

WEBARBITER: A PRINCIPLE-GUIDED REASONING PROCESS REWARD MODEL FOR WEB AGENTS

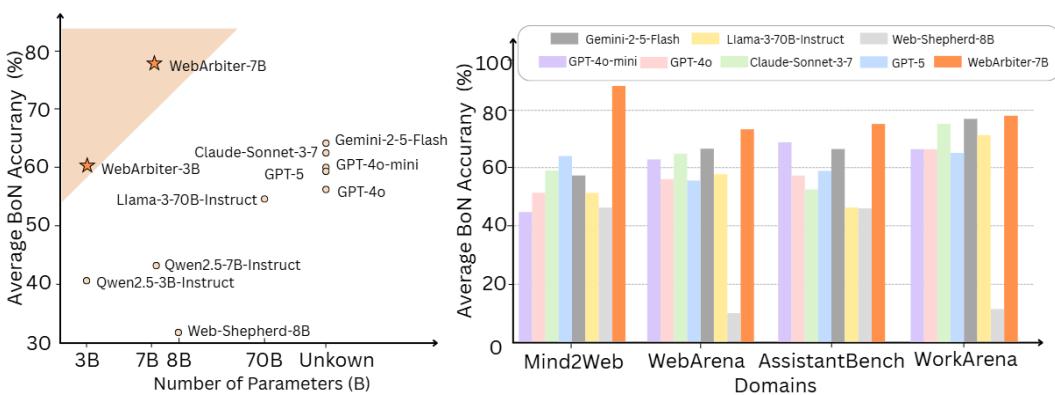
000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT

011 Web agents hold great potential for automating complex computer tasks, yet their
012 interactions involve long horizons, multi-step decisions, and actions that can be
013 irreversible. In such settings, outcome-based supervision is sparse and delayed,
014 often rewarding incorrect trajectories and failing to support inference-time scaling.
015 This motivates the use of Process Reward Models (WebPRMs) for web navigation,
016 but existing approaches remain limited: scalar WebPRMs collapse progress into
017 coarse, weakly grounded signals, while checklist-based WebPRMs rely on brittle
018 template matching that fails under layout or semantic changes and often mislabels
019 superficially correct actions as successful, providing little insight or interpretability.
020 To address these challenges, we introduce **WebArbiter**, a reasoning-first, principle-
021 inducing WebPRM that formulates reward modeling as text generation, producing
022 structured justifications that conclude with a preference verdict and identify the
023 action most conducive to task completion under the current context. Training
024 follows a two-stage pipeline: reasoning distillation equips the model with coherent
025 principle-guided reasoning, and reinforcement learning corrects teacher biases by
026 directly aligning verdicts with correctness, enabling stronger generalization. To
027 support systematic evaluation, we release WEBPRMBENCH, a comprehensive
028 benchmark spanning four diverse web environments with rich tasks and high-quality
029 preference annotations. On WEBPRMBENCH, WebArbiter-7B outperforms the
030 strongest baseline, Gemini Flash, by 10.9%. In reward-guided trajectory search on
031 WebArena-Lite, it surpasses the best prior WebPRM by up to 7.2%, underscoring
032 its robustness and practical value in real-world complex web tasks.
033
034

1 INTRODUCTION

035 Large Language Models (LLMs) (Achiam et al., 2023; Guo et al., 2025a) have demonstrated impres-
036 sive capabilities in planning (Huang et al., 2024; Zhang et al., 2025a), decision-making (Li et al.,
037



051 Figure 1: Performance comparison on WEBPRMBENCH. **Left:** Average Best-of-N Acc vs. model
052 size, showing superior efficiency despite smaller scale. **Right:** Domain-wise Avg BoN Acc, where
053 WEBARBITER achieves the best results across all environments, confirming robustness and scalability.

054 2024), and complex task execution (Xi et al., 2024; Zhang et al., 2025b). Extending these abilities
 055 with browser access enables LLM agents to perform complex web tasks similar to humans (OpenAI,
 056 2025a; Anthropic, 2024a; Adept, 2022). However, web interactions involve long horizons, multi-step
 057 decisions, and actions that can be irreversible. For example, submitting an incorrect form may not be
 058 recoverable. This requires agents to make reliable decisions throughout the interaction process, rather
 059 than relying solely on final outcomes. Traditional Outcome Reward Models (ORMs) are ill-suited:
 060 they provide only sparse and delayed feedback, may misclassify incorrect trajectories as successes,
 061 and cannot guide inference-time strategies, such as reward-guided search.

062 Recent studies on web agents (Zhang et al., 2025b; Koh et al., 2025) have introduced step-level
 063 rewards using LLM-as-judge. While such supervision can be useful, LLM-as-judge suffers from
 064 high cost, limited scalability, and susceptibility to hallucination, often rewarding fluent but incorrect
 065 actions. This motivates the development of dedicated Process Reward Models (WebPRMs) for web
 066 tasks. Existing WebPRMs largely fall into two categories: scalar WebPRM (Miao et al., 2025), which
 067 collapse progress into coarse scores with little interpretability or weak grounding; and generative
 068 WebPRM (Chae et al., 2025), which rely on checklists that are brittle under dynamic layouts and
 069 shifting semantics. Moreover, lacking explicit reasoning, generative WebPRMs remain vulnerable
 070 to surface correlations and sensitive to page changes. These limitations highlight the need for a
 071 reasoning-first WebPRM that can verify progress, resist superficial biases, and provide interpretable
 072 chains for diagnosing errors.

073 To this end, we propose **WebArbiter**, a reasoning-first, principle-inducing WebPRM. It formulates
 074 process reward modeling as text generation: given task context and candidate actions with their
 075 reasoning traces, the model produces a structured justification that concludes with a preference verdict,
 076 identifying the action most conducive to task completion. Unlike scalar scores or checklist-based
 077 methods tied to fixed templates, WebArbiter dynamically derives principles from user intent and
 078 the current state, incorporates them into reasoning chains that verify whether an action advances
 079 task completion. Training follows a two-stage pipeline: reasoning distillation equips the model
 080 with coherent principle-guided reasoning, and reinforcement learning corrects teacher biases and
 081 aligns verdicts with correctness. This design transforms reward signals from shallow correlations
 082 into auditable analyses, making judgments robust to environment and page variations, resistant to
 083 spurious cues, and accurate in credit assignment.

084 To advance the evaluation of WebPRMs, we introduce **WEBPRMBENCH**, the first comprehensive
 085 evaluation benchmark spanning diverse environments dedicated to WebPRMs. It provides 1,287
 086 step-level preference instances, each consisting of one correct action and four rejected alternatives,
 087 collected across 4 web environments: AssistantBench (Yoran et al., 2024), Mind2Web (Deng et al.,
 088 2023), WorkArena (Drouin et al., 2024; Boisvert et al., 2025), and WebArena (Zhou et al., 2023).
 089 The tasks span everyday activities such as online shopping and forum posting, as well as enterprise
 090 scenarios like updating schedules in IT management platforms. By combining scale, diversity, and
 091 fine-grained supervision, **WEBPRMBENCH** establishes a unified standard for systematic evaluation
 092 of WebPRMs, with *Pairwise* and *Best-of-N (BoN) Accuracy* as the primary metrics.

093 Extensive experiments on **WEBPRMBENCH** show that WebArbiter achieves SOTA *Avg. BoN Acc*,
 094 consistently surpassing both proprietary LLMs and previous SOTA WebPRM, WebShepherd, across
 095 all environments, and outperforming the strongest LLM baseline, Gemini Flash, by +10.9%. Beyond
 096 static evaluation, WebArbiter also proves effective in practice: in reward-guided trajectory search on
 097 WebArena-Lite (Liu et al., 2024b), it delivers substantial gains, surpassing WebShepherd by up to
 098 7.2%, further demonstrating robustness in realistic interaction settings.

099 The key contributions of this work are:

1. We propose WebArbiter, a reasoning-first, principle-inducing PRM trained with reasoning distillation and reinforcement learning, providing auditable reasoning chains and correctness-aligned signals.
2. We release **WEBPRMBENCH**, the first comprehensive evaluation benchmark to provide systematic WebPRM evaluation across 4 web environments, using *Pairwise* and *Best-of-N (BoN) Accuracy* as standard metrics.
3. We show that WebArbiter achieves SOTA performance on **WEBPRMBENCH**, surpassing both proprietary LLMs and the previous SOTA WebPRM. WebArbiter delivers up to +7.2% gains in reward-guided trajectory search on WebArena-Lite.

108 4. We analyze the training dynamics of WebArbiter, revealing how different strategies influence
 109 performance.
 110

111 **2 RELATED WORK**

114 **2.1 LLM-BASED AUTONOMOUS WEB AGENTS**

116 LLM advances have enabled browser-operating agents (Kim et al., 2024; Sun et al., 2024; Prasad
 117 et al., 2023; Fu et al., 2024; Ma et al., 2023; Zheng et al., 2023b; Tao et al., 2023). One line distills
 118 environment-specific state–action pairs from demonstrations, strong on seen states yet brittle on novel
 119 ones, with SteP as a leading example on WebArena (Sodhi et al., 2024; Zhou et al., 2023). A second
 120 line pursues open-ended exploration via reflexive evaluation and search (Pan et al., 2024; Shinn et al.,
 121 2024; Koh et al., 2024; Zhang et al., 2025b). A third direction applies reinforcement learning (Qi
 122 et al., 2025; Wei et al., 2025), yet real sites provide sparse and delayed signals, which makes value
 123 learning unstable without dense step feedback. Therefore, WebAgents require a process-level judge
 124 that assesses progress step by step and supplies auditable signals for search and planning.

125 **2.2 REWARD MODELS IN REASONING AND WEB TASKS**

127 RMs fall into two families. Scalar RMs attach a single numeric score to a response with a linear
 128 scorer and use either absolute or discriminative schemes for evaluation (Uesato et al., 2022; Ouyang
 129 et al., 2022; Liu et al., 2024a; 2025; Park et al., 2024; Wang et al., 2024a; 2023b; 2024b). Generative
 130 RMs instead produce natural–language feedback from which rewards are extracted, aligning with
 131 LLM-as-Judge and supporting both single-instance evaluation and multi-response comparison; they
 132 show promising scalability but raise reliability concerns due to bias and hallucination (Lightman et al.,
 133 2023; Wang et al., 2023a; Zhang et al., 2025c; Wu et al., 2024; Ye et al., 2025; Zhang et al., 2024;
 134 Zheng et al., 2023a). Building on these, Reasoning RMs cast judging as a deliberate process: they first
 135 generate an explicit, context-grounded chain of principle and analysis, then issue a single preference
 136 verdict, yielding adaptive test-time compute, stronger grounding, and interpretable feedback (Chen
 137 et al., 2025a; Guo et al., 2025b; Mahan et al., 2024). In web agents, action rewards have been
 138 derived by the following methods: LLM-as-Judge (Zhang et al., 2025b; Koh et al., 2025), slow
 139 and unstable during search; scalar scoring (Miao et al., 2025), which collapses progress into coarse
 140 values with little interpretability and weak grounding; and checklist-driven generative feedback (Chae
 141 et al., 2025), whose external templates are brittle under layout and semantic drift and prone to
 142 surface correlations. These limitations motivate a reasoning-first approach that turns rewards from
 143 shallow correlations into auditable analyses. WebArbiter produces structured justifications with a
 144 single preference verdict, induces principles from the current instruction and state, and is trained
 145 by reasoning distillation followed by reinforcement learning, so that judgments remain robust to
 146 environment variations, resist spurious cues, and provide accurate credit assignment while supporting
 147 inference-time scaling.

148 **3 METHODOLOGY**

150 In this section, we present the design of WebArbiter. We begin by framing web navigation as a
 151 Partially Observable Markov Decision Process (POMDP) in §3.1, then describe how we construct a
 152 pairwise-preference dataset for training in §3.2. We introduce the training pipeline of WebArbiter
 153 model in §3.3. For clarity, we summarize all notations in Appendix A.

155 **3.1 BACKGROUND**

157 We formalize web navigation as a POMDP. The environment \mathcal{E} is defined by a state space \mathcal{S} , an action
 158 space \mathcal{A} , and an observation space \mathcal{O} . $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ denotes the state transition function. At step p ,
 159 the agent receives a partial observation $o_p \in \mathcal{O}$, executes $a_p \in \mathcal{A}$, and transitions to $s_{p+1} = T(s_p, a_p)$
 160 with a new observation o_{p+1} . Following WebArena (Zhou et al., 2024), we represent observations
 161 using accessibility trees, i.e., text-only encodings of visible interactive elements and their attributes.
 Given a task instruction \mathcal{I} and the initial state $s_0 \in \mathcal{S}$, the agent aims to generate a trajectory

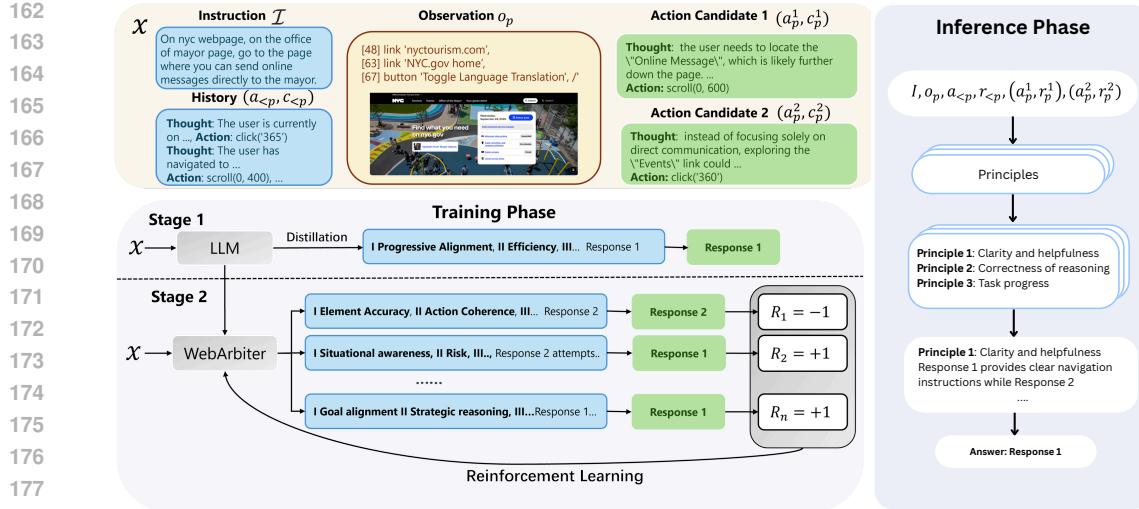


Figure 2: Overview of WebArbiter. Given an instruction \mathcal{I} , current observation o_p , and history $(a_{<p}, c_{<p})$, the model compares candidate actions (a_p^1, c_p^1) and (a_p^2, c_p^2) . In **Stage 1**, principle-guided reasoning traces are distilled from a stronger teacher LLM. In **Stage 2**, WEBARBITER is trained with reinforcement learning using verifiable rewards $R \in \{-1, +1\}$, producing structured justifications and a final verdict. During inference, the model induces principles (e.g., clarity, correctness, progress) from $(\mathcal{I}, o_p, a_{<p}, c_{<p}, (a_p^1, r_p^1), (a_p^2, r_p^2))$, applies them to candidate actions, and outputs an auditable judgment identifying the action that best advances task completion.

$\tau = (a_1, \dots, a_P)$ that completes the task. Here P is the trajectory length and $a_p \in \mathcal{A}$ denotes the action at step p . The task evaluator determines whether the task is completed based on the final state.

3.2 TRAINING DATASET CONSTRUCTION

We build on the **WEBPRM COLLECTION** (Chae et al., 2025) for training WebArbiter. Each instance consists of an instruction \mathcal{I} , a sequence of observations $O = (o_1, \dots, o_P)$, and expert-annotated trajectories. Specifically, the dataset provides a set of positive actions $A^+ = (a_1^+, \dots, a_P^+)$ taken from expert demonstrations and negative actions $A^- = (a_1^-, \dots, a_P^-)$ obtained from rejected trajectories. We convert these into pairwise preference samples where each candidate action is paired with its reasoning trace, yielding the preference dataset $\mathcal{D}_{\text{Train}}$ used for WebArbiter training.

3.3 WEBARBITER: A PRINCIPLE-INDUCING REASONING PROCESS REWARD MODEL

WebArbiter is built on a Transformer-decoder backbone and formulates process reward modeling as a text generation task. At each state, it evaluates candidate actions $\{(a_p^q, c_p^q)\}_{q=1}^Q$, where each action a_p^q is paired with a reasoning trace c_p^q explaining why the agent generated this action. Given task instruction \mathcal{I} , observation o_p , and history $(a_{<p}, c_{<p})$, the model autoregressively generates a structured justification $j = (j_1, \dots, j_L)$ of length L that concludes with a preference verdict \hat{y} selecting the most appropriate action among the candidates. The historical traces are $c_{<p} = \{c_1, \dots, c_{p-1}\}$, i.e., the per-action reasoning traces for previously executed actions. A concrete training example is provided in Appendix B. While our experiments instantiate this framework in the standard pairwise preference setting, the design is general and extends naturally to multi-candidate.

Unlike the scalar WebPRM (Miao et al., 2025) that collapses progress into opaque scores or the checklist-based WebPRM (Chae et al., 2025), WebArbiter is a reasoning-first, principle-inducing WebPRM: it dynamically derives principles from user intent and the current state, integrates them into reasoning chains that explicitly assess whether each candidate action truly advances task completion. This design moves reward signals beyond shallow correlations toward auditable analyses, yielding judgments that are robust to environment changes, resistant to spurious cues, and precise in credit assignment.

216 Formally, the preference dataset is defined as
 217

$$218 \quad \mathcal{D}_{\text{Train}} = \{(\mathcal{I}^{(i)}, o_p^{(i)}, a_{<p}^{(i)}, c_{<p}^{(i)}, (a_p^{1(i)}, c_p^{1(i)}), (a_p^{2(i)}, c_p^{2(i)}), y^{(i)})\}_{i=1}^M, \quad (1)$$

219 where $y \in \{a_p^1, a_p^2\}$ denotes the preferred action. For notational simplicity, let
 220

$$221 \quad x = (\mathcal{I}, o_p, a_{<p}, c_{<p}, (a_p^1, c_p^1), (a_p^2, c_p^2)). \quad (2)$$

223 WebArbiter π_θ is parameterized by θ and models the justification j autoregressively:
 224

$$225 \quad 226 \quad \pi_\theta(j \mid x) = \prod_{l=1}^L \pi_\theta(j_l \mid x, j_{<l}). \quad (3)$$

228 3.3.1 TRAINING OVERVIEW

230 The overall training objective is to maximize the likelihood that the predicted preference matches the
 231 ground truth:

$$232 \quad \max_{\pi_\theta} \mathbb{E}_{(x, y) \sim \mathcal{D}_{\text{Train}}, \hat{y} \sim \pi_\theta(j \mid x)} [\mathbb{1}(\hat{y} = y)]. \quad (4)$$

234 Training proceeds in two stages. The first stage, described in §3.3.2, is reasoning distillation, which
 235 equips the model with the ability to generate coherent principle-guided justifications. This stage
 236 encourages judgments to be grounded in explicit reasoning rather than surface correlations, as we
 237 later validate through ablation studies in §5.1.3.

238 Concretely, we sample K examples from $\mathcal{D}_{\text{Train}}$ to form \mathcal{D}_{SFT} for supervised distillation, while the
 239 remaining data is used as \mathcal{D}_{RL} for reinforcement learning. The second stage, detailed in §3.3.3, is
 240 reinforcement learning, which aligns the verdicts with correctness signals and produces interpretable
 241 step-level rewards for long-horizon tasks. Together, these stages enable WebArbiter to deliver robust,
 242 interpretable, and scalable supervision for web agents.
 243

244 3.3.2 STAGE 1: REASONING DISTILLATION

246 Directly prompting an instruction-tuned LLM as a reward model often yields superficial, inconsistent
 247 chains that do not justify why an action advances the task. We therefore distill principle-guided
 248 reasoning from a stronger teacher. Concretely, $\circ 3$ synthesizes structured justifications that first
 249 derive task-specific principles from the instruction and state, then ground these principles in the
 250 page, compare candidate actions against them, and finally output the preferred action. This equips
 251 WebArbiter with principles rather than surface heuristics. From ablations, we observe that removing
 252 explicit principles and using reasoning-only justifications markedly degrades performance, under-
 253 scoring the importance of principle induction for stable step-level judgments on the web. Given
 254 $(x^{(i)}, y^{(i)}) \in \mathcal{D}_{\text{SFT}}$, the teacher generates a justification $\hat{j}^{(i)} = (\hat{j}_1^{(i)}, \dots, \hat{j}_{L_i}^{(i)})$. The distillation dataset
 255 is then: $\mathcal{D}_{\text{SFT}} = \{x^{(i)}, \hat{j}^{(i)}\}_{i=1}^K$.

256 **Objective.** Reasoning distillation adjusts θ to maximize the likelihood of generating the teacher
 257 justification \hat{j} that concludes with the preferred action y given x . We minimize the standard negative
 258 log-likelihood:

$$259 \quad 260 \quad \mathcal{L}_{\text{SFT}}(\theta) = -\frac{1}{K} \sum_{i=1}^K \sum_{l=1}^{L_i} \log \pi_\theta(\hat{j}_l^{(i)} \mid x^{(i)}, \hat{j}_{<l}^{(i)}). \quad (5)$$

262 3.3.3 STAGE 2: REINFORCEMENT LEARNING

264 While distillation provides initial reasoning ability, it inherits teacher biases and may overfit to
 265 superficial patterns, limiting generalization to unseen environments. To further enhance judgment
 266 accuracy, stability, and generalization, we introduce a reinforcement learning stage. WebArbiter π_θ is
 267 treated as a policy that outputs a justification j that concludes with a final verdict \hat{y} . During rollout,
 268 π_θ generates the full justification and verdict, after which a correctness reward $R(x, \hat{y}) \in \{-1, 1\}$ is
 269 assigned solely based on whether \hat{y} matches the ground-truth preference y . The distilled model from
 §3.3.2 serves as the reference policy π_{ref} , ensuring stable optimization.

270 Table 1: Data distribution of WEBPRMBENCH, the first comprehensive evaluation benchmark
 271 spanning diverse environments for WebPRMs.
 272

273 Models	274 Mind2Web			275 WebArena	276 AssistantBench	277 WorkArena	278 Avg.
279 Models	280 Cross-Task	281 Cross-Website	282 Cross-Domain	275 WebArena	276 AssistantBench	277 WorkArena	278 Avg.
	142	148	283 417				
282 Count	142	148	283 417	371	29	180	1287

279 **Objective.** Reinforcement learning adjusts θ to maximize the expected reward while stabilizing
 280 reasoning style via KL regularization. The optimization objective is defined as:

$$281 \mathcal{L}_{\text{RL}}(\theta) = \max_{\pi_\theta} \mathbb{E}_{(x, y) \sim \mathcal{D}_{\text{RL}}, \hat{y} \sim \pi_\theta(j|x)} \left[R(x, \hat{y}) \right] - \beta \mathbb{D}_{\text{KL}}(\pi_\theta \parallel \pi_{\text{ref}}). \quad (6)$$

282 In practice, we adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024) to optimize this
 283 objective, which enables stable updates under binary verifiable rewards. Through this reinforcement
 284 learning stage, WebArbiter directly aligns its verdicts with correctness signals and converts structured
 285 justifications into reliable, interpretable step-level reward signals.

286 4 WEBPRMBENCH

289 In this section, we introduce WEBPRMBENCH, the first comprehensive evaluation benchmark
 290 spanning diverse environments for WebPRMs. Details of dataset construction and the evaluation
 291 protocol are provided below.

292 4.1 BENCHMARK CONSTRUCTION

294 WEBPRMBENCH is constructed from successful trajectories in AGENTREWARDBENCH (Lù et al.,
 295 2025), expanding beyond WEBREWARDBENCH (Chae et al., 2025), which only provides Mind2Web
 296 and limited WebArena data. We enrich WebArena with additional trajectories and incorporate Assis-
 297 tantBench and WorkArena, resulting in broader coverage of real-world tasks across four domains:
 298 Mind2Web, WebArena, AssistantBench, and WorkArena. Mind2Web emphasizes cross-task gen-
 299 eralization across heterogeneous websites. WebArena provides controlled environments such as
 300 shopping, CMS, forums, and GitLab. AssistantBench introduces open-world tasks on real websites.
 301 WorkArena focuses on enterprise workflows, including IT and HR. This diversity enables systematic
 302 evaluation across both consumer-facing and enterprise scenarios, while covering different levels of
 303 control, openness, and task complexity.

304 For each state, the action from the successful trajectory is retained as the positive label, and four
 305 rejected alternatives with associated reasoning traces are synthesized to form preference pairs. To
 306 ensure data quality, we sample negatives from diverse policy models to broaden coverage, apply rule-
 307 based filters to remove invalid or mismatched actions, discard inconsistent cases, and conduct expert
 308 verification to further ensure reliability. We also conduct targeted auditing to eliminate potential
 309 false negatives. Reasoning traces are truncated to a fixed length to minimize formatting noise. The
 310 resulting benchmark spans 1,287 preference pairs across four environments, as shown in Tab. 1.

311 4.2 EVALUATION PROTOCOL

313 Evaluating WebPRMs requires metrics that capture both local preference fidelity and global decision
 314 reliability under realistic multi-candidate settings. Inspired by RMB (Zhou et al., 2025), we adopt two
 315 complementary metrics: *Pairwise Accuracy*, which measures correctness on individual preference
 316 pairs, and *Best-of-N (BoN) Accuracy*, which evaluates robustness when ranking among multiple
 317 distractors. Compared with *Pairwise Acc*, *BoN Acc* applies a stricter criterion by requiring the correct
 318 action to outrank all distractors simultaneously, providing stronger discriminative power and better
 319 alignment with downstream agent performance.

320 **Pairwise Acc.** Given a preference pair (a^+, a^-) , where a^+ is the correct action and a^- a rejected
 321 one, the WebPRM is correct if it assigns higher preference to a^+ . Formally:

$$322 \text{Acc}_{\text{Pairwise}} = \frac{1}{|\mathcal{D}_{\text{Bench}}|} \sum_{(a^+, a^-) \in \mathcal{D}_{\text{Bench}}} \mathbb{1}[\pi_\theta(a^+) \succ \pi_\theta(a^-)]. \quad (7)$$

324
 325 Table 2: Results on WEBPRMBENCH with *Pairwise* and *BoN Acc*. \star denotes our models. Bold numbers
 326 indicate the best results, while underlined numbers denote the second best. Our WebArbiter-7B
 327 achieves the highest *BoN Acc* across all environments, with an Avg. *BoN Acc* of 77.78%, outperforming
 328 the second-best baseline, i.e., Gemini Flash, by 10.85%.

Models	Mind2Web		WebArena		AssistantBench		WorkArena		Avg.	
	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN
<i>LLM-as-judge, Proprietary Language Models</i>										
GPT-4o-mini	80.28	45.69	81.40	61.46	88.79	<u>68.97</u>	84.86	66.67	83.83	60.70
GPT-4o	80.60	52.84	79.78	56.06	83.62	<u>58.62</u>	85.69	66.67	82.42	58.55
GPT-5	82.06	62.26	76.89	55.26	76.72	58.62	80.83	65.56	79.12	60.42
Claude Sonnet	81.19	59.60	81.40	62.53	76.92	53.85	<u>88.06</u>	72.22	81.89	62.05
Gemini Flash	81.94	57.86	<u>83.89</u>	<u>67.39</u>	<u>84.62</u>	65.38	92.36	<u>77.08</u>	<u>85.70</u>	<u>66.93</u>
DeepSeek-R1	82.18	56.44	81.13	60.38	<u>75.00</u>	51.72	88.61	<u>72.78</u>	81.73	60.33
<i>LLM-as-judge, Open-source Language Models</i>										
Qwen2.5-3B-Instruct	76.89	37.66	69.54	35.04	80.17	48.28	70.56	42.22	74.29	40.80
Qwen2.5-7B-Instruct	79.02	41.00	72.44	40.16	76.72	37.93	78.06	51.67	76.56	42.69
Llama-3-70B-Instruct	79.43	50.96	77.90	56.60	78.85	46.15	87.50	70.14	80.92	55.96
<i>WebPRMs (3B)</i>										
WebShepherd-3B	37.41	21.22	20.33	9.47	36.54	17.24	10.49	2.44	26.19	12.59
\star WebArbiter-3B	<u>93.28</u>	<u>78.75</u>	83.29	56.87	76.72	44.83	84.03	60.56	84.33	60.25
<i>WebPRMs (7B+)</i>										
WebShepherd-8B	72.35	46.68	33.16	12.37	55.77	44.83	35.85	12.68	49.28	29.14
\star WebArbiter-7B	96.47	90.07	84.30	71.43	80.17	72.41	87.36	77.22	87.08	77.78

347 **BoN Acc.** For each instance $(a^+, a^{-1}, \dots, a^{-Q}) \in \mathcal{D}_{\text{Bench}}$, the WebPRM is considered correct only
 348 when a^+ is consistently ranked above all Q distractors, with $Q = 4$ in our benchmark. BoN Acc is:

$$\text{Acc}_{\text{BoN}} = \frac{1}{|\mathcal{D}_{\text{Bench}}|} \sum_{i=1}^{|\mathcal{D}_{\text{Bench}}|} \prod_{q=1}^Q \mathbb{1}[\pi_\theta(a_i^+) \succ \pi_\theta(a_i^{-q})]. \quad (8)$$

353 5 EXPERIMENTS

355 We conduct comprehensive experiments to evaluate WebArbiter on the reward modeling benchmark
 356 WEBPRMBENCH in § 5.1 and on practical applications in § 5.2.

358 5.1 WEBPRMBENCH

360 5.1.1 EXPERIMENTAL SETUP

362 **Baselines.** We compare WebArbiter against three categories of baselines. (1) Proprietary LLM-
 363 as-judge models, including GPT-4o-mini (OpenAI, 2024a), GPT-4o (OpenAI, 2024b), GPT-5 (Ope-
 364 nAI, 2025b), Claude-3.7-Sonnet (Anthropic, 2025), Gemini-2.5-Flash (Comanici et al., 2025), and
 365 DeepSeek-R1 (Guo et al., 2025a), which are prompted to act as judges by selecting the preferred ac-
 366 tion given task context. (2) Open-source LLM-as-judge models, represented by Qwen2.5-3B-Instruct
 367 and Qwen2.5-7B-Instruct (Qwen et al., 2025), and Llama-3-70B-Instruct (Grattafiori et al., 2024),
 368 providing accessible yet competitive instruction-tuned baselines. (3) WebPRMs, where we include
 369 WebShepherd (Chae et al., 2025).

370 **Implementation Details.** We train WebArbiter on WEBPRM Collection (Chae et al., 2025), which
 371 comprises 30k step-level preference pairs drawn from the Mind2Web environment. We use 10k pairs
 372 for stage-1 reasoning distillation and the remainder for stage-2 reinforcement learning. Models are
 373 initialized from Qwen2.5-3B-Instruct and Qwen2.5-7B-Instruct (Qwen et al., 2025) and fine-tuned
 374 with LoRA (Hu et al., 2022). Further implementation details are provided in the Appendix C.

376 **Evaluation Metrics.** We report results using two complementary metrics: *Pairwise Accuracy*,
 377 which measures correctness on individual preference pairs, and *Best-of-N (BoN) Accuracy*, which
 378 evaluates robustness under multi-candidate settings. Detailed definitions are provided in § 4.2.

378
 379 Table 3: Ablation results on WEBPRMBENCH with Qwen2.5-7B-Instruct as backbone. Best results
 380 are in bold and the second best underlined. WEBARBITER, combining principle-guided reasoning
 381 distillation with RL, achieves the highest overall performance.

Method	Mind2Web		WebArena		AssistantBench		WorkArena		Avg.	
	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN
Instruct (Original)	79.02	41.00	72.44	40.16	76.72	37.93	78.06	51.67	76.56	42.69
Instruct + Cold Start RL	97.63	<u>91.38</u>	67.59	43.40	71.55	34.48	73.33	55.00	77.53	56.07
Instruct + Cold Start RL + Principles	96.42	87.88	<u>84.10</u>	<u>60.65</u>	<u>79.31</u>	<u>55.17</u>	<u>83.19</u>	<u>55.56</u>	<u>85.75</u>	<u>64.81</u>
Instruct + SFT _{w/o Principles} + RL	94.26	82.39	75.34	49.87	68.97	41.38	78.61	54.44	79.30	57.02
★ WebArbiter	<u>96.47</u>	<u>90.07</u>	84.30	71.43	80.17	72.41	87.36	77.22	87.08	77.78

388 389 5.1.2 MAIN RESULTS 390

391 **WebArbiter Significantly Outperforms Baselines.** As shown in Tab. 2, WebArbiter consistently
 392 surpasses both proprietary and open-source LLMs across all environments with *BoN Acc*. While
 393 LLM-as-judge methods often maintain moderate *Pairwise Acc*, their performance drops sharply on
 394 *BoN Acc*, revealing poor robustness to hard negatives. In contrast, WebArbiter sustains strong results
 395 on both metrics, establishing its reliability under realistic multi-candidate settings.

396 **Advantage over the SOTA WebPRM.** WebShepherd (Chae et al., 2025) represents the previous
 397 SOTA WebPRMs. Trained on the same WEBPRM Collection, which was drawn from the Mind2Web
 398 environment, WebArbiter-7B achieves an *Avg. BoN Acc* of 77.78%, surpassing WebShepherd-8B
 399 by an absolute gain of 48%. Unlike WebShepherd, which relies on fragile checklists, WebArbiter
 400 employs principle-guided reasoning, yielding judgments robust to environment and page variations.
 401 Case studies illustrating these differences are provided in Appendix E.

402 **Robust Generalization Across Environments.** WebArbiter not only excels in-domain, achieving
 403 96.47% *Pairwise Acc* and 90.07% *BoN Acc* on Mind2Web, but also generalizes across diverse
 404 benchmarks. On WebArena, it outperforms the second-best baseline by nearly 4% in *BoN Acc*,
 405 gains about 3% on AssistantBench, and still edges out strong baselines on WorkArena with 77.22%.
 406 These results confirm that principle-guided reasoning supports both strong in-domain learning and
 407 robustness across heterogeneous, noisy, and enterprise-level environments.

408 409 5.1.3 ABLATION STUDY 410

411 We compare four training variants to disentangle the effects of reinforcement learning, principle
 412 guidance, and justification style. *Instruct (Original)* denotes a purely instruction-tuned model without
 413 additional optimization. *Cold Start RL* directly applies RL on top of the instruction model. *Cold Start*
 414 *RL + Principles* augments RL with principle prompting during training, enabling explicit principle
 415 induction before judgment. *SFT_{w/o Principles} + RL* performs reasoning distillation without principles,
 416 followed by RL, thereby testing whether narrative-style justifications alone are sufficient. As shown in
 417 Tab. 3, WebArbiter achieves the best performance. Explicit principles anchor judgments to progress,
 418 producing stable supervision under multi-candidate web settings.

419 **RL Alone is Unstable Across Web Environments.** *Cold Start RL* performs well on in-domain
 420 Mind2Web but collapses on out-of-domain benchmarks. This highlights that reward optimization
 421 without reasoning distillation struggles in noisy and complex environments.

423 **Principles Enable Cross-Environment Generalization.** Augmenting RL with principles boosts
 424 *Avg. BoN Acc*, especially in structurally diverse environments such as WebArena and AssistantBench.
 425 Principles provide transferable facets for reasoning, reducing reliance on brittle layout cues and
 426 improving robustness to web variability.

428 **Reasoning Without Principles is Insufficient.** *SFT_{w/o Principles} + RL*, i.e., narrative-style justifica-
 429 tions alone, improves fluency but lags behind principle-aware settings. This confirms that narrating
 430 reasoning chains without principles cannot ensure alignment with true task progress in complex,
 431 long-horizon real-world web navigation.

432
 433 Table 4: Success Rates (%) of trajectory search with GPT-4o-mini and GPT-4o as policy on WebArena-
 434 lite. * Results reported from the WebShepherd (Chae et al., 2025). Δ is relative to the *w/o Trajectory
 435 Search* baseline. Our WebArbiter consistently achieves the highest gains across both policy models.

Policy	WebPRM	Shopping	CMS	Reddit	GitLab	Avg.	Δ
GPT-4o-mini	w/o Trajectory Search*	21.74	22.86	19.05	34.38	24.51	-
	GPT-4o-mini	24.44	22.86	26.32	33.33	26.74	+2.23
	WebShepherd-8B*	26.09	45.71	23.81	40.62	34.06	+9.55
GPT-4o	★ WebArbiter-7B	37.78	42.86	36.84	46.67	41.04	+16.53
	w/o Trajectory Search*	23.91	31.43	28.57	56.25	35.04	-
	GPT-4o-mini	26.67	37.14	42.11	40.00	36.48	+1.44
WebArbiter	WebShepherd-8B*	30.43	42.86	47.62	46.88	41.95	+6.91
	★ WebArbiter-7B	44.44	42.86	52.63	56.67	49.15	+14.11

447 5.2 REWARD-GUIDED TRAJECTORY SEARCH

448 5.2.1 EXPERIMENTAL SETUP AND IMPLEMENTATIONS

449 Reward-guided trajectory search represents one of the most practical applications of PRMs, as
 450 it directly leverages fine-grained step-level supervision to improve decision quality during agent
 451 execution. To evaluate WebArbiter in this setting, we conduct experiments on WebArena-Lite¹ (Liu
 452 et al., 2024b), which contains diverse, long-horizon tasks such as online shopping and content
 453 management, closely reflecting real-world web activities. Performance is measured with Success
 454 Rate. Following WebShepherd (Chae et al., 2025), we adopt a Best-of-N sampling strategy: the
 455 policy model generates $N = 5$ candidate actions for each step, and WebArbiter selects the most
 456 promising one through a Knockout Tournament mechanism (Guo et al., 2025b). We evaluate two
 457 policies, GPT-4o-mini (OpenAI, 2024a) and GPT-4o (OpenAI, 2024b).

458 5.2.2 ANALYSIS

459 As shown in Tab. 4, WebArbiter achieves substantial average improvements under both policy models,
 460 far surpassing baselines. Its advantages arise from three main factors. First, reasoning mitigates
 461 spurious correlations that often mislead WebPRMs in domains such as Shopping and Reddit. Gains
 462 in Shopping are particularly striking, as tasks require dense semantic retrieval and inference; stronger
 463 policies can roll out more promising candidate actions, and WebArbiter’s structured reward modeling
 464 further amplifies these benefits. Second, in GitLab, tasks frequently allow multiple equivalent paths.
 465 WebShepherd is brittle under such variability, whereas WebArbiter leverages reasoning over historical
 466 trajectories and current states to evaluate action validity, enabling stronger generalization in dynamic
 467 workflows. By contrast, CMS exhibits a more template-driven structure, where actions closely
 468 follow standardized patterns. In such cases, checklist-based supervision remains comparatively
 469 effective, which narrows the relative performance margin. Overall, WebArbiter’s reasoning-first
 470 design consistently provides robust, interpretable, and scalable supervision across diverse domains.

471 6 CONCLUSION

472 We presented WEBARBITER, a reasoning-first, principle-inducing process reward model that frames
 473 reward modeling as structured text generation and produces auditable step-level judgments with
 474 rationales. Through reasoning distillation and reinforcement learning, WebArbiter transforms superfi-
 475 cial correlations into robust signals that verify genuine task progress, enforce trajectory consistency,
 476 and generalize across dynamic websites. To support systematic evaluation, we released WEBPRM-
 477 BENCH, the first comprehensive evaluation benchmark spanning diverse environments for WebPRMs
 478 in web navigation, covering four domains with diverse tasks and fine-grained step-level supervi-
 479 sion. Extensive experiments demonstrate SOTA performance on WEBPRM-BENCH and substantial
 480 improvements in reward-guided trajectory search on WebArena-Lite, establishing principle-guided
 481 reasoning WebPRMs as a robust and interpretable foundation for scalable web agents.

482 ¹We did not have access to the MAP domain during this work and therefore excluded it from our experiments.

486 REFERENCES
487

488 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
489 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.
490 *arXiv preprint arXiv:2303.08774*, 2023.

491 Adept. Act-1: Transformer for actions. adept.ai/blog/act-1/, 2022.

492 Anthropic. Introducing computer use, a new claudie 3.5 sonnet, and claudie 3.5 haiku.
493 anthropic.com/news/3-5-models-and-computer-use, 2024a.

494 Anthropic. Introducing computer use, a new claudie 3.5 sonnet, and claudie 3.5 haiku. <https://www.anthropic.com/news/3-5-models-and-computer-use>, October 2024b.

495 Anthropic. Claude 3.7 sonnet and claudie code. anthropic.com/news/claudie-3-7-sonnet,
496 2025.

497 David Anugraha, Zilu Tang, Lester James V Miranda, Hanyang Zhao, Mohammad Rifqi Farhansyah,
498 Garry Kuwanto, Derry Wijaya, and Genta Indra Winata. R3: Robust rubric-agnostic reward models.
499 *arXiv preprint arXiv:2505.13388*, 2025.

500 Léo Boisvert, Megh Thakkar, Maxime Gasse, Massimo Caccia, Thibault Le Sellier De Chezelles,
501 Quentin Cappart, Nicolas Chapados, Alexandre Lacoste, and Alexandre Drouin. Workarena++:
502 Towards compositional planning and reasoning-based common knowledge work tasks, 2025. URL
503 <https://arxiv.org/abs/2407.05291>.

504 Hyungjoo Chae, Sunghwan Kim, Junhee Cho, Seungone Kim, Seungjun Moon, Gyeom Hwangbo,
505 Dongha Lim, Minjin Kim, Yeonjun Hwang, Minju Gwak, Dongwook Choi, Minseok Kang,
506 Gwanhoon Im, ByeongUng Cho, Hyojun Kim, Jun Hee Han, Taeyoon Kwon, Minju Kim, Beong
507 woo Kwak, Dongjin Kang, and Jinyoung Yeo. Web-shepherd: Advancing prms for reinforcing
508 web agents, 2025. URL <https://arxiv.org/abs/2505.15277>.

509 Xiusi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang,
510 Denghui Zhang, Tong Zhang, Hanghang Tong, and Heng Ji. Rm-r1: Reward modeling as reasoning,
511 2025a. URL <https://arxiv.org/abs/2505.02387>.

512 Xiusi Chen, Gaotang Li, Ziqi Wang, Bowen Jin, Cheng Qian, Yu Wang, Hongru Wang, Yu Zhang,
513 Denghui Zhang, Tong Zhang, et al. Rm-r1: Reward modeling as reasoning. *arXiv preprint*
514 *arXiv:2505.02387*, 2025b.

515 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
516 Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier
517 with advanced reasoning, multimodality, long context, and next generation agentic capabilities.
518 *arXiv preprint arXiv:2507.06261*, 2025.

519 Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su.
520 Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing*
521 *Systems*, 36:28091–28114, 2023.

522 Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam H. Laradji, Manuel Del Verme, Tom
523 Marty, Léo Boisvert, Megh Thakkar, Quentin Cappart, David Vazquez, Nicolas Chapados, and
524 Alexandre Lacoste. Workarena: How capable are web agents at solving common knowledge work
525 tasks?, 2024. URL <https://arxiv.org/abs/2403.07718>.

526 Yao Fu, Dong-Ki Kim, Jaekyeom Kim, Sungryull Sohn, Lajanugen Logeswaran, Kyunghoon Bae,
527 and Honglak Lee. Autoguide: Automated generation and selection of state-aware guidelines for
528 large language model agents. *arXiv preprint arXiv:2403.08978*, 2024.

529 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
530 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
531 models. *arXiv preprint arXiv:2407.21783*, 2024.

532 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
533 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
534 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025a.

540 Jiaxin Guo, Zewen Chi, Li Dong, Qingxiu Dong, Xun Wu, Shaohan Huang, and Furu Wei. Reward
 541 reasoning model, 2025b. URL <https://arxiv.org/abs/2505.14674>.

542

543 Ilgee Hong, Changlong Yu, Liang Qiu, Weixiang Yan, Zhenghao Xu, Haoming Jiang, Qingru Zhang,
 544 Qin Lu, Xin Liu, Chao Zhang, et al. Think-rm: Enabling long-horizon reasoning in generative
 545 reward models. *arXiv preprint arXiv:2505.16265*, 2025.

546 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 547 Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

548

549 Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang,
 550 Ruiming Tang, and Enhong Chen. Understanding the planning of llm agents: A survey, 2024.
 551 URL <https://arxiv.org/abs/2402.02716>.

552 Geunwoo Kim, Pierre Baldi, and Stephen McAleer. Language models can solve computer tasks.
 553 *Advances in Neural Information Processing Systems*, 36, 2024.

554

555 Jing Yu Koh, Stephen McAleer, Daniel Fried, and Ruslan Salakhutdinov. Tree search for language
 556 model agents. *arXiv preprint arXiv:2407.01476*, 2024.

557

558 Jing Yu Koh, Stephen McAleer, Daniel Fried, and Ruslan Salakhutdinov. Tree search for language
 559 model agents, 2025. URL <https://arxiv.org/abs/2407.01476>.

560

561 Manling Li, Shiyu Zhao, Qineng Wang, Kangrui Wang, Yu Zhou, Sanjana Srivastava, Cem Gokmen,
 562 Tony Lee, Erran Li Li, Ruohan Zhang, et al. Embodied agent interface: Benchmarking llms for
 563 embodied decision making. *Advances in Neural Information Processing Systems*, 37:100428–
 100534, 2024.

564

565 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 566 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth
 567 International Conference on Learning Representations*, 2023.

568

569 Chris Yuhao Liu, Liang Zeng, Jiacai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang
 570 Liu, and Yahui Zhou. Skywork-reward: Bag of tricks for reward modeling in llms, 2024a. URL
 571 <https://arxiv.org/abs/2410.18451>.

572

573 Xiao Liu, Tianjie Zhang, Yu Gu, Iat Long Iong, Yifan Xu, Xixuan Song, Shudan Zhang, Hanyu Lai,
 574 Xinyi Liu, Hanlin Zhao, Jiadai Sun, Xinyue Yang, Yu Yang, Zehan Qi, Shuntian Yao, Xueqiao
 575 Sun, Siyi Cheng, Qinkai Zheng, Hao Yu, Hanchen Zhang, Wenyi Hong, Ming Ding, Lihang Pan,
 576 Xiaotao Gu, Aohan Zeng, Zhengxiao Du, Chan Hee Song, Yu Su, Yuxiao Dong, and Jie Tang.
 577 Visualagentbench: Towards large multimodal models as visual foundation agents, 2024b. URL
 578 <https://arxiv.org/abs/2408.06327>.

579

580 Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, and Juanzi Li. Pairjudge rm: Perform best-of-n
 581 sampling with knockout tournament, 2025. URL <https://arxiv.org/abs/2501.13007>.

582

583 Xing Han Lù, Amirhossein Kazemnejad, Nicholas Meade, Arkil Patel, Dongchan Shin, Alejandra
 584 Zambrano, Karolina Stańczak, Peter Shaw, Christopher J. Pal, and Siva Reddy. Agentrewardbench:
 585 Evaluating automatic evaluations of web agent trajectories, 2025. URL <https://arxiv.org/abs/2504.08942>.

586

587 Kaixin Ma, Hongming Zhang, Hongwei Wang, Xiaoman Pan, and Dong Yu. Laser: Llm agent with
 588 state-space exploration for web navigation. *arXiv preprint arXiv:2309.08172*, 2023.

589

590 Dakota Mahan, Duy Van Phung, Rafael Rafailov, Chase Blagden, Nathan Lile, Louis Castricato,
 591 Jan-Philipp Fränken, Chelsea Finn, and Alon Albalak. Generative reward models, 2024. URL
 592 <https://arxiv.org/abs/2410.12832>.

593

594 Bingchen Miao, Yang Wu, Minghe Gao, Qifan Yu, Wendong Bu, Wenqiao Zhang, Yunfei Li,
 595 Siliang Tang, Tat-Seng Chua, and Juncheng Li. Boosting virtual agent learning and reasoning:
 596 A step-wise, multi-dimensional, and generalist reward model with benchmark, 2025. URL
 597 <https://arxiv.org/abs/2503.18665>.

594 OpenAI. Gpt-4o mini: advancing cost-efficient intelligence. openai.com/gpt-4o-mini,
 595 2024a.

596

597 OpenAI. Gpt-4o. platform.openai.com/gpt-4o, 2024b.

598

599 OpenAI. Introducing operator. openai.com/introducing-operator, 2025a.

600

601 OpenAI. Gpt-5 is here. openai.com/gpt-5, 2025b.

602

603 Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong
 604 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton,
 605 Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and
 606 Ryan Lowe. Training language models to follow instructions with human feedback, 2022. URL
<https://arxiv.org/abs/2203.02155>.

607

608 Jiayi Pan, Yichi Zhang, Nicholas Tomlin, Yifei Zhou, Sergey Levine, and Alane Suhr. Autonomous
 609 evaluation and refinement of digital agents. *arXiv preprint arXiv:2404.06474*, 2024.

610

611 Junsoo Park, Seungyeon Jwa, Meiyi Ren, Daeyoung Kim, and Sanghyuk Choi. Offsetbias:
 612 Leveraging debiased data for tuning evaluators, 2024. URL <https://arxiv.org/abs/2407.06551>.

613

614 Archiki Prasad, Alexander Koller, Mareike Hartmann, Peter Clark, Ashish Sabharwal, Mohit Bansal,
 615 and Tushar Khot. Adapt: As-needed decomposition and planning with language models. *arXiv
 616 preprint arXiv:2311.05772*, 2023.

617

618 Zehan Qi, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Wenyi Zhao, Yu Yang, Xinyue
 619 Yang, Jiadai Sun, Shuntian Yao, Tianjie Zhang, Wei Xu, Jie Tang, and Yuxiao Dong. Webrl:
 620 Training llm web agents via self-evolving online curriculum reinforcement learning, 2025. URL
<https://arxiv.org/abs/2411.02337>.

621

622 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 623 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 624 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 625 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi
 626 Tang, Tingyu Xia, Xingzheng Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan,
 627 Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL
<https://arxiv.org/abs/2412.15115>.

628

629 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 630 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
 631 mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

632

633 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
 634 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint
 635 arXiv: 2409.19256*, 2024.

636

637 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:
 638 Language agents with verbal reinforcement learning. *Advances in Neural Information Processing
 639 Systems*, 36, 2024.

640

641 Paloma Sodhi, S. R. K. Branavan, Yoav Artzi, and Ryan McDonald. Step: Stacked llm policies for
 642 web actions, 2024.

643

644 Haotian Sun, Yuchen Zhuang, Lingkai Kong, Bo Dai, and Chao Zhang. Adapanner: Adaptive
 645 planning from feedback with language models. *Advances in Neural Information Processing
 646 Systems*, 36, 2024.

647

Heyi Tao, Sethuraman TV, Michal Shlapentokh-Rothman, Derek Hoiem, and Heng Ji. Webwise:
 648 Web interface control and sequential exploration with large language models. *arXiv preprint
 649 arXiv:2310.16042*, 2023.

648 Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia
 649 Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and
 650 outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.

651

652 Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui.
 653 Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv preprint*
 654 *arXiv:2312.08935*, 2023a.

655

656 Tianlu Wang, Ilia Kulikov, Olga Golovneva, Ping Yu, Weizhe Yuan, Jane Dwivedi-Yu,
 657 Richard Yuanzhe Pang, Maryam Fazel-Zarandi, Jason Weston, and Xian Li. Self-taught evaluators,
 658 2024a. URL <https://arxiv.org/abs/2408.02666>.

659

660 Zhilin Wang, Yi Dong, Jiaqi Zeng, Virginia Adams, Makesh Narsimhan Sreedhar, Daniel Egert,
 661 Olivier Delalleau, Jane Polak Scowcroft, Neel Kant, Aidan Swope, and Oleksii Kuchaiev. Helpsteer:
 662 Multi-attribute helpfulness dataset for steerlm, 2023b. URL <https://arxiv.org/abs/2311.09528>.

663

664 Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J. Zhang,
 665 Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. Helpsteer2: Open-source dataset for training
 666 top-performing reward models, 2024b. URL <https://arxiv.org/abs/2406.08673>.

667

668 Zhepei Wei, Wenlin Yao, Yao Liu, Weizhi Zhang, Qin Lu, Liang Qiu, Changlong Yu, Puyang Xu,
 669 Chao Zhang, Bing Yin, Hyokun Yun, and Lihong Li. Webagent-r1: Training web agents via
 670 end-to-end multi-turn reinforcement learning, 2025. URL <https://arxiv.org/abs/2505.16421>.

671

672 Tianhao Wu, Weizhe Yuan, Olga Golovneva, Jing Xu, Yuandong Tian, Jiantao Jiao, Jason Weston,
 673 and Sainbayar Sukhbaatar. Meta-rewarding language models: Self-improving alignment with
 674 llm-as-a-meta-judge, 2024. URL <https://arxiv.org/abs/2407.19594>.

675

676 Zhiheng Xi, Yiwen Ding, Wenxiang Chen, Boyang Hong, Honglin Guo, Junzhe Wang, Dingwen
 677 Yang, Chenyang Liao, Xin Guo, Wei He, et al. Agentgym: Evolving large language model-based
 678 agents across diverse environments. *arXiv preprint arXiv:2406.04151*, 2024.

679

680 Ziyi Ye, Xiangsheng Li, Qiuchi Li, Qingyao Ai, Yujia Zhou, Wei Shen, Dong Yan, and Yiqun Liu.
 681 Learning llm-as-a-judge for preference alignment. In *The Thirteenth International Conference on*
682 Learning Representations, 2025.

683

684 Ori Yoran, Samuel Joseph Amouyal, Chaitanya Malaviya, Ben Beglin, Ofir Press, and Jonathan
 685 Berant. Assistantbench: Can web agents solve realistic and time-consuming tasks?, 2024. URL
<https://arxiv.org/abs/2407.15711>.

686

687 Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal.
 688 Generative verifiers: Reward modeling as next-token prediction. *arXiv preprint arXiv:2408.15240*,
 689 2024.

690

691 Yao Zhang, Chenyang Lin, Shijie Tang, Haokun Chen, Shijie Zhou, Yunpu Ma, and Volker Tresp.
 692 Swarmagenic: Towards fully automated agentic system generation via swarm intelligence. *arXiv*
693 preprint arXiv:2506.15672, 2025a.

694

695 Yao Zhang, Zijian Ma, Yunpu Ma, Zhen Han, Yu Wu, and Volker Tresp. Webpilot: A versatile and
 696 autonomous multi-agent system for web task execution with strategic exploration. In *Proceedings*
697 of the AAAI Conference on Artificial Intelligence, volume 39, pp. 23378–23386, 2025b.

698

699 Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu,
 700 Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical
 701 reasoning. *arXiv preprint arXiv:2501.07301*, 2025c.

702

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

702 Longtao Zheng, Rundong Wang, Xinrun Wang, and Bo An. Synapse: Trajectory-as-exemplar
 703 prompting with memory for computer control. In *The Twelfth International Conference on*
 704 *Learning Representations*, 2023b.

705
 706 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyuan Luo, Zhangchi Feng, and
 707 Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Pro-*
 708 *ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3:*
 709 *System Demonstrations)*, Bangkok, Thailand, 2024. Association for Computational Linguistics.
 710
 URL <http://arxiv.org/abs/2403.13372>.

711 Enyu Zhou, Guodong Zheng, Binghai Wang, Zhiheng Xi, Shihan Dou, Rong Bao, Wei Shen,
 712 Limao Xiong, Jessica Fan, Yurong Mou, Rui Zheng, Tao Gui, Qi Zhang, and Xuanjing Huang.
 713 Rmb: Comprehensively benchmarking reward models in llm alignment, 2025. URL <https://arxiv.org/abs/2410.09893>.
 714

715 Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng,
 716 Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. Webarena: A realistic web environment for building
 717 autonomous agents. *arXiv preprint arXiv:2307.13854*, 2023.

718 Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng,
 719 Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic
 720 web environment for building autonomous agents, 2024. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2307.13854)
 721 2307.13854.
 722

723
 724
 725
 726
 727
 728
 729
 730
 731
 732
 733
 734
 735
 736
 737
 738
 739
 740
 741
 742
 743
 744
 745
 746
 747
 748
 749
 750
 751
 752
 753
 754
 755

756	CONTENTS	
757		
758		
759	1 Introduction	1
760		
761	2 Related Work	3
762	2.1 LLM-Based Autonomous Web Agents	3
763	2.2 Reward Models in Reasoning and Web Tasks	3
764		
765		
766	3 Methodology	3
767	3.1 Background	3
768	3.2 Training Dataset Construction	4
769	3.3 WebArbiter: a Principle-Inducing Reasoning Process Reward Model	4
770	3.3.1 Training Overview	5
771	3.3.2 Stage 1: Reasoning Distillation	5
772	3.3.3 Stage 2: Reinforcement Learning	5
773		
774		
775		
776		
777	4 WEBPRMBENCH	6
778	4.1 Benchmark Construction	6
779	4.2 Evaluation Protocol	6
780		
781		
782	5 Experiments	7
783	5.1 WEBPRMBENCH	7
784	5.1.1 Experimental Setup	7
785	5.1.2 Main Results	8
786	5.1.3 Ablation Study	8
787	5.2 Reward-Guided Trajectory Search	9
788	5.2.1 Experimental Setup and Implementations	9
789	5.2.2 Analysis	9
790		
791		
792		
793		
794		
795	6 Conclusion	9
796		
797	A Notation Summary	17
798		
799	B Example of Preference Dataset	17
800		
801		
802	C Training Details	18
803		
804		
805	D Prompt Repository	18
806		
807	E Case Study: WebArbiter vs. WebShepherd	19
808		
809	F Benchmark Construction	20

810 **G General-Domain Generative Reward Models and Their Transferability to Web Trajec-
811 tories**

20

812

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 **A NOTATION SUMMARY**
865866 For clarity, we summarize the main notations used throughout this paper:
867

- 868 • \mathcal{E} : web environment, defined by state space \mathcal{S} , action space \mathcal{A} , and observation space \mathcal{O} .
869
- 870 • T : state transition function $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$.
871
- 872 • \mathcal{I} : task instruction.
873
- 874 • s_p, o_p, a_p : state, observation, and action at step p .
875
- 876 • c_p : reasoning trace associated with action a_p .
877
- 878 • $c_{<p}$: reasoning traces of all previously executed actions.
879
- 880 • $\tau = (a_1, \dots, a_P)$: trajectory of length P .
881
- 882 • $j = (j_1, \dots, j_L)$: structured justification of length L , consisting of explicit reasoning and a
883 final verdict.
884
- 885 • π_θ : WebArbiter model parameterized by θ .
886
- 887 • \hat{y} : predicted preference verdict.
888
- 889 • $\mathcal{D}_{\text{Train}}, \mathcal{D}_{\text{SFT}}, \mathcal{D}_{\text{RL}}$: training datasets for supervised distillation and reinforcement learning.
890
- 891 • $\mathcal{D}_{\text{Bench}}$: evaluation dataset for WEBARBITER-BENCH.
892
- 893 • a^+ : action from a successful trajectory (positive label).
894
- 895 • a^- or a^{-q} : rejected action ($q = 1, \dots, Q$).
896
- 897 • Q : number of distractors in WEBARBITER-BENCH (set to 4).
898
- 899 • M : dataset size (training or evaluation).
900
- 901 • $R(x, \hat{y})$: verifiable reward signal.
902

903 **B EXAMPLE OF PREFERENCE DATASET**
904905 To illustrate the construction of a preference pair, consider a task instruction from WEBARENA:
906

- 907 • **Instruction \mathcal{I} :** “Find the 2026 conference submission page on the ICLR website.”
908
- 909 • **Current observation o_p :** [356] banner ‘header’, role=‘banner’ [359] link ‘Home’ [380]
910 button ‘Select Year (2026)’ [386] button ‘Dates’ [391] button ‘Calls’ [396] button ‘Guides’
911 [401] button ‘Organization’ [403] heading ‘ICLR 2026’
912
- 913 • **History** ($a_{<p}, c_{<p}$):
914
 - 915 – a_1 : Search “ICLR”, c_1 : “I need to find the official ICLR website first. Let me search
916 for ‘ICLR’ to locate it.”
917
 - 918 – a_2 : Click link “ICLR homepage”, c_2 : “I can see the ICLR official website link in the
919 search results. Clicking on it will take me to the ICLR homepage where I can find the
920 conference submission information.”
921
- 922 • **Candidate actions and reasoning traces:**
923
 - 924 – (a_p^1, c_p^1) : Click link “Call for Papers”; c_p^1 : “I can see a ‘Call for Papers’ link on the
925 ICLR homepage. This link would likely lead to the submission details page, which
926 should contain information about the 2026 conference submission process that I’m
927 looking for.”
928
 - 929 – (a_p^2, c_p^2) : Click “About” link; c_p^2 : “I can see an ‘About’ link on the ICLR homepage.
930 Since I need to find the 2026 conference submission page, the ‘About’ section might
931 contain conference overview information including links to submission details or
932 important dates for the 2026 conference.”
933
- 934 • **Label y :** a_p^1 is preferred.
935

936 This example is represented in the dataset as:
937

$$(\mathcal{I}, o_p, a_{<p}, c_{<p}, (a_p^1, c_p^1), (a_p^2, c_p^2), y = a_p^1).$$

918 C TRAINING DETAILS
919920 All training is conducted on 8 NVIDIA A100-80GB GPUs with fixed random seeds. Our training
921 framework is based on LLama-Factory (Zheng et al., 2024) and VERL (Sheng et al., 2024)
922923 **Distillation Stage.** We train the model for 5 epochs with a learning rate of 8e-4, using
924 LoRA with a rank of 128. We apply a cosine learning rate scheduler with a warmup ratio of 0.1. We
925 set the batch size to 256 and the maximum sequence length to 8,192 tokens.
926927 **RLVR Stage.** We employ the VERL framework for GRPO training. The learning rate is
928 set to 7.0×10^{-6} . The training uses a fixed batch size of 1,024 with mini-batch size of 128, and
929 adopts Fully Sharded Data Parallel (FSDP) for enhanced memory efficiency. For rollout generation,
930 we deploy vLLM with tensor parallelism of 4 and GPU memory utilization limited to 0.4. Response
931 sampling uses standard parameters (temperature=1.0, top-p=1.0), generating 7 candidate responses
932 per prompt. We apply KL regularization with a coefficient of 1.0×10^{-3} and clip ratio of 0.2. The
933 maximum input sequence length is 8,192 tokens, and the maximum response length is 4,096 tokens.
934935 D PROMPT REPOSITORY
936937 WebArbiter
938

939 You are a skilled expert at evaluating assistant responses. You
940 should evaluate given responses based on the given judging
941 criteria.\nGiven the context of the conversation and two
942 responses from the Assistant, you need to determine the better
943 response. Provide an overall comprehensive comparison upon them.

944 ##### Intent #####
945 {intent}
946 ##### AXTREE #####
947 Note: [bid] is the unique alpha-numeric identifier at the
948 beginning of lines for each element in the AXTree. Always use
949 bid to refer to elements in your actions.
950 {observation}
951 ##### Trajectory #####
952 Note: The trajectory contains the sequence of previous actions and
953 their corresponding thoughts. Each entry reflects the agent's
954 internal reasoning ('thought') and the concrete operation it
955 performed ('action').
956 {trajectory}
957 ##### start url #####
958 {start_url}
959 ##### current url #####
960 The URL provides clues about the user's position in the
961 application flow. Use both the path and query parameters to
962 infer page type (e.g., homepage, search results, product
963 detail, cart, checkout).
964 {current_url}
965 ##### Assistant Responses #####
966 [The Begin of Response 1]
967 THOUGHT:
968 {thought1}
969 ACTION:
970 {action1}
971 [The End of Response 1]
972 [The Begin of Response 2]
973 THOUGHT:
974 {thought2}
975 ACTION:
976 {action2}
977 [The End of Response 2]

```

972
973     ### Output Instructions ###
974     Format your output strictly using the following XML-style tags:
975     <State>Summarize the current state based on the URL, AXTree, and
976     previous actions. Include what page the user is currently on,
977     and what relevant UI elements or information are
978     visible.</State>
979     <Criteria>Other potential criteria specific to the query and the
980     context, and the weights of each criteria.</Criteria>
981     <Analysis>Compare Response 1 and Response 2 in detail according to
982     the <State> and <Criteria>.</Analysis>
983     <Answer>Response 1 or Response 2</Answer>
984     Rules for <Answer>:
985     - If Response 1 is better, output exactly: <Answer>Response
986         1</Answer>
987     - If Response 2 is better, output exactly: <Answer>Response
988         2</Answer>
989     Important Notes:
990     - Be objective and base your evaluation strictly on the content of
991         the responses.
992     - Do not let the response order, length bias your judgment.
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

```

E CASE STUDY: WEBARBITER VS. WEBSHEPHERD

To further illustrate the differences between WebArbiter and WebShepherd, we present a representative example from WEBARENA. The task instruction is:

“What is the rating of Ugreen lightning to 3.5mm cable? Round to the nearest whole number.”

At the current step, the agent observes a search-results page listing the target product. The snippet already shows a “65%” rating in the result list, but the product detail page has not been opened yet. Two candidate actions are considered:

- a_p^1 : send_msg_to_user with “65%”.
 c_p^1 : “The list view already shows a 65% rating, so answer directly.”
- a_p^2 : Click the target product entry to open its detail page, then extract and (if needed) round the rating.
 c_p^2 : “Verify the rating on the product page (correct product, correct field) before responding.”

WebShepherd. WebShepherd evaluates candidates using *checklist*-style templates that are *precompiled* before the next observation. These checklists typically include predicates such as “verify on the product page”; hence, even when a rating is already visible in an earlier search-results snippet, the checklist *still requires verification*, and thus tends to favor a_p^2 . When the actual page deviates from the ex-ante forecast (e.g., an early results page surfaces), the precompiled predicates become stale: they enforce a verification path that may be unnecessary or even brittle under interstitials or layout drift.

WebArbiter. WebArbiter derives principles such as “*answer directly when the objective’s required field is already unambiguously satisfied by the current observation*,” “*ensure correct rounding*,” and “*avoid redundant navigation when the answer is already grounded*.” It performs *dynamic expectation alignment*: (i) it forms expectation about what evidence is needed, (ii) compares the actual page with that expectation, and (iii) revises principle weights and candidate scoring accordingly. Concretely, upon seeing a clear “65%” rating in the snippet, it downweights “must verify on product page” and upweights “answer directly with proper rounding,” issuing a preference verdict for a_p^1 and correctly completing the task with minimal steps.

Discussion. This case illustrates a key limitation of precompiled, open-loop checklists: they conflate *procedural requirements* (“must navigate to detail page”) with *goal satisfaction* and thus underperform

1026 when early observations already satisfy the objective. In contrast, WebArbiter grounds decisions
 1027 in explicit, principle-guided reasoning *and* closed-loop, dynamic expectation alignment (predict →
 1028 observe → compare → revise), enabling it to act on already-sufficient evidence and remain robust to
 1029 goal–observation mismatches.

1031 F BENCHMARK CONSTRUCTION

1032 **Positive samples.** We construct WEBPRMBENCH using the successful trajectories from AGENTREWARDBENCH, a human-verified evaluation suite that aggregates over a thousand trajectories
 1033 generated by multiple LLM-based web agents across diverse real-world environments. Each trajectory in AGENTREWARDBENCH is annotated for success and execution quality by expert annotators,
 1034 providing a reliable source of environment-grounded optimal behavior. From this dataset, we select
 1035 only those trajectories that complete each task with the minimum number of steps. Each trajectory is
 1036 independently reviewed by annotators to ensure monotonic progress and to verify that no redundant
 1037 or detour actions are present. When deviations are identified, annotators revise the trajectory to
 1038 recover the shortest valid execution path consistent with successful task completion. For consistency,
 1039 missing reasoning traces are completed to ensure that every state–action pair is paired with a coherent
 1040 rationale. The resulting actions from these validated minimal-step trajectories serve as positive labels,
 1041 reflecting actions empirically verified to succeed in the real web environment.

1042 **Negative samples.** For each state, we sample four alternative actions and their associated reasoning
 1043 from a diverse ensemble of policy models, covering both open-source and proprietary LLMs. The
 1044 pool includes high-capacity instruction-tuned models such as Qwen2.5-7B / 72B-Instruct (Qwen et al.,
 1045 2025), Llama-3.3-8B / 70B-Instruct (Grattafiori et al., 2024), as well as frontier commercial models
 1046 including GPT-4o / 4o-mini (OpenAI, 2024a;b), Claude-3.5-Haiku / Claude-3.7-Sonnet (Anthropic,
 1047 2024b; 2025), and Gemini-2.5-Flash / Gemini-2.5-Pro (Comanici et al., 2025). This ensures that
 1048 alternative actions exhibit broad stylistic and policy diversity rather than reflecting any single model’s
 1049 reasoning behavior. Since alternative actions may still succeed under certain web interfaces, we
 1050 apply a rule-based filtering procedure to remove actions that remain potentially valid. We retain
 1051 only actions that are clearly invalid or non-progressing, ensuring that negative samples correspond
 1052 to failures under the actual environment dynamics rather than differences in reasoning style. To
 1053 ensure consistency and avoid false negatives, the filtered actions are manually reviewed, and any
 1054 remaining actions that appear potentially valid are discarded. If more than four valid rejected actions
 1055 remain after filtering, we randomly sample a subset to maintain a consistent number of action pairs
 1056 per instance. All rationales are truncated to a fixed length to reduce formatting noise while preserving
 1057 semantic content.

1058 The final benchmark consists of 1,287 step-level preference pairs across four environments, each
 1059 containing one environment-verified positive action and four rule-filtered negative alternatives.

1060 G GENERAL-DOMAIN GENERATIVE REWARD MODELS AND THEIR 1061 TRANSFERABILITY TO WEB TRAJECTORIES

1062 This section presents additional evaluations of general-domain generative reward models, including
 1063 RM-R1 (Chen et al., 2025b), RRM (Guo et al., 2025b), Think-RM (Hong et al., 2025), and R3 (Anu-
 1064 graha et al., 2025). Although these models represent SOTA approaches within preference-based
 1065 reward modeling, they are trained primarily on static QA, dialogic reasoning, mathematical problem
 1066 solving, and related preference datasets. Consequently, their training objectives do not incorporate
 1067 key structural properties of interactive web environments, such as AXTree-grounded observations.

1068 To enable comparison, we adapt each model’s preference interface for step-level scoring on
 1069 WEBPRMBENCH. As shown in Tab. 5, all general-domain reward models achieve substantially lower
 1070 *Avg. Pairwise Acc* and *Avg. BoN Acc* than WebArbiter. Because these models are trained exclusively
 1071 on static, text-only preference corpora, their learned reward functions emphasize linguistic plausibility
 1072 and abstract reasoning rather than the procedural validity required for web actions. They do not model
 1073 the environment-dependent factors that govern real web interaction, such as action executability under
 1074 the current state, UI structural changes, and whether an action produces measurable task progress,
 1075 making them fundamentally mismatched to process-level reward modeling. Overall, these results

1080
 1081 Table 5: Results on WEBPRMBENCH with *Pairwise* and *BoN Acc*. \star denotes our models. Bold numbers
 1082 indicate the best results, while underlined numbers denote the second best. Our WebArbiter-7B
 1083 achieves the highest *BoN Acc* across all environments, with an Avg. *BoN Acc* of 77.78%, outperforming
 1084 the second-best baseline, i.e., Gemini Flash, by 10.85%.

1085 1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133	1084									
	1085 1086 1087 1088 1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133									
Models	Mind2Web		WebArena		AssistantBench		WorkArena		Avg.	
	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN	Pairwise	BoN
<i>LLM-as-judge, Proprietary Language Models</i>										
GPT-4o-mini	80.28	<u>45.69</u>	81.40	61.46	88.79	<u>68.97</u>	84.86	66.67	83.83	60.70
GPT-4o	80.60	52.84	79.78	56.06	83.62	58.62	85.69	66.67	82.42	58.55
GPT-5	82.06	62.26	76.89	55.26	76.72	58.62	80.83	65.56	79.12	60.42
Claude Sonnet	81.19	59.60	81.40	62.53	76.92	53.85	<u>88.06</u>	72.22	81.89	62.05
Gemini Flash	81.94	57.86	<u>83.89</u>	67.39	<u>84.62</u>	65.38	92.36	<u>77.08</u>	<u>85.70</u>	<u>66.93</u>
DeepSeek-R1	82.18	56.44	81.13	60.38	75.00	51.72	88.61	72.78	81.73	60.33
<i>LLM-as-judge, Open-source Language Models</i>										
Qwen2.5-3B-Instruct	76.89	37.66	69.54	35.04	80.17	48.28	70.56	42.22	74.29	40.80
Qwen2.5-7B-Instruct	79.02	41.00	72.44	40.16	76.72	37.93	78.06	51.67	76.56	42.69
Llama-3-70B-Instruct	79.43	50.96	77.90	56.60	78.85	46.15	87.50	70.14	80.92	55.96
<i>Generative RMs</i>										
RM-R1-Qwen2.5-Instruct-7B	69.11	23.77	63.68	20.22	73.28	34.48	62.36	18.33	67.11	24.20
RRM-7B	82.28	48.51	74.60	48.25	86.21	68.97	77.78	55.56	80.22	55.32
Think-RM-3B	70.93	28.57	60.44	26.42	62.93	24.14	68.75	33.89	65.76	28.25
Think-RM-8B	75.45	45.54	74.33	48.79	77.59	51.72	83.47	63.33	77.71	52.35
R3-Qwen3-4B-LoRA-4k	78.32	43.56	76.48	48.25	83.19	62.78	83.19	62.78	80.30	54.34
<i>WebPRMs (3B)</i>										
WebShepherd-3B	37.41	21.22	20.33	9.47	36.54	17.24	10.49	2.44	26.19	12.59
★ WebArbiter-3B	<u>93.28</u>	<u>78.75</u>	83.29	56.87	76.72	44.83	84.03	60.56	84.33	60.25
<i>WebPRMs (7B+)</i>										
WebShepherd-8B	72.35	46.68	33.16	12.37	55.77	44.83	35.85	12.68	49.28	29.14
★ WebArbiter-7B	96.47	90.07	84.30	71.43	80.17	72.41	87.36	77.22	87.08	77.78

1106
 1107 show that general-domain Generative RMs do not generalize to procedural, state-dependent web
 1108 tasks and highlight the need for domain-grounded reasoning and environment-verified supervision in
 1109 WebPRMs.

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188

I Instruction
What is the rating of Ugreen lightning to 3.5mm cable?
Round to the nearest whole number.

($a_{<p}, r_{<p}$) Trajectory
Thought: On One Stop Market's homepage, the search ...
Action: type [2826] "Ugreen lightning to 3.5mm cable"
Thought: The current Page is on

(a_p^1, r_p^1) Action Candidate 1
Thought: The page lists a 65% rating on the search results snippet.
Action: send_msg_to_user ["65%"]

(a_p^2, r_p^2) Action Candidate 2
Thought: Open the product page to verify the rating, then round to nearest whole.
Action: click [6368]

[My Account](#)
[My Wish List](#)
[Sign in](#)
Welcome to One Stop Market
[Create an Account](#)

One Stop Market

Search

[Beauty & Personal Care](#)
[Sports & Outdoors](#)
[Clothing, Shoes & Jewelry](#)
[Home & Kitchen](#)
[Office Products](#)

[Tools & Home Improvement](#)
[Health & Household](#)
[Patio, Lawn & Garden](#)
[Electronics](#)
[Cell Phones & Accessories](#)

[Video Games](#)
[Grocery & Gourmet Food](#)

[Home](#)
Search results for: 'Ugreen lightning to 3.5mm cable'

Search results for: 'Ugreen lightning to 3.5mm cable'

Shop By
Items 1-12 of 8232
Sort By

Did you mean
ugreen lightning to 35mm cable

Related search terms
ugreen



UGREEN Lightning to 3.5mm Adapter Apple MFI Certified Headphone Adapter for iPhone, in-Ear Lightning to Audio Aux Jack Dongle Converter Compatible with iPhone 13 Mini 12 Pro SE 11 Pro Max X XR XS 8 8P; Black

★★★★★ 12 Reviews



UGREEN Wired Lightning Headphones MFI Certified for iPhone, in-Ear Earphones with Microphone and Volume Control Wired Lightning Earbuds Compatible with iPhone 13 Mini Pro Max 11 Pro Max, Black

★★★★★ 3 Reviews



UGREEN 3.5mm Audio Cable Stereo Auxiliary AUX Cord Gold-Plated Male to Male Braided Cable for Car Home Stereo, Headphone, Speaker, MP3, Compatible with Beats, iPhone, iPod, iPad, Sony (Black, 1.5ft)

\$6.99



UGREEN RCA Cable 2RCA Male to 2RCA Male Stereo Audio Cable Gold Plated Adapter Dual Shielded RCA Cord Compatible with Home Theater Amplifier Hi-Fi System Car Audio Speaker, 15FT

★★★★★ 12 Reviews

WebShepherd Analysis

✗ Action Candidate 1

Checklist 1: Search for Product (Yes)
Results for 'Ugreen lightning to 3.5mm cable' are visible; the list view already shows a 65% rating, so I can answer directly.

Checklist 2: Extract Rating (No)
Outputs '65%' from the results snippet without locating a rating field or confirming the rating on the product page.

Checklist 3: Verify Page (No)
Does not open the product detail page; no on-page verification or source evidence is performed.

✓ Action Candidate 2

Checklist 1: Search for Product (Yes)
Results are visible, but the correct next step is to open the product page to verify the rating.

Checklist 2: Extract Rating (In Progress)
Plan to locate the rating on the product page and then round, the action moves toward the rating field, but the value has not yet been located or parsed at this step.

Checklist 3: Verify Page (No)
Does not open the product detail page; no on-page verification or source evidence is performed.

WebArbiter Analysis

Principle

1. Correctness of the rating value (60%)
2. Correct rounding to the nearest whole number (20%)
3. Reference to the right product (not a different product) (10%)
4. Action appropriateness (should not click a wrong element) (10%)

Comparing against the principle:

1. Correct rating – both fine.
2. Rounding – both fine.
3. Correct product
 - Action Candidate 1 implicitly correct
 - Action Candidate 2 may be off.
4. Action appropriateness
 - Action Candidate 1 minimal

Conclusion:

✓ Action Candidate 1 — more accurate and progressive (Chosen)

✗ Action Candidate 2 — may click the wrong element

Figure 3: Case study on product rating in WebArena-Lite. The snippet shows “65%” before opening the product page. WebShepherd, constrained by fixed checklists, may enforce redundant verification. WebArbiter, using principle-guided reasoning, recognizes the snippet as sufficient and selects the correct action.