreset...')  
| selected_indices = torch.where(mask > 0.5)[0]  
| selected_data = torch.utils.data.Subset(train_subset,  
| selected_indices)  
| selected_loader = torch.utils.data.DataLoader(selected_data,  
| batch_size=64,  
| shuffle=True)  
|  
| criterion = torch.nn.CrossEntropyLoss()  
| optimizer = torch.optim.Adam(model.parameters(), lr=0.001)  
|  
| model.train()  
| for epoch in range(2):  
|     for inputs, labels in selected_loader:  
|         optimizer.zero_grad()  
|         outputs = model(inputs)  
|         loss = criterion(outputs, labels)  
|         loss.backward()  
|         optimizer.step()  
|  
| print('Training completed!')  
|  
| # Test evaluation  
| print('Testing evaluation...')  
| model.eval()  
| correct = 0
```

```

| total = 0
| with torch.no_grad():
|     for inputs, labels in test_loader:
|         outputs = model(inputs)
|         _, predicted = outputs.max(1)
|         total += labels.size(0)
|         correct += predicted.eq(labels).sum().item()
|
| acc = 100. * correct / total
| print(f'Test accuracy: {acc:.2f}%')
| print('All tests passed!')
| " 2>&1
+-----+
+----- Tool Output: bash -----+
| Testing baseline method...
| Testing uniform sampling...
| Mask sum: 100.0
| Testing training on coreset...
| Training completed!
| Testing evaluation...
| Test accuracy: 27.38%
| All tests passed!
+-----+
.
.
.
(Omit Code Refinement & Experiments for 387 Assistant Turns)
.
.
.
(12 Hours Timeout and End the Task)

```

Listing 2: One Bad Case of Claude 45 Sonnet (Thinking) for bridging-data-gaps Task

```

.
.
.
+----- Assistant -----+
| The classifier training is running but taking a while. Let me update |
| the
| ablation study to run faster for demos:
| bash
| cd /home/submission && pkill -f python 2>/dev/null; sleep 2; echo
| "Processes cleaned"
+-----+
(The Main Process is Killed)

```