

# WebSTAR: Scalable Data Synthesis for Computer Use Agents with Step-Level Filtering

Yifei He<sup>\*†1</sup>, Pranit Chawla<sup>†2</sup>, Yaser Souri<sup>2</sup>, Subhojit Som<sup>2</sup>, Xia Song<sup>2</sup>

<sup>1</sup>University of Illinois Urbana-Champaign <sup>2</sup>Microsoft

 Website  Dataset  Code

## Abstract

Computer use agents (CUAs) can operate real-world digital interfaces but remain difficult to train due to the high cost of graphical user interface (GUI) interaction and the scarcity of high-quality trajectory data. Existing datasets rely on human demonstrations, limiting scalability. A natural alternative is to synthesize data from strong CUAs, yet their rollouts are highly noisy, with incorrect or suboptimal actions consisting a large proportion of the steps, making naive imitation ineffective. To tackle this challenge, we introduce a scalable data synthesis pipeline that transforms noisy rollouts into reliable supervision without human annotation. The core idea is step-level filtering, which evaluates actions individually to retain only correct steps, complemented by reasoning augmentation for improved planning. Using this pipeline, we construct WebSTAR, a dataset of 13.3K trajectories and 267K graded, reasoning-rich steps synthesized from OpenAI’s computer-use-preview model. We train Qwen-2.5-VL-Instruct models (7B and 32B) on WebSTAR. On WebVoyager, our 7B model surpasses SoTA open-source CUA model UITARS-1.5-7B by more than 15% with only supervised finetuning. Building on step-level grading, we further create WebSCORE, a dataset of graded step-level actions, and train StepRM, a 7B multimodal process reward model distilled from o4-mini, which matches its grading quality while being far more efficient to deploy at scale. Our results establish step-level filtering as a key principle for scalable CUA training, and we introduce two new datasets (WebSTAR, WebSCORE) along with a lightweight process reward model (StepRM) as practical tools to advance robust and efficient CUAs.

## 1 Introduction

Autonomous agents are increasingly capable of interpreting task instructions, perceiving environments, reasoning and planning, and executing precise actions. Among them, computer use agents (CUAs) (OpenAI, 2025; Anthropic, 2024; Agashe et al., 2025; Qin et al., 2025; Wang et al., 2025a,b) have emerged as a powerful class of models that can operate real-world digital interfaces such as web browsers (Deng et al., 2023; He et al., 2024a; Xue et al., 2025) and desktop environments (Xie et al., 2024; Bonatti et al., 2025; Wu et al., 2025). Despite recent progress, training CUAs remains fundamentally challenging due to two key issues: the high cost of interaction with GUI environments, and the scarcity of high-quality training data.

Live interaction with GUIs typically incurs high infrastructure costs (e.g., virtual machines or cloud environments (Bonatti et al., 2025)), which limits the scale at which interaction data can be collected. As a result, most CUA training is conducted *offline*, where models are trained on pre-collected trajectories using supervised finetuning (SFT) (Xu et al., 2025b; Qin et al., 2025). However, *offline CUA trajectory data is extremely limited*. Existing data sources are either (i) manually labeled (Deng et al., 2023; Xu et al., 2025b), incurring high labor costs, or (ii) curated using general-purpose vision-language models (VLMs) such as GPT-4o (He et al., 2024a; Pahuja et al., 2025), which are not grounded in agentic tasks and often rely on auxiliary information such as accessibility trees that are not available in general deployment settings.

A promising alternative is to synthesize trajectory data using strong CUAs themselves, such as OpenAI CUA (OpenAI, 2025) or Claude Computer Use (Anthropic, 2024). This approach resembles hard distillation in language model training, where a teacher model generates outputs and a student model learns to imitate them. However, unlike tex-

<sup>\*</sup>Work done during an internship at Microsoft. Correspondence email: yifeihe3@illinois.edu.

<sup>†</sup>Equal contribution

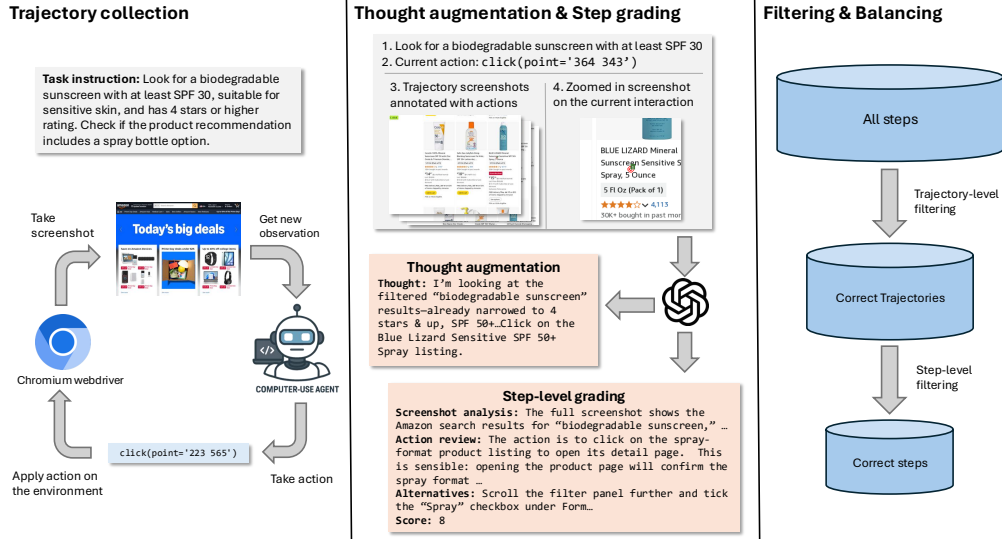


Figure 1: **Overview of our data synthesis pipeline.** Our approach consists of three main stages: **(i) Trajectory collection:** We begin by rolling out a teacher CUA in a Chromium environment, executing actions based on user instructions and capturing observations and screenshots. **(ii) Thought augmentation & Step grading:** For each step, the action and screenshots trajectories are passed through a model to generate an intermediate thought to guide the action. A grading model also receives the same input to evaluate the current step, assigning a score from 0 to 10. **(iii) Step-level Filtering:** We retain only high-scoring steps to ensure that the agent only learns from high-quality actions.

tual LLM outputs, *CUA trajectories are inherently noisy*. We observe that even state-of-the-art CUAs produce incorrect or suboptimal actions in more than half of the steps. Training directly on such noisy rollouts leads to degraded performance in student models, underscoring the need for more principled filtering and augmentation strategies to transform raw rollouts into high-quality training data.

In this work, we aim to answer the research question: *Given a set of task instructions and access to a capable CUA, how to most effectively synthesize high-quality training data to improve agent performance?* We propose a scalable three-step pipeline for generating high-quality CUA training data from noisy teacher rollouts (illustrated in Figure 1):

- 1. Trajectory collection:** We first roll out a teacher CUA on live websites, executing actions based on user instructions and recording screenshots. This process provides the raw but noisy rollouts that serve as the basis for data synthesis.
- 2. Thought augmentation:** Proprietary CUAs do not produce explicit reasoning traces. Thus, inspired by prior work on reasoning-augmented agents (Yao et al., 2023; Xu et al., 2025b; Yang et al., 2024), we insert an intermediate thought before each action to strengthen the agent’s reasoning and planning

abilities. These thoughts provide interpretable, step-by-step guidance that grounds the action in a structured decision-making process.

- 3. Step-level filtering:** Instead of accepting full trajectories, we apply a grading model to evaluate each action individually, retaining only those that are contextually correct and likely to contribute to task success. This ensures that the agent is only trained on high-quality steps, and we empirically find step-level filtering crucial to the agent performance.

Together, these steps allow us to transform noisy rollouts into clean, diverse, and reasoning-rich training data, enabling more reliable offline training of computer use agents.

To evaluate the effectiveness of our data synthesis pipeline, we synthesize data with OpenAI CUA<sup>1</sup> using queries from OpenWebVoyager (He et al., 2024b). From these rollouts, we construct WebSTAR (**WebVoyager Step-Level Trajectories with Augmented Reasoning**), a dataset of 13.3K trajectories and 267K steps enriched with reasoning traces and step-level scores. To complement it, we further derive WebSCORE (**WebVoyager Step-level CORrectness Evaluation**), a dedicated step-level

<sup>1</sup>OpenAI CUA corresponds to the OpenAI computer-use-preview API 2025-03-01 release. We use Azure OpenAI for all API calling.

grading dataset on 50K steps for training process reward models. We train Qwen-2.5-VL-Instruct-7B and 32B (Bai et al., 2025b) on WebSTAR to assess the pipeline’s effectiveness. Our contributions are summarized as follows:

1. **A scalable data synthesis pipeline with step-level filtering:** We propose a three-stage pipeline of trajectory collection, thought augmentation, and step-level filtering, which converts noisy rollouts into clean, reasoning-rich supervision. A key insight is that step-level filtering is essential for effective offline training, consistently outperforming trajectory-level filtering even under simple SFT.
2. **Empirical gains:** On live web tasks such as WebVoyager, our 7B model trained with step-level filtering surpasses the SoTA general-purpose CUA of similar size (UI-TARS-1.5-7B) by more than 15%.
3. **WebSTAR & WebSCORE datasets:** We construct WebSTAR for trajectory training with reasoning traces, and WebSCORE for process reward modeling, providing high-quality resources for future CUA research.
4. **StepRM:** We further train a lightweight 7B process reward model on WebSCORE, achieving grading quality comparable to o4-mini at a fraction of the cost, enabling scalable step-level filtering and paving way for RL-based optimization.

## 2 Preliminaries

A CUA operates by receiving a task instruction and interacting with a graphical user interface to complete the task. At each step  $i$ , the agent observes the current state of the interface as a screenshot  $o_i$ , and executes an action  $a_i$  in response. Given apriori knowledge of task difficulty, a maximum number of steps  $T$  is defined, within which the agent must complete the task. To enhance the agent’s reasoning and planning capabilities, it is common to adopt the ReAct framework (Yao et al., 2023), which augments each step with an intermediate thought  $t_i$  to guide the action  $a_i$ . To ensure broad generalizability, we do not provide the agent with auxiliary inputs such as accessibility trees or specialized annotations like Set-of-Marks used in prior work (He et al., 2024a; Pahuja et al., 2025).

The agent is trained to generate the next thought and action  $(t_n, a_n)$ , conditioned on the task instruc-

Table 1: Action space for computer use agents.

Action	Definition
click(x,y)	Clicks at coordinates $(x, y)$
left_double(x,y)	Double-clicks at $(x, y)$
right_single(x,y)	Right-clicks at $(x, y)$
drag(x1,y1,x2,y2)	Drags from $(x_1, y_1)$ to $(x_2, y_2)$
scroll(x,y,dir)	Scrolls at $(x, y)$ in given direction
type(content)	Types text
hotkey(keys)	Presses hotkey
wait()	Pauses 5s
finished(content)	Complete task with final answer

tion  $\mathcal{I}$ , the sequence of past thoughts and actions  $\{(t_{n-i}, a_{n-i})\}_{i=1}^{n-1}$ , and a truncated window of the most recent  $w$  screenshots  $\{o_{n-i}\}_{i=1}^w$ . Only the image inputs are truncated due to their high token cost, which often dominates the input length in multimodal models. The action space is specified Table 1, where we generally follow common desktop operations defined in UI-TARS (Qin et al., 2025).

During supervised finetuning (SFT), the training objective is to predict the correct thought and action in each step. Formally, given a dataset of trajectories:  $\mathcal{D} = \{(\mathcal{I}, \{o_i, t_i, a_i\}_{i=1}^T)\}$ , the SFT objective minimizes the negative log-likelihood of the ground-truth thought–action pairs:

$$\mathcal{L}_{\text{SFT}}(\theta) = -\mathbb{E}_{\mathcal{D}} \left[ \sum_{n=1}^T \log \pi_{\theta}(t_n, a_n \mid \mathcal{I}, \{t_{n-i}, a_{n-i}\}_{i=1}^{n-1}, \{o_{n-i}\}_{i=1}^w) \right]. \quad (1)$$

This formulation makes clear that *training operates at the step level*, aligning the supervision objective with the recursive nature of CUA execution. Unlike trajectory-level learning, which treats entire rollouts as atomic units, this step-level objective allows the agent to focus on producing locally correct behaviors at each decision point.

## 3 Method

### 3.1 Teacher CUA Rollout

We use queries from the training split of OpenWebVoyager (He et al., 2024b) and collect trajectories by rolling out the OpenAI computer-use-preview API. For each query, we generate 16 rollouts with a maximum horizon of 100 steps, resulting in 21K raw trajectories across 1,326 queries. Following prior works (He et al., 2024a; Xue et al., 2025), we perform first-round trajectory-level filtering to remove failed rollouts by providing GPT-4o (Hurst et al., 2024) with the full action sequence and associated screenshots and

prompting it to judge task success. After this initial grading, 13,338 successful trajectories are retained.

### 3.2 Step-Level Filtering

Previous attempts at data synthesis for CUAs often only retain the trajectories that successfully complete the task (Murty et al., 2024; Pahuja et al., 2025). However, the trajectory-level approach does not align well with the *step-level* nature of CUA training, as in Eq. 1. The training objective focuses on predicting the next thought and action at each individual step, rather than modeling the entire trajectory end-to-end. The key limitation of trajectory-level filtering is that even successful trajectories often contain suboptimal or incorrect intermediate steps. Directly using these full trajectories for SFT exposes the model to flawed decision points, which can bias it toward replicating incorrect behaviors and ultimately degrade downstream performance.

**Grading** To address the limitations of trajectory-level filtering, we propose a step-level filtering strategy that further evaluates each individual action in the successful trajectories. Instead of retaining entire trajectories, our method assigns a quality score to each step based on its expected contribution to task success. This evaluation is performed by a grading model that reasons over the context including the task instruction, the current and past actions, a window of the past  $w$  screenshots annotated with actions following Chen et al. (2024), and a zoomed-in crop of the current screenshot centered on the action target, illustrated in Figure 1. We use o4-mini as our grading model due to its strong reasoning capability and relatively lightweight inference cost. Importantly, o4-mini provides structured intermediate analysis, which enables more consistent judgments across diverse task types. The grading process proceeds in four stages:

1. **Screenshot analysis:** The grading model receives the past  $w$  screenshots along with action annotations, and a zoomed-in crop of the latest screen focused on the relevant web element. This provides contextual grounding for understanding the task state.
2. **Proposed action review:** The model rephrases the candidate action in natural language and judges whether it meaningfully contributes to the goal. Misaligned or unhelpful actions (e.g., misclicks) receive lower scores.
3. **Alternative analysis:** The model proposes

three alternative actions and simulates their likely outcomes. If any alternative action is strictly more effective<sup>2</sup> than the proposed one, the proposed action is penalized.

4. **Evaluation:** Based on correctness (from step 2) and optimality (from step 3), the model assigns an integer score from 0 to 10. A score of 0 indicates an irreversible error or guaranteed failure; 5 denotes a partially correct or suboptimal action; and 10 reflects an unambiguously helpful step with no superior alternative.

An example structure of the response is provided in the step-level grading box in Figure 1, and the detailed prompts describing grading criteria is presented in Appendix F.

While o4-mini offers a practical balance between accuracy and efficiency, it remains a proprietary model. To make grading scalable, we construct WebSCORE, a dedicated step-level dataset with 200K steps where the inputs mirror those provided to o4-mini, and the labels are its outputs. In Section 4.4, we introduce an efficient alternative, StepRM, a 7B model trained on WebSCORE, which achieves grading accuracy on par with o4-mini while being significantly cheaper to deploy, making it suitable for large-scale data synthesis.

**Filtering** To integrate step-level grading into SFT, we retain all steps in the trajectory as contextual information but restrict the loss computation to only high-quality steps. Let  $s_n \in \{0, 1, \dots, 10\}$  denote the score assigned by the grading model to step  $n$ . We define a binary mask  $m_n = \mathbb{1}\{s_n > 5\}$ , where  $m_n = 1$  indicates that the step is judged correct and  $m_n = 0$  otherwise. The training input to the policy  $\pi_\theta$  remains the full sequence of past thoughts, actions, and screenshots, ensuring that the model is exposed to realistic contexts that may include suboptimal decisions. However, *the SFT objective only applies to the correct steps*:

$$\mathcal{L}_{\text{SFT}}(\theta) = - \sum_{n=1}^T m_n \log \pi_\theta \left( t_n, a_n \mid \mathcal{I}, \{t_{n-i}, a_{n-i}\}_{i=1}^{n-1}, \{o_{n-i}\}_{i=1}^w \right).$$

This formulation ensures that incorrect actions remain part of the conditioning context, allowing the agent to learn recovery strategies, while preventing them from directly contributing to the optimization signal. In practice, we find that this

<sup>2</sup>For instance, if the answer is already in view, reporting the answer directly is a strictly better action than scrolling.

Table 2: **Comparison of trajectory-level vs. step-level filtering in SFT.** Step-level filtering trains only on correct steps within successful trajectories, resulting in about half the data volume but consistently higher performance.

Trajectory	Step	# data	Allrecipes	Amazon	Apple	Arxiv	BBCNews	Coursera	ESPN	GitHub	Huggingface	Wolfram	Average
Correct	All	97K	48.9	26.8	32.6	16.3	28.6	45.2	25.0	19.5	19.0	37.0	29.9
Correct	Correct	46K	<b>60.0</b>	<b>34.1</b>	<b>41.9</b>	<b>20.9</b>	<b>47.6</b>	<b>61.9</b>	<b>29.5</b>	<b>31.7</b>	<b>20.9</b>	<b>47.8</b>	<b>39.6</b>

selective masking greatly improves training stability and downstream task performance compared to trajectory-level filtering.

**Less is more for CUA SFT** Surprisingly, in the set of successful trajectories we collect, *fewer than half of the steps are actually correct*. To quantify the impact of removing incorrect intermediate steps, Table 2 compares trajectory-level filtering (training on all steps from correct trajectories) with step-level filtering (training only on the correct steps within those trajectories). Although the step-level filtered dataset is roughly half the size of the trajectory-level one, it consistently yields substantial performance gains across all tasks. This result underscores that removing noisy steps is more beneficial than increasing dataset size for SFT.

### 3.3 Thought Augmentation

The raw trajectories collected from OpenAI CUA consist solely of sequences of observations and actions, *without any explicit reasoning traces*. However, prior work has shown that intermediate reasoning steps are crucial for enhancing agents’ planning and decision-making capabilities, especially in complex, multi-step tasks (Yao et al., 2023).

To address this, we follow previous approaches (Xu et al., 2025b; Yang et al., 2024; Qin et al., 2025) and generate thoughts in a post-hoc manner. Each thought is designed to bridge perception and action, and is structured around three key components: **i) Situation description:** A brief description of the current observation, especially focusing on the region where the action takes place. **ii) Reasoning:** An explanation of the rationale behind the action and how it contributes to completing the overall task. **iii) Actionable instruction:** A natural language command that describes the action in an interpretable way, serving as a textual bridge between thought and execution. The detailed prompt is presented in Appendix F.

**Thought augmentation enhances robustness.** The results in Table 3 show that incorporating thought augmentation yields consistent benefits, especially under the Pass@4 metric. In line with

Table 3: **Comparison of training with and without thought augmentation.** The metrics are Pass@1 (Pass@4). Thought augmentation yields notably higher Pass@4, suggesting more reliable performance.

Model	Task	Correct steps w/o thought	Correct steps w/ thought
7B	WebVoyager	<b>44.8</b> (57.2)	43.3 ( <b>64.2</b> )
	Mind2Web-Live	<b>23.0</b> (28.4)	17.0 ( <b>38.6</b> )
	Online-Mind2Web	<b>22.8</b> (27.7)	17.0 ( <b>29.4</b> )
32B	WebVoyager	<b>48.6</b> (60.6)	45.7 ( <b>65.8</b> )
	Mind2Web-Live	<b>24.8</b> (37.5)	20.4 ( <b>39.7</b> )
	Online-Mind2Web	23.2 (31.3)	<b>23.8</b> ( <b>38.1</b> )

prior findings (Xu et al., 2025b), the Pass@1 scores with and without thought traces remain close, which is expected since Pass@1 is a brittle indicator of reliability: CUA policies are stochastic, and in long-horizon tasks, a single error can derail an entire trajectory. By contrast, Pass@4 provides a more faithful measure of performance, reflecting whether the agent can succeed across multiple sampled rollouts rather than relying on a single attempt.

Across benchmarks and model sizes, Pass@4 consistently improves with thought augmentation, indicating that reasoning traces substantially improve robustness. By explicitly linking perception to action through situation descriptions, rationales, and actionable instructions, thought augmentation helps reduce variance in decision-making. This results in more stable execution, more consistent task completion, and greater reliability when deploying CUAs in real-world environments.

## 4 Experiments

### 4.1 Setup

**Data synthesis** We roll out OpenAI CUA on queries from the OpenWebVoyager (He et al., 2024b) dataset. The original query set comes from the training split of both WebVoyager (He et al., 2024a) and Mind2Web (Deng et al., 2023), totaling 2,111 queries. However, due to the dynamic nature of websites, many tasks become outdated or invalid, caused by issues such as websites access errors, added human verification steps, or updated URLs. After filtering out these queries, we retain

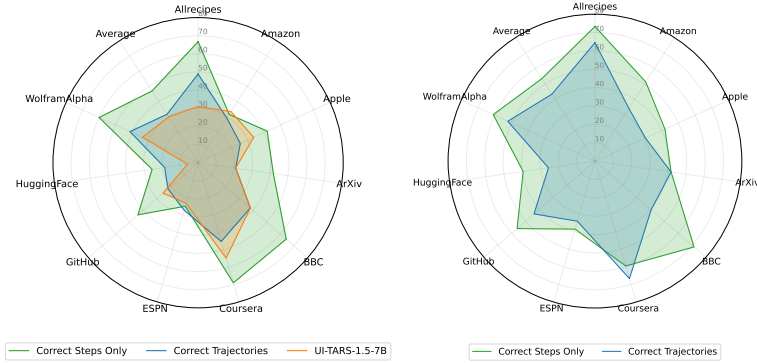


Figure 2: **WebVoyager results for different configurations.** Step-level filtering consistently outperforms trajectory-level filtering and even surpasses the SoTA general-purpose CUA UI-TARS-1.5-7B, despite using only SFT.

1,326 valid ones. Details about those website failures are discussed in Appendix B.

To interact with live websites, we adopt the WebVoyager (He et al., 2024a) codebase. However, the original implementation only supports Set-of-Mark (SoM) prompting, which is incompatible with the more common raw coordinate interaction. Thus, we modify the implementation with Playwright to enable more robust interactions with websites. For each query, we cap the rollout length at 100 steps (i.e.,  $T = 100$ ), terminating the episode at step 100 regardless of whether the task is completed.

**Training** We use Qwen-2.5-VL-7B-Instruct and Qwen-2.5-VL-32B-Instruct (Bai et al., 2025b) as base models. Following common practice (Wang et al., 2024; Xu et al., 2025b), we freeze the vision tower and only train the language components. All experiments are conducted on NVIDIA A100 GPUs with 80 GB memory, using LlamaFactory (Zheng et al., 2024b) for multimodal SFT. We set the sequence length to 8,192 and feed in only the most recent screenshot ( $w = 1$ ) following Xu et al. (2025b), as this configuration yields the most stable training in our setting. We use an initial learning rate of  $1e-5$  with warmup ratio of 0.03 and cosine decay to 0. The global batch size is set to 128. Training on WebSTAR for 2 epochs takes approximately 256 GPU hours for 7B models and 1024 GPU hours for 32B models.

**Evaluation** To test the capabilities of CUA, we deploy three online evaluation benchmarks: WebVoyager (He et al., 2024a), Mind2Web-Live (Pan et al., 2024) and Online Mind2Web (Xue et al., 2025). The online evaluation is particularly challenging due to the everchanging nature of websites, requiring the model to generalize the static data

Table 4: **Performance on Mind2Web suites with Pass@1 (Pass@4).** Step-level filtering substantially improves success rates across both Mind2Web-Live and Online-Mind2Web, with larger gains for 32B models. This confirms that cleaner step-level supervision benefits models of all scales.

Model Size	Tasks	All steps	Correct steps
7B	Mind2Web-Live	15.0 (25.0)	<b>17.0 (38.6)</b>
	Online-Mind2Web	12.0 (23.0)	<b>17.0 (29.4)</b>
32B	Mind2Web-Live	14.5 (26.1)	<b>20.4 (39.7)</b>
	Online-Mind2Web	14.5 (24.8)	<b>23.8 (38.1)</b>

during offline training to dynamically changing environments. Following He et al. (2024a), we provide the full action trajectory and corresponding screenshots of the agent to GPT-4o (Hurst et al., 2024) to automatically judge the task success.

## 4.2 Main Results

The results in Figure 2 and Table 4 highlight the clear advantages of step-level filtering for CUA SFT. In the WebVoyager benchmark, training on only the correct steps in correct trajectories consistently yields stronger performance across almost all domains compared to training on all steps within correct trajectories. Under this setting, our 7B step-level model achieves an average success rate of 47.0%, substantially outperforming the SoTA general-purpose CUA UI-TARS-1.5-7B (30.0%) despite using only SFT. This demonstrates that carefully filtered training data can rival and even surpass the performance of stronger baselines trained with more complex methods. The full numerical results and statistical significance test are presented in Appendix A.

The same pattern holds on the Mind2Web evaluations. Across both Mind2Web-Live and Online-Mind2Web, step-level filtering substantially improves success rates for both 7B and 32B models. The gains are larger for the 32B models, indicating that larger models can capitalize more effectively on clean supervision once noisy steps are removed, further amplifying the benefits of filtering.

Together, these findings demonstrate that step-level filtering is critical for effective offline training of CUAs. By removing incorrect or suboptimal intermediate actions while preserving correct steps from successful trajectories, the resulting training data better aligns with the step-level prediction ob-




Figure 3: **Scaling behavior of step-level vs. trajectory-level filtering.** Step-level filtering consistently outperforms trajectory-level filtering for both 7B and 32B models, with clear gains in low-data settings and stronger improvements as dataset size increases.

jective. This leads to significant improvements in downstream performance, offering a simple yet powerful strategy for advancing the robustness and efficiency of CUA training.

**Performance with data scaling** Figure 3 shows the trend of average performance across the three benchmarks with increasing amount of data under different model sizes. The results highlight that step-level filtering not only improves performance but also scales more effectively with additional data compared to trajectory-level filtering. In the low-data regime (around 25K examples), models trained on correct steps already achieve substantially higher success rates than those trained on full trajectories. As the number of examples increases, the advantage of step-level filtering becomes even more pronounced: both the 7B and 32B models continue to gain steadily, while trajectory-level filtering plateaus early, particularly for the 7B model.

The scalability benefits are more pronounced for the 32B model, showing that larger models can make even better use of clean supervision with more data. By contrast, the trajectory-level baseline falls short in capitalizing on additional data, highlighting the inefficiency of learning from noisy trajectories. Overall, these results show that step-level filtering provides a dual advantage: it yields strong gains when data is scarce and continues to deliver improvements as the dataset grows, making it a robust and scalable strategy for training CUAs.

### 4.3 Ablations

In this section, we further validate the consistency and reliability of o4-mini grading.

Table 5: Confusion matrix for grading.

	o4-mini $\geq 5$	o4-mini $< 5$
Human $\geq 5$	36	16
Human $< 5$	11	37

Table 6: Ablation on score cutoff threshold on 7B model with metrics Pass@1 (Pass@4).

Cutoff	2	4	5	6	8
WebVoyager	29.7 (53.7)	40.2 (64.9)	43.3 (64.2)	41.3 (64.2)	42.5 (62.7)

**Human score validations** To verify that o4-mini’s step-level scores reflect meaningful correctness signals, we manually annotated 100 randomly sampled steps and compared them against the model’s judgments. Using a binary agreement criterion (both assigning scores  $\geq 5$  or  $< 5$ ), we observed a 73% agreement rate. We find that the main source of disagreement occurs when the model penalizes exploratory-but-not-wrong steps more harshly than a human annotator would. This behavior is expected: the grading prompt is deliberately strict to prefer clearly goal-directed actions, and we refined this prompt through extensive iterations to achieve the most consistent downstream performance. The confusion matrix (Table 5) also shows that most disagreements are false negatives. Overall, the agreement analysis supports that the step-level scores are aligned with human intuition and do not exhibit systematic bias that would undermine the filtering process.

**Grading consistency** To assess the stability of the grader, we randomly sampled 230 steps from our training data and queried o4-mini five times per step. We then computed score variance across the repeated samples. On average, scores varied by only 0.81 points, with a median standard deviation of 0.49. The coefficient of variation was similarly low (median  $< 6\%$ ), indicating that o4-mini’s outputs are generally stable and consistent across repeated evaluations.

**Score threshold** In our experiment, we choose 5 as the step-grading cutoff, as a score of 5 is explicitly defined as the boundary case for a partially correct step. To further justify this choice, we conducted an ablation over grading thresholds  $\{2, 4, 6, 8\}$ . For each threshold, we performed filtering and kept the same final dataset size (100K steps) for a fair comparison. The full results are shown in the table below; we report Pass@1 and Pass@4 in parentheses. We can see that threshold = 5 achieves com-

petitive performance. Note that using higher thresholds results in lower initial retention rates, requiring many more rollouts to reach 100K steps, making them significantly less scalable. Overall, these results indicate that threshold = 5 strikes a practical balance between data quality (removing noisy steps) and data efficiency (reasonable retention).

#### 4.4 StepRM

As discussed in Section 3.2, step-level filtering is central to our pipeline: each action in a trajectory is graded individually to decide whether it provides useful supervision. In our primary setup, we relied on o4-mini as the grading model, given its strong reasoning ability and structured outputs. However, o4-mini is proprietary and relatively expensive to query, especially since step-level grading requires evaluating up to  $T = 100$  actions per trajectory, each with screenshot context and action history.

To alleviate this bottleneck, we introduce StepRM, a lightweight multimodal process reward model based on Qwen2.5-VL-7B. StepRM is explicitly trained to replicate the step-level grading behavior of o4-mini, effectively distilling its judgments into a compact, open model. The training is conducted on the WebSCORE dataset, which contains 200K graded steps randomly sampled from the raw rollout trajectories collected in our data synthesis pipeline. To ensure balanced supervision, the dataset contains 50% steps with score  $> 5$  and 50% with score  $\leq 5$ . Each sample includes the same multimodal inputs used in our grading pipeline (the latest annotated screenshot, recent history of actions, and task instruction) and predicts a scalar quality score for the current action. In this way, StepRM functions as a direct substitute for o4-mini in step-level filtering, while being significantly cheaper to deploy at scale. StepRM uses the same hyperparameters and compute setup as our main policy models.

For testing the performance of StepRM, we take a separate subset of 50K steps. As shown in Table 7, both o4-mini and StepRM filtering yield large improvements over trajectory-level filtering, confirming that removing noisy steps is crucial for CUA training. For a fair comparison, both graders start from an independently collected 200K-step dataset distinct from the WebSCORE training set and perform filtering on this shared pool before SFT. Under this controlled setup, StepRM achieves performance on par with o4-mini across multiple benchmarks, while being far more efficient to

Table 7: **Step-level filtering performance from StepRM with Pass@1 (Pass@4)**. Compared with o4-mini as a grading model, StepRM achieves competitive quality while being substantially more efficient, demonstrating that lightweight learned process reward models can replace expensive LLM evaluations for step-level filtering.

Model Size	Tasks	All steps	Correct steps o4-mini	Correct steps StepRM
7B	WebVoyager	27.0 (49.1)	<b>38.4 (63.7)</b>	35.4 (57.9)
	Mind2Web-Live	6.3 (19.4)	<b>18.5 (36.0)</b>	13.2 (28.4)
	Online-Mind2web	9.5 (20.3)	<b>17.3 (31.2)</b>	13.0 (24.4)
32B	WebVoyager	33.4 (58.3)	45.4 (66.3)	<b>45.7 (67.4)</b>
	Mind2Weblive	8.9 (25.0)	<b>20.5 (36.4)</b>	20.0 (38.6)
	OnlineMind2web	13.1 (27.3)	<b>16.5 (29.2)</b>	15.3 (29.1)

deploy at scale.

Beyond serving as a cost-effective substitute for step-level filtering, StepRM also has the potential to act as a critic in reinforcement learning frameworks such as PPO (Schulman et al., 2017). In this setting, StepRM can provide dense, step-level feedback to the agent, helping to overcome two central challenges in CUA training: sparse rewards, where task success is only known at the end of long trajectories, and credit assignment, where it is difficult to determine which intermediate actions led to success or failure. By offering fine-grained step evaluations, StepRM could enable more sample-efficient RL training of CUAs in interactive environments. We leave the integration of StepRM into RL pipelines as an exciting direction for future work.

## 5 Related Works

**Computer use agent datasets** Computer use agents (CUAs) have recently emerged as a powerful class of autonomous systems that can operate directly on graphical user interfaces (OpenAI, 2025; Anthropic, 2024; Agashe et al., 2025; Qin et al., 2025; Wang et al., 2025a,b). Training these agents is particularly challenging because it requires grounding visual input into actionable affordances, reasoning and planning over long horizons, and executing precise actions in dynamic environments. While direct online interaction with GUIs would be the most natural way to acquire these capabilities, it is prohibitively expensive due to the need for virtual machines and complex infrastructure. As a result, most prior work relies on offline training with pre-collected trajectories.

Human-annotated datasets provide valuable supervision but are expensive and limited in scale. UI-TARS (Qin et al., 2025) collected human demon-

Table 8: Comparison with existing datasets.

Datasets	Avg. Steps	# Traj	Thought	Step Grading
MM-Mind2Web	7.7	1K	✗	✗
GUIAct	6.7	2.5K	✗	✗
AgentTrek	12.1	10.4K	✓	✗
WebSTAR (Ours)	22.1	13.3K	✓	✓

strations along with curated self-reflection examples, and UI-TARS-2 (Wang et al., 2025a) went further by requiring annotators to verbalize their reasoning traces during trajectory collection. Others, including MM-Mind2Web (Zheng et al., 2024a) and GUIAct (Chen et al., 2024), provide only short action trajectories without reasoning traces or step-level grading, while AgentTrek (Xu et al., 2025a) includes reasoning but lacks explicit step quality control. Earlier datasets (Liu et al., 2018; Yao et al., 2022; Koh et al., 2024) rely on Set-of-Marks or accessibility trees, which simplifies grounding but deviates from realistic deployment settings where only screenshots are available. These limitations motivate the need for scalable, reasoning-rich, and coordinate-based CUA data.

As summarized in Table 8, WebSTAR is the first to combine scale, long-horizon trajectories, structured reasoning traces, and step-level grading, offering realistic and high-quality supervision without human annotation or auxiliary signals.

**Data filtering** With large-scale but noisy synthetic rollouts, filtering becomes crucial for ensuring that models learn from high-quality supervision. The dominant strategy in CUAs has been trajectory-level filtering, which retains only trajectories that reach successful task completion (Sun et al., 2025). While straightforward, this approach is mismatched to the step-wise nature of CUA training, as successful trajectories often include suboptimal or even incorrect intermediate steps.

Broader works in reinforcement learning have explored step-level filtering, most notably filtered behavior cloning (FBC) (Gulcehre et al., 2023; Singh et al.; Shao et al., 2024), where the model selectively trains on high-quality actions. However, in CUAs, step-level filtering is especially challenging because rewards are inherently sparse: only final success is observed, and Monte Carlo credit assignment is high-variance. Actor-critic methods attempt to address this by learning a critic to provide intermediate feedback (Bai et al., 2024, 2025a), but these methods are often unstable and costly to train (Parisi et al., 2019).

Our work advances this line by introducing a scalable step-level filtering strategy that relies on a separate grading model. Compared with actor-critic methods, our filtering is stable and cost-effective, while still aligning supervision with the step-level nature of CUA training. This makes our method both practically scalable and more faithful to the underlying training objective.

## 6 Conclusions

In this work, we present a scalable data synthesis pipeline for training computer use agents, addressing the limitations of costly human annotation and noisy trajectory-level filtering. Our key insight is that step-level filtering better aligns with the supervision objective, enabling agents to learn from correct actions even within imperfect trajectories. Combined with thought augmentation, which strengthens planning and stability, this approach yields substantial improvements across multiple benchmarks, surpassing strong baselines with only SFT. To support broader research, we introduce two complementary datasets: WebSTAR, a trajectory dataset enriched with reasoning and step-level scores, and WebSCORE, a grading dataset tailored for training process reward models. Building on WebSCORE, we introduced StepRM, a lightweight process reward model that matches o4-mini in grading quality at a lower cost, offering a practical tool for future reinforcement learning based optimization. Together, these contributions establish step-level filtering, reasoning-rich augmentation, and efficient reward modeling as core principles for robust and scalable CUA training.

## Limitations

While our work demonstrates that step-level filtering and thought augmentation substantially improve the scalability and reliability of computer-use agents (CUAs), several limitations remain.

Firstly, the scalability of our pipeline is partly constrained by the quality of the teacher CUA used for data collection. Even state-of-the-art CUAs introduce noise and suboptimal behaviors, which limits the diversity of high-quality steps that can be synthesized. As CUAs themselves improve, we expect this bottleneck to diminish, enabling higher-quality data generation at scale.

Secondly, our approach relies exclusively on supervised finetuning (SFT). Although effective, SFT alone cannot fully address issues such as long-horizon credit assignment or exploration in unseen environments. Reinforcement learning (RL) has the potential to improve these aspects, and our StepRM critic model offers a natural pathway toward integrating step-level filtering with RL in future work.

## References

- Saaket Agashe, Jiuzhou Han, Shuyu Gan, Jiachen Yang, Ang Li, and Xin Eric Wang. 2025. Agent s: An open agentic framework that uses computers like a human. In *The Thirteenth International Conference on Learning Representations*.
- Anthropic. 2024. Developing a computer use model. <https://www.anthropic.com/news/developing-computer-use>.
- Hao Bai, Yifei Zhou, Li Erran Li, Sergey Levine, and Aviral Kumar. 2025a. Digi-q: Learning VLM q-value functions for training device-control agents. In *The Thirteenth International Conference on Learning Representations*.
- Hao Bai, Yifei Zhou, Jiayi Pan, Mert Cemri, Alane Suhr, Sergey Levine, and Aviral Kumar. 2024. Digi-girl: Training in-the-wild device-control agents with autonomous reinforcement learning. *Advances in Neural Information Processing Systems*, 37:12461–12495.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Siboz Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuezhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Rogério Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Buckner, and 1 others. 2025. Windows agent arena: Evaluating multi-modal os agents at scale. In *Forty-second International Conference on Machine Learning*.
- Wentong Chen, Junbo Cui, Jinyi Hu, Yujia Qin, Junjie Fang, Yue Zhao, Chongyi Wang, Jun Liu, Guirong Chen, Yupeng Huo, and 1 others. 2024. Guicourse: From general vision language models to versatile gui agents. *arXiv preprint arXiv:2406.11317*.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36:28091–28114.
- Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, and 1 others. 2023. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*.
- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. 2024a. Webvoyager: Building an end-to-end web agent with large multimodal models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6864–6890.
- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Hongming Zhang, Tianqing Fang, Zhenzhong Lan, and Dong Yu. 2024b. Openwebvoyager: Building multimodal web agents via iterative real-world exploration, feedback and optimization. *arXiv preprint arXiv:2410.19609*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Russ Salakhutdinov, and Daniel Fried. 2024. Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 881–905.
- Evan Zheran Liu, Kelvin Guu, Panupong Pasupat, Tianlin Shi, and Percy Liang. 2018. Reinforcement learning on web interfaces using workflow-guided exploration. In *International Conference on Learning Representations (ICLR)*.
- Shikhar Murty, Christopher D Manning, Peter Shaw, Mandar Joshi, and Kenton Lee. 2024. Bagel: Bootstrapping agents by guiding exploration with language. In *International Conference on Machine Learning*, pages 36894–36910. PMLR.

- OpenAI. 2025. Introducing operator. <https://openai.com/index/introducing-operator/>.
- Vardaan Pahuja, Yadong Lu, Corby Rosset, Boyu Gou, Arindam Mitra, Spencer Whitehead, Yu Su, and Ahmed Awadallah. 2025. Explorer: Scaling exploration-driven web trajectory synthesis for multi-modal web agents. *arXiv preprint arXiv:2502.11357*.
- Yichen Pan, Dehan Kong, Sida Zhou, Cheng Cui, Yifei Leng, Bing Jiang, Hangyu Liu, Yanyi Shang, Shuyan Zhou, Tongshuang Wu, and 1 others. 2024. Web-canvas: Benchmarking web agents in online environments. *arXiv preprint arXiv:2406.12373*.
- Simone Parisi, Voot Tangkaratt, Jan Peters, and Mohammad Emtiyaz Khan. 2019. Td-regularized actor-critic methods. *Machine Learning*, 108(8):1467–1501.
- Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, and 1 others. 2025. Uitars: Pioneering automated gui interaction with native agents. *arXiv preprint arXiv:2501.12326*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Xavier Garcia, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, and 1 others. Beyond human data: Scaling self-training for problem-solving with language models. *Transactions on Machine Learning Research*.
- Qiushi Sun, Kanzhi Cheng, Zichen Ding, Chuanyang Jin, Yan Wang, Fangzhi Xu, Zhenyu Wu, Chengyou Jia, Liheng Chen, Zhoumianze Liu, Ben Kao, Guohao Li, Junxian He, Yu Qiao, and Zhiyong Wu. 2025. OS-genesis: Automating GUI agent trajectory construction via reverse task synthesis. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5555–5579, Vienna, Austria. Association for Computational Linguistics.
- Haoming Wang, Haoyang Zou, Huatong Song, Jiazhan Feng, Junjie Fang, Junting Lu, Longxiang Liu, Qinyu Luo, Shihao Liang, Shijue Huang, and 1 others. 2025a. Ui-tars-2 technical report: Advancing gui agent with multi-turn reinforcement learning. *arXiv preprint arXiv:2509.02544*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Xinyuan Wang, Bowen Wang, Dunjie Lu, Junlin Yang, Tianbao Xie, Junli Wang, Jiaqi Deng, Xiaole Guo, Yiheng Xu, Chen Henry Wu, and 1 others. 2025b. Opencua: Open foundations for computer-use agents. *arXiv preprint arXiv:2508.09123*.
- Qianhui Wu, Kanzhi Cheng, Rui Yang, Chaoyun Zhang, Jianwei Yang, Huiqiang Jiang, Jian Mu, Baolin Peng, Bo Qiao, Reuben Tan, and 1 others. 2025. Gui-actor: Coordinate-free visual grounding for gui agents. *arXiv preprint arXiv:2506.03143*.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh J Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, and 1 others. 2024. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *Advances in Neural Information Processing Systems*, 37:52040–52094.
- Yiheng Xu, Dunjie Lu, Zhennan Shen, Junli Wang, Zekun Wang, Yuchen Mao, Caiming Xiong, and Tao Yu. 2025a. Agenttrek: Agent trajectory synthesis via guiding replay with web tutorials. In *The Thirteenth International Conference on Learning Representations*.
- Yiheng Xu, Zekun Wang, Junli Wang, Dunjie Lu, Tianbao Xie, Amrita Saha, Doyen Sahoo, Tao Yu, and Caiming Xiong. 2025b. Aguis: Unified pure vision agents for autonomous gui interaction. In *Forty-second International Conference on Machine Learning*.
- Tianci Xue, Weijian Qi, Tianneng Shi, Chan Hee Song, Boyu Gou, Dawn Song, Huan Sun, and Yu Su. 2025. An illusion of progress? assessing the current state of web agents. *arXiv preprint arXiv:2504.01382*.
- Zonghan Yang, Peng Li, Ming Yan, Ji Zhang, Fei Huang, and Yang Liu. 2024. React meets actre: When language agents enjoy training data autonomy. *arXiv preprint arXiv:2403.14589*.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- Boyuan Zheng, Boyu Gou, Jihyung Kil, Huan Sun, and Yu Su. 2024a. Gpt-4v (ision) is a generalist web agent, if grounded. In *International Conference on Machine Learning*, pages 61349–61385. PMLR.

Yaowei Zheng, Richong Zhang, Junhao Zhang, YeYanhan YeYanhan, and Zheyang Luo. 2024b. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 400–410.

## A More Results

**Full numerical results** We present the full numerical results of WebVoyager in Table 9 and Table 10. To reduce the variance in the evaluation process, we further report the Pass@4 metric beyond Pass@1. Under both model sizes, step-level filtering significantly outperforms trajectory level filtering, validating the effectiveness of our filtering strategy. OAI CUA still obtains the strongest performance, but our 32B model results demonstrated in Table 10 already closes the gap significantly. With the effective scaling of both data (shown in Figure 3 and model, we believe our data synthesis pipeline can yield a recipe that approach OAI CUA performance.

There are domains where smaller models fail to close the gap with the teacher CUA, such as HuggingFace, and our analysis shows this domain is more challenging than others. The queries here typically require deep technical knowledge, multi-hop navigation, and precise extraction of information from code-heavy documentation (e.g., implementation details of TRL modules or tokenizer parameters). The teacher model (OAI CUA) is far more sample-efficient in locating the correct documentation sections, whereas the student often spends many turns navigating intermediate pages. A key observation is that failure is dominated by running out of turns, not by incorrect reasoning. Across HuggingFace queries, the student exhausts the 100-step limit in 65% of tasks, compared to 28% for OAI CUA.

**Statistical significance** We include the 95% confidence interval for our main results on the 7B models in Table 11. To compute error bounds, we run each query 4 times. These intervals show that the improvements we report, such as the gains from step-level filtering and the performance differences among compared methods, remain statistically significant under reasonable uncertainty estimates. In addition, in the paper, we already report Pass@4, which is a more stable and reliable metric for long-horizon CUA tasks and further reduces variance compared to Pass@1.

**Score distribution** The CDF in Figure 4 illustrates the cumulative distribution of step-level scores. The distribution is computed only over the successful trajectories, as our final goal is to only retain the correct steps in the successful trajectories. Even within those successful trajectories,

Domain	All steps	Correct steps	UI-TARS-1.5-7B	OAI CUA
Allrecipes	48.9 (80.0)	<b>66.7 (86.7)</b>	30.9 (56.8)	78.2 (91.1)
Amazon	29.3 (46.3)	31.7 (56.1)	<b>33.5 (65.9)</b>	54.3 (92.7)
Apple	25.6 (46.5)	<b>41.9 (55.8)</b>	33.7 ( <b>60.5</b> )	68.6 (79.1)
Arxiv	20.9 (34.9)	<b>41.9 (58.1)</b>	20.9 (39.5)	71.5 (90.7)
BBCNews	38.1 (61.9)	<b>64.3 (73.8)</b>	38.1 ( <b>76.2</b> )	67.1 (81.0)
Coursera	45.2 (66.7)	<b>69.0 (78.6)</b>	54.8 (76.2)	79.8 (92.9)
ESPN	<b>27.3 (47.7)</b>	25.0 (50.0)	23.3 (43.2)	56.8 (75.0)
GitHub	21.9 (75.6)	<b>43.9 (82.9)</b>	25.6 (48.8)	66.5 (85.4)
Huggingface	18.6 (25.6)	<b>25.6 (34.9)</b>	5.8 (11.6)	61.6 (79.1)
Wolfram	41.3 ( <b>65.2</b> )	<b>60.0 (65.2)</b>	33.7 (47.8)	85.3 (93.5)
<b>Average</b>	31.9 (55.1)	<b>47.0 (64.2)</b>	30.0 (52.4)	69.2 (86.0)

Table 9: **Performance comparison on 7B models.** The reported metrics are Pass@1 and Pass@4 (in parentheses). The bold numbers show the best performing metric among the 7B models, where OAI CUA serves as the oracle. Step-level filtering (“Correct steps”) substantially improves over trajectory-level filtering (“All steps”), and further outperforms UI-TARS-1.5-7B across domains. OpenAI CUA (OAI CUA) remains strongest, but our method narrows the gap significantly.

Domain	All steps	Correct steps
Allrecipes	64.4 (84.4)	<b>73.3 (91.1)</b>
Amazon	34.1 (41.5)	<b>48.8 (51.2)</b>
Apple	30.2 (44.2)	<b>41.9 (62.8)</b>
Arxiv	41.9 (44.2)	<b>41.9 (53.5)</b>
BBCNews	40.5 (61.9)	<b>71.4 (76.2)</b>
Coursera	<b>66.7 (76.2)</b>	59.5 ( <b>85.7</b> )
ESPN	34.1 (43.2)	<b>38.6 (50.0)</b>
GitHub	43.9 (73.2)	<b>56.1 (75.6)</b>
Huggingface	25.6 (30.2)	<b>39.5 (46.5)</b>
Wolfram	52.2 ( <b>69.6</b> )	<b>60.9 (67.4)</b>
<b>Average</b>	43.5 (56.9)	<b>53.5 (65.8)</b>

Table 10: **Performance comparison for 32B models.** Step-level filtering (“Correct steps”) consistently improves over trajectory-level filtering (“All steps”), yielding large gains across nearly all domains.




Figure 4: **Step-level score distribution from o4-mini.** The CDF shows the proportion of actions with scores above each threshold, highlighting that a large share of steps are graded as suboptimal.

a majority of actions receive low scores: over half fall below 5. This confirms that trajectory-level filtering, which accepts whole rollouts as long as the final task succeeds, inevitably includes many

flawed intermediate steps. Such noisy supervision can mislead agents during training. Step-level filtering directly addresses this issue by isolating the small fraction of high-quality actions (scores close to 10) while discarding the majority of low-quality ones. This reinforces our key insight that filtering at the step granularity is essential for constructing cleaner and more reliable training data for CUAs.

## B Evaluation Details

In Table 12, we list the websites where we have access issues caused by anti-scraping or anti-robot measures as well as bugs within websites. Due to the invalidity, we exclude tasks involving interactions with those websites during training and testing. We also want to clarify that this table only represents the error happening at the time of our testing (August 2025). It is possible that certain tasks are fixed after this time frame, and become valid again. Similarly, we do not include google related tasks in the WebVoyager evaluation, as they have strict human verification requirements.

## C Data Examples

An example datapoint in WebSTAR is shown in Figure 5. Each datapoint contains components: the action from the agent, the thought augmented by o4-mini, the step-level score provided by o4-mini, and the rationale for giving this score. Each step also comes with a screenshot that captures the observation of the current state. Since we do not use any auxiliary source of information such as set of marks, there is no annotation on the screenshot. We

Domain	All steps	Correct steps	UI-TARS-1.5-7B	OAI CUA
Allrecipes	49.4 ± 9.9	<b>59.4 ± 9.9</b>	30.9 ± 10.0	78.2 ± 9.9
Amazon	25.0 ± 9.7	<b>34.8 ± 11.6</b>	33.5 ± 9.9	54.3 ± 7.7
Apple	26.7 ± 10.9	<b>36.0 ± 11.7</b>	33.7 ± 10.3	68.6 ± 12.5
Arxiv	24.4 ± 11.2	<b>34.3 ± 11.5</b>	20.9 ± 9.7	71.5 ± 9.6
BBCNews	29.8 ± 9.1	<b>45.8 ± 10.6</b>	38.1 ± 9.2	67.1 ± 11.6
Coursera	48.2 ± 12.6	<b>64.9 ± 12.0</b>	54.8 ± 11.0	79.8 ± 9.1
ESPN	19.9 ± 7.8	<b>31.2 ± 10.9</b>	23.3 ± 9.1	56.8 ± 11.8
GitHub	40.9 ± 10.5	<b>48.8 ± 10.5</b>	25.6 ± 9.3	66.5 ± 10.9
Huggingface	12.2 ± 7.5	<b>20.3 ± 10.0</b>	5.8 ± 5.5	61.6 ± 11.7
Wolfram	47.8 ± 12.1	<b>57.1 ± 13.2</b>	33.7 ± 12.0	85.3 ± 8.1
<b>WebVoyager</b>	32.6 ± 3.5	<b>43.4 ± 3.8</b>	30.0 ± 3.3	69.2 ± 3.4
<b>Mind2Web-Live</b>	14.8 ± 6.2	17.2 ± 4.6	<b>18.1 ± 6.0</b>	50.3 ± 8.5
<b>Online-Mind2Web</b>	12.3 ± 3.9	<b>17.0 ± 5.7</b>	14.6 ± 4.5	48.5 ± 6.3

Table 11: **Performance with 95% confidence intervals on 7B models.** Values are Pass@1 in percentages.

Table 12: Websites with access problems in Mind2Web queries.

Error	Websites
Human verification	bbb.org, redfin.com, extraspace.com, seatgeek.com, google.com/shopping, sec.gov, expedia.com, discogs.com
Access denied	americashealthrankings.org, adoptapet.com, reddit.com, qatarairways.com, fedex.com, macys.com, gamestop.com, disney.com, kbb.com, accuweather.com
Website issue	nba.com, target.com, phys.org, ups.com, kfc.com, cars.com, bestbuy.com, apartments.com, kohls.com

have attached more data samples with full trajectories in the attached data folders.

## D Potential Risks

This paper presents work whose goal is to advance the field of NLP. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

## E Licenses

Both WebVoyager and OpenWebVoyager are under Apache-2.0 license. Both Online-Mind2Web Mind2Web-Live are under cc-by-4.0 license.

## F Prompts for Thought Augmentation and Grading

We attach the full prompts we use for thought augmentation and grading as follows. We use o4-mini to generate the answers for both purposes.



IN THE KITCHEN

### The 15 Best Summer Recipes for Your 9x13 Dish

From dreamy lemon bars to an easy Mexican casserole, it's worth turning on your

### The Latest

IN THE KITCHEN | 4 MINUTES AGO  
**PSA: Grocery Stores Cover Potatoes at Night—Here's Why You Should Too**

GROCERY | 1 HOUR AGO  
**M&M's Just Launched a Brand-New Flavor Fans Are Already Calling Their 'Favorite Snack'**

GROCERY | 2 HOURS AGO  
**Aldi Brought Back the Snack Customers Finish in the Parking Lot**

KITCHEN TOOL & UTENSIL REVIEWS

```
"conversations": [
  {
    "role": "user",
    "content": "Now given a task: Locate a chicken stir-fry recipe on Allrecipes with over 50 reviews, hinting that it can be cooked in 30 minutes or less. Specify the type of vegetables included in the recipe. Please interact with https://www.allrecipes.com/ and get the answer. Once you have found the answer, please reply by providing the answer as finished(content='<your_answer>') and exit the task.<image>Please analyze the attached screenshot and give the Thought and Action."
  },
  {
    "from": "assistant",
    "value": "Thought: I see the Allrecipes homepage with the logo at top left, a prominent 'Find a recipe or ingredient' search box in the header, and featured content beneath, but no stir-fry recipes are visible. To locate a chicken stir-fry dish with the required criteria, I need to activate the search bar and enter my query. I'll click on the search field at the top of the page.\nAction: click(point='<point>461 33</point>')",
    "score": 8,
    "judge": "1. Latest Screenshot Analysis \n - We see the Allrecipes homepage with the top navigation bar containing the 'Find a recipe or ingredient' search input in the center. On the right are 'Log In' and a heart icon. Below is a banner and a featured article image, followed by recipe highlights. \n - A green 'click' label appears in the top-left, denoting the assistant intends to click. A red circle is centered directly on the search input box. \n - A zoomed inset confirms the red circle is neatly over the text field where you type your query. \n\n2. Success and Rejection Criteria \n Success criteria: \n - We search for 'chicken stir fry' (or similar) on Allrecipes. \n - We locate a recipe titled or clearly labeled 'Chicken Stir-Fry' (or close variant). \n - That recipe shows at least 50 reviews. \n - The recipe states a total time or cook time of 30 minutes or less. \n - We identify and report the specific vegetables listed in the ingredients. \n\n Rejection criteria: \n - Failing to find a recipe with ≥50 reviews. \n - Selecting a recipe with >30 minutes total time. \n - Not reporting the vegetables. \n\n3. Progress Analysis \n - Step 1 (this action): Assistant clicked the search bar. No criteria met yet, but this is the correct preparatory step. \n - Remaining: type query, submit search, pick suitable recipe, verify reviews/time, extract veggies. \n\n4. Proposed Action Review \n a. The assistant is focusing the search input box so it can type the search terms. \n b. This is exactly the right next move: we need to search before anything else. \n c. The red circle is correctly centered on the search field. \n d. This is a fully correct action and advances the task. Score should be >5. \n\n5. Simulation of Outcomes \n - Best-case (very likely): The search box gains focus, allowing typing of 'chicken stir fry.' \n - Worst-case (very unlikely): The click misses and no field is focused, but this can be retried. No harm done. \n\n6. Alternatives Analysis \n a. An alternative might be scrolling down to find a 'Stir-Fry' category, but that is roundabout and slower. \n b. The search box is visible and the most direct approach. \n c. Typing immediately after focusing is faster than scanning page content. \n d. The proposed action is strictly better than alternatives. \n\n7. Evaluation \n Clicking the search input is the essential next step to find the recipe. It directly supports the user's goal with no better immediate options. \n\n8. Expected Value \n\nExpected value: 8"
```

Figure 5: Example step for Allrecipe tasks.

## Thought augmentation prompt

You are given the action and thought of the previous several steps of a web browsing agent, the current screenshot annotated with the current action, and an action dictionary describing the specific parameters of the current action. Optionally, you will be given previous screenshots about previous states. The goal is to output the detailed thought process that leads to the current action.

The action space of the assistant includes:

- click(x, y) where (x, y) are the coordinates of the element to click on.
- scroll(x, y, scroll\_x, scroll\_y) where (x, y) are the coordinates of the element to scroll on and (scroll\_x, scroll\_y) are the scroll amounts in pixels.
- keypress(keys) where keys is a string of keys to press.
- type(text) where text is a string to type. To type in a searchbox, the assistant needs to first click on it, then type.
- wait
- screenshot
- final\_answer(answer) where answer is the final answer to the user task.

Your response should contain the following components:

1. **Situation Description**: Describe your observation of the screenshot in detail. Do not only focus on the regions where the action takes place. Rather, identify key areas and elements that contribute to the decision-making process, such as relevant text, images, or layout features that inform the next steps. Then, arrive at the decision to interact with certain area of interest, and relate it to the goal to achieve.
2. **Reasoning Alignment**: Ensure your reasoning aligns with the current action and how it contributes to achieving the goal, but avoid using the current action or the annotation on the screenshot as reasoning support, as they represent hindsight rather than predictive insight. Think about the previous steps and how they lead to the current action. Do not output completely equivalent reasoning to the current action even if it is the same action as previous actions. Rather, think about why this action is repeated, and how it is different from previous attempts.
3. **Actionable Instruction**: Conclude with a clear, actionable instruction in one sentence, but no need to use any specific format. Make sure that the instruction matches the provided action.

Important notes:

1. Aim to reason through the task as if solving it, rather than simply reflecting on the outcome. Use the first-person perspective to represent the annotator's thought process.
  2. The actions are not necessarily optimal or correct. **STICK TO THE GIVEN CURRENT ACTION AND DO NOT COME UP WITH YOUR OWN ACTION** or impose your idea about what is the correct action in the current step. Instead, focus on finding rationale and reasoning for the given action.
  3. The screenshots are annotated with the actions. On the top left corner, there is the action label, such as click, scroll, wait, type etc. For clicking, there is a red target dot with green label, indicating the clicking position. Do not confuse it with other red elements in the screenshot. For scroll, there is a red target dot for the scrolling centroid, and a red arrow points to the scrolling direction. For drag, there is a red arrow pointing towards the direction of drag, with the dragging start point annotated by a green label.
  4. The assistant can only perform one action in each step.
  5. If the assistant is on the first step, provide a brief overview of the task and decompose it into smaller steps.
  6. The effect of the current action is NOT reflected in the latest screenshot, as the assistant has not executed the action yet. The latest screenshot is just a snapshot of the current state before the action is executed.
  7. For final answer, the action is "finished", but the answer shown on the annotated screenshot may be truncated, so focus on the content of the answer provided in text. The agent will finish the task after this step, so do not plan for any further steps or actions.
- There is no need to explicitly state the titles of the components, just write them in one paragraph.

## Grading prompt

You are a critic helping evaluate the next action of a computer use agent. Your goal is to judge the expected value of the proposed next action based on:

1. Whether it meaningfully contributes to successful completion of the user's task.
2. Whether it is the most promising possible next action from the assistant's available action space – i.e., a "no-regret" choice.

You will be given:

- A USER\_TASK: what the user wants the assistant to do.
- A sequence of prior screenshots (including any zoomed-in crops), each corresponding to one earlier assistant action.
- A single PROPOSED\_NEXT\_ASSISTANT\_ACTION.
- The latest full-screen screenshot and zoomed-in image (if any), annotated with red/green visual guides for where the proposed action is targeting.

Assistant action space includes: - 'click(x, y)': click at coordinates (x, y)

- 'scroll(x, y, scroll\_x, scroll\_y)': scroll at (x, y) by the pixel amounts (scroll\_x, scroll\_y)
- 'keypress(keys)': press keys like "Enter", "Ctrl+A", etc.
- 'type(text)': type a string (must click on an input first)
- 'wait': wait for page to change or update
- 'screenshot': take a new screenshot
- 'final\_answer(answer)': output the final answer to the user (no screenshot will be given with this)

SCREENSHOT ANNOTATION CONVENTIONS: - The top-left green label (e.g. "click") is the **type of action**, not where it occurs.

- The **red circle** is the exact target position for a click, scroll, or drag – the assistant's mouse cursor will land there. It must be on the correct element.
- For **scroll actions**, a red arrow shows scroll direction, and red circle marks scroll origin.
- For **drag actions**, red circle marks the start point, red arrow marks direction.
- For sliders, the specific knob dragged and direction both matter. Always examine drag annotations carefully.

IMPORTANT: The assistant has **not yet executed** the proposed action. You must judge its value **before it runs**, based on what is visible on screen.

Your task is to follow the 8 steps below, output your analysis for each of the steps, then return an integer score in the format 'Expected value: <int>', where:

- 0 = guaranteed task failure or irreversible error
- 10 = guaranteed task success and no better alternative action exists
- 5 = a borderline step that is either only partially correct or may be outperformed by a better next action

1. **Latest Screenshot Analysis** - Then analyze the **latest full screenshot**: describe relevant UI elements visible, current screen state, and the annotation overlays (red circles/arrows, green labels). - Describe any **zoomed-in** image(s). For each, examine where the red circle is at: if it's not directly on an interactive element, say so explicitly (e.g. "not centered on any element"). If a web element (such as a search button) is small, it could be partially obscured by the red circle, pay close attention to such details. - If red annotations suggest a different intent than the textual action description, update the interpretation accordingly.

2. **Success and Rejection Criteria** - Break down the USER\_TASK into **specific, verifiable success conditions**. - Define what would count as incorrect or incomplete (rejection criteria).

3. **Progress Analysis** - Go through EACH screenshot and earlier action. - For each step, infer what the assistant likely did and how the screen changed. - Mark which success criteria have already been completed, and which remain.

4. **Proposed Action Review** a. Rephrase in plain English what the assistant is trying to do. b. Judge whether this makes sense given the current context and progress. c. If the red circle is off-target or the action does not help, state that the score should be  $\leq 5$ . d. If the action is fully correct and contributes meaningfully to task completion or fixes a past mistake, state that the score should be  $> 5$ . Specifically, if the current state is already wrong, and the assistant fixes a previous mistake, the score should be higher than 5. **EXPLICITLY** think about whether the action is fixing a previous mistake. Notes for action judgement: - For clicking, it does not need to be exactly centered on the element, as long as it is reasonably close. - For dragging on sliders, if the knob and direction are correct but the dragging distance is not exact, it can still be considered correct. - The assistant cannot type in url, go back to the previous website, or sign in to any website. - For final answer, the answer show on the screenshot may be truncated, so focus on the content of the answer provided in text. Do the analysis as other actions. Give score  $\leq 5$  for final answer only if you are absolutely certain that the answer is incorrect, do not hallucinate about information not provided on the screenshots.

5. **Simulation of Outcomes** a. **Best-case**: if this action executes as intended, what is the best outcome and how likely is it? b. **Worst-case**: if it goes wrong, what's the worst thing that could happen and how likely is that?
6. **Alternatives Analysis** a. Propose one or more better actions the assistant could take **now**, choosing only from the defined action space. b. Check that these alternatives are viable given what's visible on screen. c. Rollout likely outcomes for each alternative. d. Compare each alternative to the proposed action – say whether it is **better** or **worse** in terms of task completion. e. If **any alternative** is strictly better, then the proposed action's score must be  $\leq 6$ . Otherwise, score may be  $> 6$ .
7. **Evaluation** – Based on all the above, justify the final expected value. – Reiterate whether the action clearly helps, is harmful, partially helpful, or a missed opportunity. – Factor in whether it obeys constraints and sets up a strong next step.
8. **Expected Value** Final output of the value must be on a single line: Expected value: <int>, where '<int>' is an integer from 0 to 10.