


# DynaWeb: Model-Based Reinforcement Learning of Web Agents


Hang Ding<sup>1</sup> Peidong Liu<sup>2</sup> Junqiao Wang<sup>1</sup> Ziwei Ji<sup>3</sup> Meng Cao<sup>4 8</sup> Rongzhao Zhang<sup>5</sup>  
Lynn Ai<sup>6</sup> Eric Yang<sup>6</sup> Tianyu Shi<sup>6</sup> Lei Yu<sup>7</sup>

<sup>1</sup>Shanghai Jiao Tong University <sup>2</sup>Sichuan University <sup>3</sup>Hong Kong University of Science and Technology  
<sup>4</sup>McGill University <sup>5</sup>Shanghai AI Lab <sup>6</sup>Gradient  
<sup>7</sup>University of Toronto <sup>8</sup>Mila - Quebec AI Institute

The development of autonomous web agents, powered by **Live Lough Moves** (LLMs) and reinforcement learning (RL), represents a significant step towards general-purpose AI assistants. However, training these agents is severely hampered by the challenges of interacting with the live internet, which is inefficient, costly, and fraught with risks. Model-based reinforcement learning (MBRL) offers a promising solution by learning a *world model* of the environment to enable simulated interaction. This paper introduces DynaWeb, a novel MBRL framework that trains web agents through interacting with a web world model trained to predict naturalistic web page representations given agent actions. This model serves as a synthetic web environment where an agent policy can *dream* by generating vast quantities of rollout action trajectories for efficient online reinforcement learning. Beyond free policy rollouts, DynaWeb incorporates real expert trajectories from training data, which are randomly interleaved with on-policy rollouts during training to improve stability and sample efficiency. Experiments conducted on the challenging WebArena and WebVoyager benchmarks demonstrate that DynaWeb consistently and significantly improves the performance of state-of-the-art open-source web agent models. Our findings establish the viability of training web agents through imagination, offering a scalable and efficient way to scale up online agentic RL.

 **Date:** Jan 30, 2026

 **Correspondence:** jadeleyiu@cs.toronto.edu, ty.shi@mail.utoronto.ca

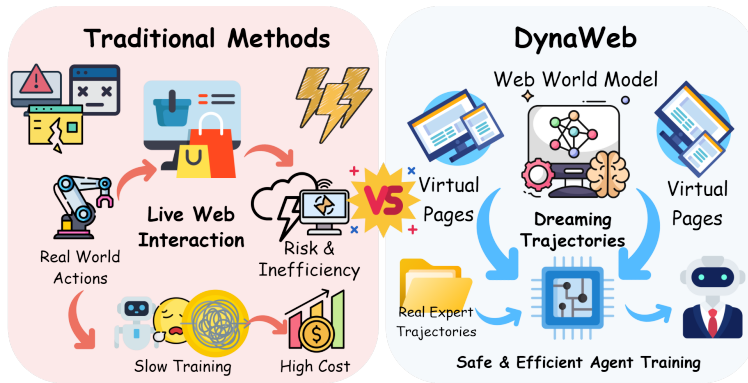
 **Project Leader:** ty.shi@mail.utoronto.ca

 **Project Page:** [Coming-Soon](#)

## 1 Introduction

The paradigm of artificial intelligence is rapidly shifting toward proactive, agentic systems that can autonomously execute complex, long-horizon tasks in open-ended environments. Large Language Models (LLMs) have emerged as a powerful backbone for such agents, enabling rich reasoning, flexible action generation, and natural language interaction. In the web domain, LLM-based agents have demonstrated strong capabilities in navigating real websites and accomplishing user-specified goals through multi-step interaction, fueled by advances in prompting, structured reasoning, and action abstractions [Yao et al., 2023, Zhou et al., 2024a, He et al., 2024]. Beyond imitation learning, recent work has shown that training web agents with online reinforcement learning (RL) can substantially improve robustness, exploration, and long-horizon decision-making, allowing agents to surpass static demonstrations and adapt through trial-and-error interaction [Wei et al., 2025, Qi et al., 2025, Zhuang et al., 2025].

Despite its promise, the effectiveness of online RL for web agents is fundamentally constrained by the cost and risk of real-environment interaction. Collecting large-scale, on-policy experience requires agents to interact directly with the live internet, which is inefficient, expensive, and difficult to control at scale.



**Figure 1.** Comparison between traditional web agent training via live web interaction and DynaWeb. By replacing risky and inefficient real-world interaction with a learned web world model, DynaWeb enables imagination-driven training using virtual pages and dreamed trajectories, optionally augmented with real expert data, resulting in safer and more efficient agent optimization.

rogate that can faithfully approximate web dynamics. To this end, recent work has begun to explore *web world models*—learned simulators of web environments. So far, however, their role has been largely auxiliary. Some approaches employ world models at inference time for short-horizon look-ahead and action evaluation, treating them as decision-time reasoning tools rather than learning substrates [Gu et al., 2024, Wei et al., 2025]. Others use world models to synthesize offline trajectories for supervised fine-tuning or imitation-style training, decoupling model-generated experience from on-policy optimization [Fang et al., 2025a, Pahuja et al., 2025]. What remains missing is an *online* model-based reinforcement learning (MBRL) paradigm for web agents, in which imagined rollouts serve as first-class experience that directly augments on-policy policy optimization.

In this work, we revisit classical model-based reinforcement learning through the lens of modern web agents. Inspired by the Dyna architecture [Sutton, 1991] and imagination-based learning frameworks such as Dreamer [Hafner et al., 2020], we propose **DynaWeb**, an MBRL framework that elevates a web world model from a planning or data-generation tool to a core component of online reinforcement learning. DynaWeb treats the world model as a controllable synthetic web environment that can replace or augment costly real interaction. Concretely, an LLM-based web world model functions as a learned web server: conditioned on the current page representation and an agent action, it predicts realistic next-state page representations and provides task-level feedback signals for policy optimization. By training web agents on a mixture of real and imagined experience, DynaWeb enables scalable, on-policy reinforcement learning through imagination while preserving the benefits of interactive learning.

Crucially, DynaWeb combines two complementary sources of training experience. In addition to policy-driven imagined rollouts generated by the web world model, we directly incorporate fully real expert trajectories sampled from existing training data. These expert trajectories are entirely independent of the world model and correspond to ground-truth web interactions. During training, real expert trajectories are randomly interleaved with imagined rollouts and real on-policy interaction, allowing the agent to benefit from high-quality demonstrations while substantially reducing reliance on costly live web exploration. This simple but effective interleaving strategy preserves the on-policy learning signal and enables efficient online reinforcement learning with significantly fewer real-environment interactions.

Our contributions are summarized as follows:

- We train a web world model that predicts naturalistic web page state transitions in the form of structured accessibility tree representations, enabling realistic simulation of web environment dynamics

During training, agents may trigger irreversible actions, such as unintended purchases, account modifications, or data submissions, while also facing non-deterministic page dynamics, transient failures, and external interference. These challenges severely limit the practicality of pure online RL, rendering large-scale policy optimization both costly and hazardous in real-world web environments [Zhou et al., 2024a, Qi et al., 2025]. As a result, a central open question emerges: how can we retain the benefits of online reinforcement learning for web agents while dramatically reducing reliance on direct interaction with the live web? A natural direction is to replace expensive and risky real-environment interaction with a learned, controllable sur-

without live web interaction.

- We show that policy-driven imagined rollouts generated by a learned web world model can be directly used as on-policy training experience for reinforcement learning of web agents.
- We propose **DynaWeb**, a model-based reinforcement learning framework that integrates fully real expert trajectories from training data with imagined rollouts from a world model, enabling effective policy optimization without interacting with the live web.
- Through extensive experiments and analysis, we demonstrate consistent performance improvements on WebArena and WebVoyager, and provide empirical insights into the design and use of world-model-based imagination for training web agents.

## 2 Related work

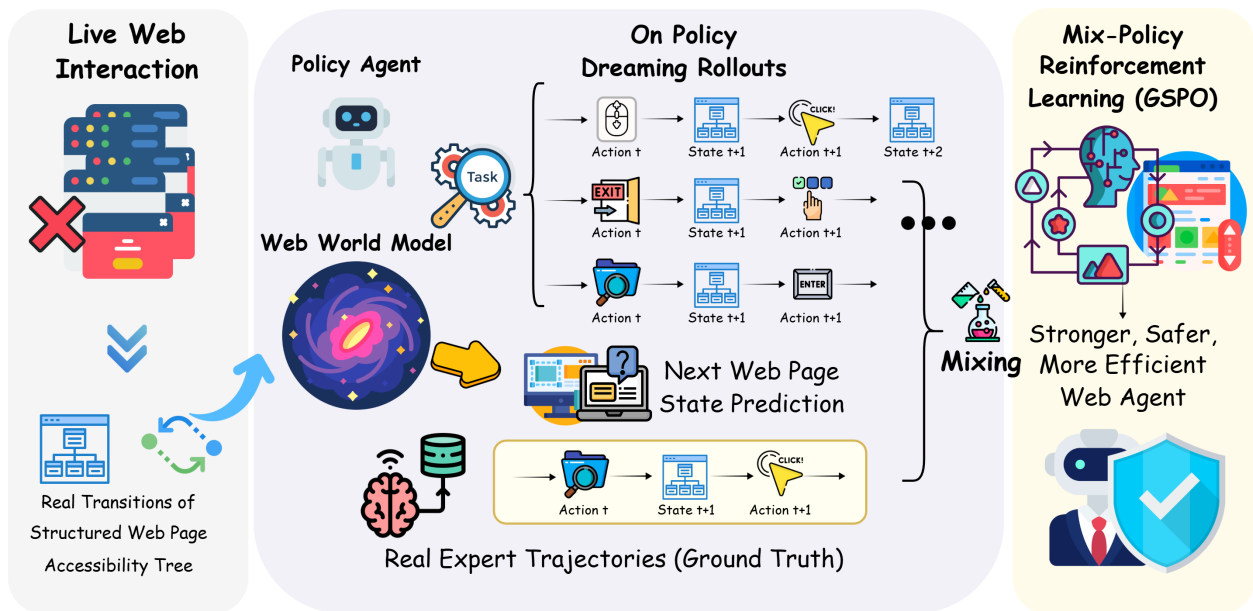
**Web Agent** Recent advances in web agents are largely driven by (multimodal) large language models (LLMs) serving as the core decision-making backbone [Dubey et al., 2024, Jia et al., 2024, OpenAI, 2023, Anthropic, 2025]. On top of these models, reasoning and interaction frameworks such as ReAct Yao et al. [2023], MCP [Anthropic, 2024], and Cognitive Kernel [Zhang et al., 2024a] enable structured multi-step web actions. Web agents are commonly evaluated on interactive benchmarks including WebShop [Yao et al., 2022], Mind2Web [Deng et al., 2023], WebArena [Zhou et al., 2024a], VisualWebArena [Koh et al., 2024a], WebVoyager [He et al., 2024], WebWalker [Wu et al., 2025a], and MMInA [Zhang et al., 2024b].

Beyond off-the-shelf prompting, a broad line of work improves web agents via data scaling and stronger agent training pipelines. Data-centric efforts such as Explorer [Pahuja et al., 2025], NNetNav [Murty et al., 2025], and InSTA [Trabucco et al., 2025] collect or synthesize high-quality interaction data, while recent agent foundation models and deep research agents push toward more general web interaction at scale, e.g., WebThinker [Li et al., 2025a], WebDancer [Wu et al., 2025b], WebSailor [Li et al., 2025b], WebShaper [Tao et al., 2025], Cognitive Kernel-Pro [Fang et al., 2025b], MiroFlow [MiroMindAI, 2025], and multimodal research agents such as WebWatcher [Geng et al., 2025]. In parallel, inference-time optimization (e.g., tree search and self-reflection) further improves decision-making without additional training [Koh et al., 2024b, Putta et al., 2024, Yu et al., 2024, Zhou et al., 2024b, Zhang et al., 2024c, Shinn et al., 2023].

**World Models** World models originate from model-based reinforcement learning [Ha and Schmidhuber, 2018] and have recently been revisited as a powerful abstraction for agent reasoning and planning [Valevski et al., 2024, Alonso et al., 2024, Smith and Wellman, 2023]. With the emergence of large language models (LLMs), several works treat LLMs as implicit world models capable of simulating future states. For general reasoning, RAP [Hao et al., 2023] integrates LLM-based world modeling with Monte Carlo Tree Search to explore hypothetical trajectories, while WKM [Qiao et al., 2024] distills structured world knowledge from agent trajectories to guide planning.

In web environments, this paradigm has been adapted by methods such as WebDreamer [Gu et al., 2024] and WMA [Chae et al., 2025], which use LLMs to predict the outcomes of web actions via natural language simulation. However, these approaches primarily employ world models at inference time, functioning as planning or prompting modules to assist action selection. Even when multi-step simulation is performed, the imagined rollouts are used only to guide immediate decisions rather than to improve the agent policy itself through training.

In contrast, our work positions the web world model as a core component of the learning process. By training a dedicated world model and using it to generate multi-step imagined trajectories for policy optimization,



**Figure 2. Overview of DynaWeb.** DynaWeb trains web agents via imagination-driven, model-based reinforcement learning. A learned web world model serves as a synthetic environment, enabling the agent to generate multi-step imagined rollouts without interacting with the live web. These imagined trajectories are mixed with a small fraction of real expert trajectories to stabilize learning. The agent policy is optimized using sequence-level policy optimization, allowing efficient and robust credit assignment for long-horizon web tasks with sparse terminal rewards.

DynaWeb moves beyond inference-time reasoning and enables genuine model-based reinforcement learning for web agents. This design directly addresses the data inefficiency and cost of real-environment interaction highlighted in prior RL-based web agents, and aligns with our goal of scalable, imagination-driven agent training.

**Reinforcement Learning of Agents** Reinforcement learning (RL) provides a direct mechanism to improve agent policies from interactive feedback, which is particularly appealing for web environments with long-horizon, multi-step decision making [Wei et al., 2025, Qi et al., 2025, Wu et al., 2025b, Zhuang et al., 2025]. A representative end-to-end online RL approach is WebAgent-R1 [Wei et al., 2025], which optimizes multi-turn web interaction policies using outcome-based rewards and scalable trajectory sampling (e.g., multi-group GRPO [Shao et al., 2024]). Complementary to direct end-to-end optimization, WebRL [Qi et al., 2025] emphasizes self-evolving curriculum design and result-supervised feedback to continually generate training tasks and improve agent robustness.

Other methods adapt RL objectives or combine RL with additional supervision to better shape reasoning and planning behaviors. WebDancer [Wu et al., 2025b] leverages structured supervision from QA-style signals and policy optimization (e.g., DAPO [Yu et al., 2025]) to internalize actionable reasoning traces. WorkForceAgent-R1 [Zhuang et al., 2025] integrates behavior cloning with GRPO-style optimization to jointly enhance single-step reasoning and multi-step planning. Beyond pure online RL, joint training pipelines that combine SFT and RL have also shown effectiveness. For instance, AutoWebGLM [Lai et al., 2024] integrates SFT, RL, and rejection sampling fine-tuning, with curriculum learning to bootstrap basic web navigation skills before subsequent RL refinement.

Despite their strong results, these RL-based agents typically rely on large amounts of real-environment interaction, which can be expensive, unstable, and risky at scale—a key motivation for our imagination-driven training framework.

## 3 Method

### 3.1 Problem Formulation

We formulate the web agent task as a Partially Observable Markov Decision Process (POMDP)  $(\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R})$ . Given a natural language query  $q$  and a target website  $w$ , the agent is required to complete a multi-step web interaction to satisfy the task objective. The state space  $\mathcal{S}$  represents the full underlying web environment, which is not directly observable by the agent. Instead, at each time step  $t$ , the agent receives a partial observation  $o_t \in \mathcal{O}$  obtained via an observation function  $o_t = \Omega(s_t)$ , where  $\Omega(\cdot)$  extracts visible information such as the current URL and web elements from state  $s_t$ .

The action space  $\mathcal{A}$  consists of atomic browser operations, including `click`, `type`, `goback`, `scroll up/down`, and `stop`. Executing an action  $a_t \in \mathcal{A}$  deterministically advances the environment according to the transition function  $\mathcal{T}$ :

$$s_{t+1} = \mathcal{T}(s_t, a_t), \quad o_{t+1} = \Omega(s_{t+1}). \quad (1)$$

Following prior work [Zhang et al., 2024a], we represent each observation  $o_t$  using an accessibility tree that captures the structured layout of visible web elements. The agent is parameterized by a large language model (LLM) policy  $\pi_\theta$ , which conditions on the task query  $q$ , the interaction history, and system instructions  $I$  to generate reasoning traces and actions. At time step  $t$ , the agent samples a chain-of-thought  $h_t$  together with an action  $a_t$  as

$$(h_t, a_t) \sim \pi_\theta(\cdot \mid I, q, o_{1:t}, h_{1:t-1}, a_{1:t-1}). \quad (2)$$

An interaction episode produces a trajectory  $\tau = (o_1, h_1, a_1, \dots, o_T, h_T, a_T)$ , where  $T$  denotes the total number of interaction steps. Upon termination, the trajectory is assigned a scalar reward  $\hat{r}(\tau, q) \in [0, 1]$  based on task completion, obtained via self-assessment or environment-specific evaluation.

### 3.2 Web World Model

We introduce a web world model that approximates the transition dynamics of web environments by predicting how visible web states evolve in response to agent actions. Given the partial observation  $o_t$  represented as an accessibility tree and an action  $a_t \in \mathcal{A}$ , the world model predicts the next observation  $\hat{o}_{t+1}$ :

$$\hat{o}_{t+1} \sim p_\phi(\cdot \mid o_t, a_t, I), \quad (3)$$

where  $p_\phi$  is a large language model parameterized by  $\phi$ . Rather than learning latent dynamics, the model operates directly in the observation space, generating naturalistic web page representations that can be directly consumed by the agent policy during imagined rollouts.

As shown in previous studies, simply using textual observations to represent web environment states and use them as training objectives would be problematic, as state transitions in websites often involve modifying only a part of the previous observation, so  $o_{t+1}$  often remains highly similar to  $o_t$ , and predicting the entire next textual observation may result in low information gain. We therefore decompose the task of next web state prediction into two subtasks: first, we ask the world model to predict a set of free-from natural language description of *state changes*  $\Delta(o_t, o_{t+1})$  resulted by the action  $a_t$ , and then apply  $\Delta(o_t, o_{t+1})$  on the current state  $o_t$  to get  $o_{t+1}$ . We could therefore only train the world model on the first reasoning subtask of predicting state change descriptions, and leverage its instruction following capability to perform the second subtask of state altering during inference.

The world model is trained using real web interaction trajectories collected from the StanfordNLP/NNetNav [Murty et al., 2025] dataset. We apply a data cleaning pipeline to filter out erroneous or incomplete trajectories, including missing observations, invalid actions, or inconsistent state transitions. We then prompt GPT-oss-120b with each valid transition  $(I, o_t, a_t, o_{t+1})$  to obtain state change descriptions  $\Delta(o_t, o_{t+1})$  as well as the model reasoning trace  $r$  before generating  $\Delta$ . The world model LLM is trained as a reasoning model to predict the ground-truth reasoning trace as well as the subsequent state changes, conditioned on the current accessibility tree and the executed action:

$$\mathcal{L}_\phi = \sum_{(I, o_t, a_t, r, \Delta)} -\log p_\phi(r, \Delta | I, o_t, a_t) \quad (4)$$

This learned world model can therefore serve as a reusable simulator for generating multi-step imagined trajectories, which we leverage for model-based reinforcement learning in the following section.

### 3.3 DynaWeb: Model-based RL of Web Agents

DynaWeb trains web agents through model-based reinforcement learning by relying on imagined rollouts generated by a learned web world model, eliminating the need for live web interaction during training. Given a task query  $q$  and an initial observation  $o_1$ , the agent policy  $\pi_\theta$  interacts with the world model as a synthetic environment to produce multi-step trajectories. At each step  $t$ , the agent samples an action

$$a_t \sim \pi_\theta(\cdot | o_{1:t}, h_{1:t-1}, a_{1:t-1}, q), \quad (5)$$

and the world model predicts the next observation

$$\hat{o}_{t+1} \sim p_\phi(\cdot | \hat{o}_t, a_t, q), \quad \hat{o}_1 = o_1. \quad (6)$$

Iterating this process yields an imagined trajectory

$$\hat{\tau} = (\hat{o}_1, h_1, a_1, \dots, \hat{o}_T, h_T, a_T), \quad (7)$$

without interacting with the live web environment.

To enable reinforcement learning on imagined trajectories, a task-level reward signal is obtained via model-based self-assessment conditioned on the task query. After trajectory termination, a task-level completion reward  $\hat{r}(\hat{\tau}, q) \in \{0, 1\}$  is obtained via model-based self-assessment with respect to the task query, indicating whether the task objective has been achieved.

We treat the return as  $G(\hat{\tau}) = \hat{r}(\hat{\tau}, q)$  and use the resulting imagined rollouts as valid training data for policy gradient optimization, substantially reducing reliance on real-environment interaction while preserving task-level supervision.

Beyond policy-driven imagined rollouts, DynaWeb additionally incorporates fully real expert trajectories sampled from training data to improve training stability and sample efficiency. These expert trajectories correspond to ground-truth web interactions and are entirely independent of the world model. In practice, real expert trajectories are randomly interleaved with policy-driven imagined rollouts during training, allowing the agent to benefit from high-quality demonstrations while learning under a simulated environment.

To optimize the agent policy from this mixture of real and imagined experiences, we adopt Group Sequence Policy Optimization (GSPO) [Zheng et al., 2025]. Concretely, GSPO lifts importance sampling from the token level to the rollout (sequence) level by assigning each trajectory a single sequence-level ratio, which is then applied uniformly to all tokens in the trajectory. Let  $\{\hat{\tau}^i\}_{i=1}^G$  denote a group of rollouts sampled from  $\pi_{\theta_{\text{old}}}$  (either imagined via the world model or real expert trajectories), and let  $y^i$  be the serialized token

sequence produced by the policy along  $\hat{\tau}^i$  (i.e., concatenating all generated reasoning traces and actions). GSPO optimizes the clipped objective

$$\mathcal{J}_{\text{GSPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{\hat{\tau}^i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \left[ \frac{1}{G} \sum_{i=1}^G \min\left(s^i(\theta) \hat{A}^i, \text{clip}(s^i(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}^i\right) \right] \quad (8)$$

where  $\hat{A}^i$  denotes the (trajectory-level) advantage computed from the terminal return  $G(\hat{\tau}^i)$ . The sequence-level ratio  $s^i(\theta)$  is defined as the geometric mean of token-wise likelihood ratios:

$$s^i(\theta) = \left( \frac{\pi_{\theta}(y^i | q, o_1)}{\pi_{\theta_{\text{old}}}(y^i | q, o_1)} \right)^{1/|y^i|} = \exp\left(\frac{1}{|y^i|} \sum_{k=1}^{|y^i|} \log r_k^i(\theta)\right), \quad (9)$$

with token-level ratios

$$r_k^i(\theta) = \frac{\pi_{\theta}(y_k^i | q, o_1, y_{<k}^i)}{\pi_{\theta_{\text{old}}}(y_k^i | q, o_1, y_{<k}^i)}. \quad (10)$$

## 4 Experiments

### 4.1 Setups

**Benchmarks and metrics** We evaluate our DynaWeb agents against baselines on two challenging benchmarks of web navigation tasks: 1) WebArena [Zhou et al., 2024a], which includes 812 real-life tasks in simulated web environments across five different websites and four application domains; 2) WebVoyager [He et al., 2024], a semi-automatically generated dataset comprising 643 open-ended web tasks from 15 commonly accessed live websites. The main evaluation metric, Success Rate (SR), is calculated as the percentage of the user instructions that are successfully accomplished by the generated agent trajectory. For WebArena, to ensure fair comparison and reproducibility, we conduct our experiments using the official environment hosted on an Amazon Web Services (AWS) EC2 instance pre-configured with a Docker environment. For WebVoyager, since some live websites became no longer accessible in our evaluation web environment, either due to geographical locations or IP blocks, we filter out these failed tasks for our experiments. To ensure robustness, we follow [Fang et al., 2025a] by running our experiments roughly at the same time window twice and report the average results.

**Web Agents** We use a state-of-the-art open-source NNetNav web agent family by [Murty et al., 2025] fine-tuned from a Llama-3.1-8B as the backbone of our DynaWeb agent. NNetNav provides two web agent models trained live and simulated web environments via supervised fine-tuning datasets of web navigation trajectory. For each of our two evaluation tasks, we take the corresponding NNetNav agent as the initialization checkpoint, and train on the DynaWeb model-based RL environment simulated by a WWM. During training, we implement expert-conditioned rollouts by randomly replacing 50% of dreamed trajectories with gold-standard trajectories sampled from the NNetNav supervised fine-tuning datasets.

**World Models** For web world models, we fine-tune GPT-oss-120b using the curated next web state prediction datasets as explained in section 3, resulting in a live WWM for WebVoyager and a simulated WWM for WebArena. Each WWM is trained via supervised fine-tuning for one epoch using the LlamaFactory library. During GSPO rollouts, to reduce the effect of WWM hallucinations, we generate synthetic multi-step trajectories up to 5 steps, with early termination if the world model generates a terminal state. To

**Table 1.** Success rate (%) on self-hosted websites in WebArena. Results are reported for five representative website domains (Reddit, GitLab, Maps, CMS, and Shopping). DynaWeb achieves the highest average success rate and consistently outperforms strong supervised, RL, and inference-time lookahead baselines across most domains. Best results are in bold, and second-best are underlined.

<i>Self-hosted Websites (WebArena)</i>						
<i>Method</i>	Reddit	Gitlab	Map	CMS	Shopping	<i>Average SR</i>
Llama3.1-8B-Instruct CoT	0.0	3.3	3.3	2.9	11.1	4.1
GPT-4o	10.5	10.0	<b>20.0</b>	20.0	11.1	14.3
Qwen2.5-32B	10.5	20.0	<u>19.2</u>	19.2	17.8	17.3
NNetNav	16.4	8.7	11.4	16.7	25.9	15.8
Go-Browse	30.7	15.3	17.9	25.3	22.4	22.3
WebRL	<u>35.6</u>	<u>23.5</u>	19.1	<u>30.6</u>	24.7	<u>26.7</u>
ITL	26.5	18.0	12.6	23.6	<u>31.3</u>	22.4
<b>DynaWeb-8B (Ours)</b>	<b>43.8</b>	<b>28.7</b>	17.8	<b>31.5</b>	<b>33.2</b>	<b>31.0</b>

further ensure that DynaWeb training covers web world states across various stages of navigation tasks, we generate synthetic trajectories with initial states that are randomly sampled from the trajectory data in the NNetNav-live or NNetNav-wa SFT dataset (i.e., including both initial and intermediate states of gold-standard trajectories).

**Baselines** We evaluated DynaWeb against several representative baselines, described as follows. 1) *Base models*: We include an open-source base agent built on Llama-3.1-8B-Instruct with vanilla chain-of-thought prompting, which serves as a minimal reasoning-enabled baseline without additional training or planning mechanisms. 2) *Proprietary model*: We report results from GPT-4o as a strong closed-source baseline, representing the performance of state-of-the-art proprietary web agents under the same evaluation protocol. *Supervised fine-tuning (SFT)*: We consider supervised fine-tuning on offline successful web navigation trajectories. This includes the NNetNav agents of [Murty et al., 2025], trained on expert-collected demonstrations, as well as Go-Browse[Gandhi and Neubig, 2025], where successful trajectories are automatically collected via structured exploration and used for standard supervised training without reinforcement learning.

4) *Offline reinforcement learning*: The offline-RL baseline refers to Llama-3.1-8B-based agents trained using the WebRL framework by [Qi et al., 2025]. In this setting, the agent first interacts with online web environments to collect exploration trajectories, which are then scored by a trained outcome-based reward model and used for offline reinforcement learning. 5) *Inference-time lookahead (ITL)*: The ITL baseline is an Llama-3.1-8B-based agent equipped with a web world model at inference time. The agent performs inference-time policy optimization by simulating multiple candidate future states at each action step and selecting the action with the highest expected advantage.

## 4.2 Main Results

As shown in Table 1, **DynaWeb** achieves the strongest overall performance on WebArena, improving the average success rate from 26.7 (Offline-RL) to **31.0**, corresponding to a relative gain of **16.1%**. Notably, these gains are broad rather than localized: DynaWeb attains the best results on Reddit, Gitlab, CMS, and Shopping, indicating that its advantages are not tied to a particular site template or interaction pattern, but generalize across heterogeneous interfaces and task objectives. Compared with Offline-RL and ITL, whose

**Table 2.** Success rate (%) on live websites in WebVoyager. Results are reported across a diverse set of commonly accessed real-world websites. DynaWeb achieves the highest overall success rate and consistently outperforms supervised, RL, and inference-time lookahead baselines on most websites. Best results are in bold, and second-best are underlined.

Method	Live Websites (WebVoyager)												
	All Rec	Ama	Apple	ArXiv	BBC News	Camb Dict	Cour sera	ESPN	Git Hub	G. Map	GSea rch	HF	Wfram α
Llama3.1-8B-Instruct+CoT	2.4	2.5	0.0	9.5	7.6	11.8	0.0	0.0	1.2	4.3	0.0	2.1	5.2
GPT-4o	31.1	25.8	29.7	27.9	<u>32.6</u>	41.9	35.3	<b>27.3</b>	<b>26.4</b>	33.5	<u>5.3</u>	19.4	37.1
NNetNav	28.9	24.5	25.2	40.1	32.2	56.8	35.8	15.6	14.1	44.4	0.0	12.9	39.4
Go-Browse	30.4	24.1	26.1	41.8	31.0	54.9	35.9	17.6	16.2	42.1	1.8	18.7	41.2
WebRL	<u>32.6</u>	20.9	27.3	<b>50.4</b>	29.7	<u>57.5</u>	<u>36.6</u>	<u>19.8</u>	<u>18.4</u>	40.6	2.7	<u>20.1</u>	<u>43.7</u>
ITL	25.6	<u>28.7</u>	<u>30.6</u>	43.7	25.1	39.0	35.8	15.4	15.0	<u>46.5</u>	0.0	1.6	18.8
<b>DynaWeb-8B (Ours)</b>	<b>38.7</b>	<b>43.8</b>	<b>34.2</b>	<u>45.7</u>	<b>39.8</b>	<b>61.6</b>	<b>38.3</b>	<u>19.8</u>	11.2	<b>49.5</b>	<b>9.2</b>	<b>22.6</b>	<b>45.4</b>

strengths tend to manifest under specific exploration regimes, DynaWeb more consistently transforms partial and imperfect exploration signals into effective learning progress. This suggests that integrating imagined experience into on-policy optimization provides a more reliable training signal for complex, multi-step web interactions.

As shown in Table 2, this trend largely carries over to the live, open-world setting of WebVoyager. DynaWeb achieves the best performance on All Rec and outperforms competing methods on a wide range of individual websites, including Amazon, Apple, BBC News, Cambridge Dict, Coursera, Google Map, Google Search, HuggingFace, and Wolfram $\alpha$ . These results demonstrate that the benefits of model-based, imagination-driven training persist even when the environment is dynamic, partially observable, and less controlled. At the same time, the remaining gaps are informative: DynaWeb underperforms the strongest baselines on a few sites such as ArXiv and GitHub, which typically require longer-horizon planning with highly branching actions and rapidly evolving page states. Overall, the results indicate that DynaWeb offers a robust and transferable improvement for web agents, while highlighting long-horizon world modeling and highly dynamic UI handling as the primary challenges for further progress.

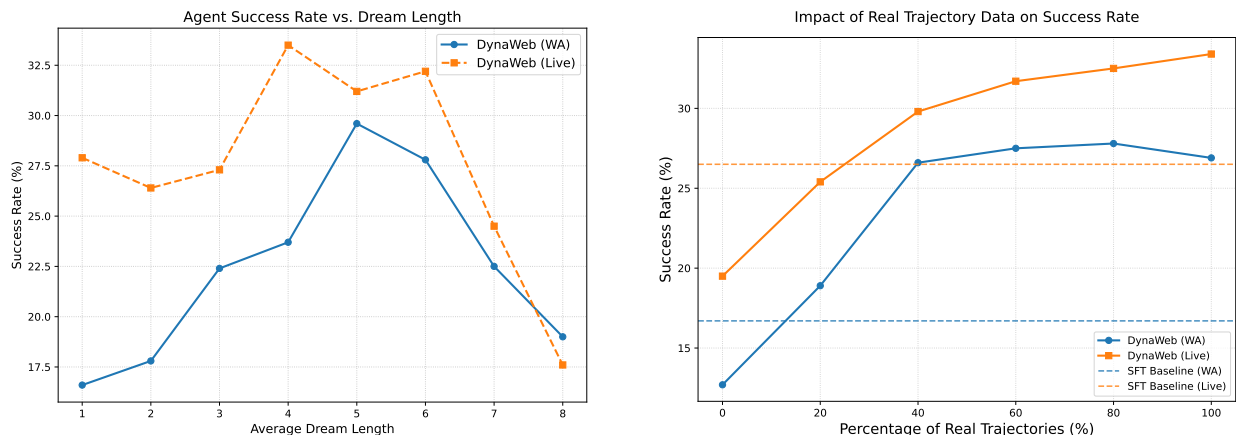
## 5 Analysis

We further analyze DynaWeb to better understand how key design choices in model-based dreaming affect training stability and overall agent performance.

### 5.1 WM Dream Length Matters

As discussed in Section 3, the length of synthetic trajectories generated through agent–world model interaction (referred to as the *dream length*) plays a central role in determining learning effectiveness. From the perspective of sequential decision making, dream length controls the effective rollout depth, i.e., the number of imagined interaction steps over which the agent can accumulate credit and propagate reward signals. This induces a fundamental trade-off: trajectories that are too short limit the agent’s ability to learn multi-step action sequences required for non-trivial web tasks, while excessively long rollouts amplify compounding world model errors and hallucinations.

To characterize this trade-off, we perform an ablation study in which DynaWeb agents are trained with



(a) Agent success rate vs. dream length. A moderate length (4–5) performs best; shorter trajectories under-complete, while longer ones compound hallucinations.

(b) Impact of real trajectory data on success rate. Adding 40% ground-truth real data significantly outperforms the SFT baseline, with diminishing returns for higher percentages.

varying maximum dream lengths. As shown in Figure 3a, we observe a clear performance peak when the average dream length is around 4–5 steps. Shorter rollouts fail to capture sufficient interaction depth for task completion, whereas longer rollouts suffer from degraded training signals due to accumulated model inaccuracies. These results validate our design choice of restricting dream depth during training, and highlight that effective model-based web RL requires balancing rollout depth against simulation fidelity. Improving the robustness of world models to support longer-horizon, high-fidelity imagined trajectories remains an important direction for future work.

## 5.2 Effect of Real Interaction Data Percentage

We analyze the role of real-environment interaction data in stabilizing and regularizing model-based web reinforcement learning. While DynaWeb relies heavily on imagined rollouts generated by the web world model, incorporating a portion of ground-truth trajectories from real environments can anchor learning to reliable state transitions and reward signals.

Figure 3b reports the average success rate of DynaWeb agents under varying proportions of real interaction data. Two clear trends emerge. First, agents trained exclusively on simulated trajectories tend to underperform the SFT baseline, indicating that world model hallucinations and compounding errors can degrade learning when unregularized. Second, introducing approximately 40% real trajectories leads to a substantial performance improvement over SFT, after which additional real data yields diminishing returns.

These results suggest that real interaction data serves as a critical regularizer for imagination-driven training: a moderate amount is sufficient to stabilize learning and correct systematic model biases, while the majority of training signal can still be efficiently obtained from simulated rollouts. This balance highlights the practical advantage of DynaWeb in reducing reliance on costly real-environment interaction without sacrificing performance.

### 5.3 Essential Role of WM Training

We assess whether explicitly training a web world model is necessary for effective model-based reinforcement learning. To isolate the contribution of world model training, we keep the agent architecture, RL algorithm, and training protocol fixed, and replace the learned DynaWeb world model with a frozen GPT-oss-120b that is prompted to predict next-step observations during imagined rollouts.

**Figure 4.** Effect of world model training on downstream agent performance. Success rate (%) comparing a supervised task-specific world model (DynaWeb WM) and a frozen general-purpose LLM (GPT-oss-120b).

Benchmark	DynaWeb WM	GPT-oss-120b
WebArena (Sim.)	<b>31.0</b>	20.9
WebVoyager (Live)	<b>35.4</b>	28.6

dynamics. Crucially, this result suggests that the effectiveness of DynaWeb does not arise merely from increased model capacity or prompting, but from grounding policy optimization in a world model explicitly trained to capture web interaction dynamics.

As shown in Table 4, agents trained with the supervised world model achieve substantially higher downstream success rates on both WebArena and WebVoyager. The large performance gap indicates that strong general-purpose LLM priors alone are insufficient to function as a reliable simulator for imagination-driven RL in web environments without learning environment-specific transition

## 6 Conclusion

We introduce **DynaWeb**, a model-based reinforcement learning framework that trains web agents through imagination rather than live web interaction. By learning a web world model and using it to generate policy-driven imagined rollouts mixed up with real expert trajectories, DynaWeb enables effective on-policy optimization while avoiding the cost and risk of real-environment interaction. Across WebArena and WebVoyager, DynaWeb consistently improves strong open-source web agents.

Our analysis reveals key principles for imagination-driven training, including the importance of rollout length, the regularizing role of real expert data, and the necessity of explicitly trained world models. These results suggest that model-based imagination offers a practical and scalable path for training web agents, and point toward world-model-centric learning as a promising direction for long-horizon decision making in complex environments.

## References

- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL [https://openreview.net/forum?id=WE\\_vluYUL-X](https://openreview.net/forum?id=WE_vluYUL-X).
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web environment for building autonomous agents. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024a. URL <https://openreview.net/forum?id=oKn9c6ytLx>.

- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. Webvoyager: Building an end-to-end web agent with large multimodal models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 6864–6890. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.371. URL <https://doi.org/10.18653/v1/2024.acl-long.371>.
- Zhepei Wei, Wenlin Yao, Yao Liu, Weizhi Zhang, Qin Lu, Liang Qiu, Changlong Yu, Puyang Xu, Chao Zhang, Bing Yin, Hyokun Yun, and Lihong Li. Webagent-r1: Training web agents via end-to-end multi-turn reinforcement learning, 2025. URL <https://arxiv.org/abs/2505.16421>.
- Zehan Qi, Xiao Liu, Iat Long Ong, Hanyu Lai, Xueqiao Sun, Wenyi Zhao, Yu Yang, Xinyue Yang, Jiadai Sun, Shuntian Yao, Tianjie Zhang, Wei Xu, Jie Tang, and Yuxiao Dong. Webrl: Training llm web agents via self-evolving online curriculum reinforcement learning, 2025. URL <https://arxiv.org/abs/2411.02337>.
- Yuchen Zhuang, Di Jin, Jiaao Chen, Wenqi Shi, Hanrui Wang, and Chao Zhang. Workforceagent-r1: Incentivizing reasoning capability in llm-based web agents via reinforcement learning, 2025. URL <https://arxiv.org/abs/2505.22942>.
- Yu Gu, Boyuan Zheng, Boyu Gou, Kai Zhang, Cheng Chang, Sanjari Srivastava, Yanan Xie, Peng Qi, Huan Sun, and Yu Su. Is your LLM secretly a world model of the internet? model-based planning for web agents. *CoRR*, abs/2411.06559, 2024. doi: 10.48550/ARXIV.2411.06559. URL <https://doi.org/10.48550/arXiv.2411.06559>.
- Tianqing Fang, Hongming Zhang, Zhisong Zhang, Kaixin Ma, Wenhao Yu, Haitao Mi, and Dong Yu. Webevolver: Enhancing web agent self-improvement with coevolving world model. *arXiv preprint arXiv:2504.21024*, 2025a.
- Vardaan Pahuja, Yadong Lu, Corby Rosset, Boyu Gou, Arindam Mitra, Spencer Whitehead, Yu Su, and Ahmed Awadallah. Explorer: Scaling exploration-driven web trajectory synthesis for multimodal web agents. *CoRR*, abs/2502.11357, 2025. doi: 10.48550/ARXIV.2502.11357. URL <https://doi.org/10.48550/arXiv.2502.11357>.
- Richard S Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM Sigart Bulletin*, 2(4):160–163, 1991.
- Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination, 2020. URL <https://arxiv.org/abs/1912.01603>.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. The llama 3 herd of models. *CoRR*, abs/2407.21783, 2024. doi: 10.48550/ARXIV.2407.21783. URL <https://doi.org/10.48550/arXiv.2407.21783>.

- Mengzhao Jia, Wenhao Yu, Kaixin Ma, Tianqing Fang, Zhihan Zhang, Siru Ouyang, Hongming Zhang, Meng Jiang, and Dong Yu. Leopard: A vision language model for text-rich multi-image tasks. *CoRR*, abs/2410.01744, 2024. doi: 10.48550/ARXIV.2410.01744. URL <https://doi.org/10.48550/arXiv.2410.01744>.
- OpenAI. Gpt-4 technical report. Technical Report, March 2023. URL <https://arxiv.org/abs/2303.08774>. A large multimodal model capable of processing image and text inputs and producing text outputs. Achieves human-level performance on various professional benchmarks including passing a simulated bar exam in the top 10
- Anthropic. Claude 3.7 sonnet: Hybrid reasoning model. <https://www.anthropic.com/news/claude-3-7-sonnet>, 2025. Accessed: 2025-04-18.
- Anthropic. Model context protocol. Open-source protocol, November 2024. URL <https://github.com/modelcontextprotocol>.
- Hongming Zhang, Xiaoman Pan, Hongwei Wang, Kaixin Ma, Wenhao Yu, and Dong Yu. Cognitive kernel: An open-source agent system towards generalist autopilots. *CoRR*, abs/2409.10277, 2024a. doi: 10.48550/ARXIV.2409.10277. URL <https://doi.org/10.48550/arXiv.2409.10277>.
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-world web interaction with grounded language agents. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL [http://papers.nips.cc/paper\\_files/paper/2022/hash/82ad13ec01f9fe44c01cb91814fd7b8c-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/82ad13ec01f9fe44c01cb91814fd7b8c-Abstract-Conference.html).
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. Mind2web: Towards a generalist agent for the web. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL [http://papers.nips.cc/paper\\_files/paper/2023/hash/5950bf290a1570ea401bf98882128160-Abstract-Datasets\\_and\\_Benchmarks.html](http://papers.nips.cc/paper_files/paper/2023/hash/5950bf290a1570ea401bf98882128160-Abstract-Datasets_and_Benchmarks.html).
- Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Russ Salakhutdinov, and Daniel Fried. Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pages 881–905. Association for Computational Linguistics, 2024a. doi: 10.18653/V1/2024.ACL-LONG.50. URL <https://doi.org/10.18653/v1/2024.acl-long.50>.
- Jialong Wu, Wenbiao Yin, Yong Jiang, Zhenglin Wang, Zekun Xi, Runnan Fang, Linhai Zhang, Yulan He, Deyu Zhou, Pengjun Xie, and Fei Huang. Webwalker: Benchmarking llms in web traversal. *CoRR*, abs/2501.07572, 2025a. doi: 10.48550/ARXIV.2501.07572. URL <https://doi.org/10.48550/arXiv.2501.07572>.
- Ziniu Zhang, Shulin Tian, Liangyu Chen, and Ziwei Liu. Mmina: Benchmarking multihop multimodal internet agents. *CoRR*, abs/2404.09992, 2024b. doi: 10.48550/ARXIV.2404.09992. URL <https://doi.org/10.48550/arXiv.2404.09992>.
- Shikhar Murty, Hao Zhu, Dzmitry Bahdanau, and Christopher D. Manning. Nnetnav: Unsupervised learning of browser agents through environment interaction in the wild. *CoRR*, 2025. URL <https://arxiv.org/abs/2410.02907>.
- Brandon Trabucco, Gunnar A. Sigurdsson, Robinson Piraamuthu, and Ruslan Salakhutdinov. Towards internet-scale training for agents. *CoRR*, abs/2502.06776, 2025. doi: 10.48550/ARXIV.2502.06776. URL <https://doi.org/10.48550/arXiv.2502.06776>.

- Xiaoxi Li, Jiajie Jin, Guanting Dong, Hongjin Qian, Yutao Zhu, Yongkang Wu, Ji-Rong Wen, and Zhicheng Dou. Webthinker: Empowering large reasoning models with deep research capability. *arXiv preprint arXiv:2504.21776*, 2025a.
- Jialong Wu, Baixuan Li, Runnan Fang, Wenbiao Yin, Liwen Zhang, Zhengwei Tao, Dingchu Zhang, Zekun Xi, Gang Fu, Yong Jiang, et al. Webdancer: Towards autonomous information seeking agency. *arXiv preprint arXiv:2505.22648*, 2025b.
- Kuan Li, Zhongwang Zhang, Huifeng Yin, Liwen Zhang, Litu Ou, Jialong Wu, Wenbiao Yin, Baixuan Li, Zhengwei Tao, Xinyu Wang, et al. Websailor: Navigating super-human reasoning for web agent. *arXiv preprint arXiv:2507.02592*, 2025b.
- Zhengwei Tao, Jialong Wu, Wenbiao Yin, Junkai Zhang, Baixuan Li, Haiyang Shen, Kuan Li, Liwen Zhang, Xinyu Wang, Yong Jiang, et al. Webshaper: Agentically data synthesizing via information-seeking formalization. *arXiv preprint arXiv:2507.15061*, 2025.
- Tianqing Fang, Zhisong Zhang, Xiaoyang Wang, Rui Wang, Can Qin, Yuxuan Wan, Jun-Yu Ma, Ce Zhang, Jiaqi Chen, Xiyun Li, et al. Cognitive kernel-pro: A framework for deep research agents and agent foundation models training. *arXiv preprint arXiv:2508.00414*, 2025b.
- MiroMindAI. Miroflow: A consistent agent framework with reproducible performance. <https://github.com/MiroMindAI/MiroFlow>, 2025. Accessed: [Insert date of access here].
- Xinyu Geng, Peng Xia, Zhen Zhang, Xinyu Wang, Qiuchen Wang, Ruixue Ding, Chenxi Wang, Jialong Wu, Yida Zhao, Kuan Li, Yong Jiang, Pengjun Xie, Fei Huang, and Jingren Zhou. Webwatcher: Breaking new frontier of vision-language deep research agent, 2025. URL <https://arxiv.org/abs/2508.05748>.
- Jing Yu Koh, Stephen McAleer, Daniel Fried, and Ruslan Salakhutdinov. Tree search for language model agents. *CoRR*, abs/2407.01476, 2024b. doi: 10.48550/ARXIV.2407.01476. URL <https://doi.org/10.48550/arXiv.2407.01476>.
- Pranav Putta, Edmund Mills, Naman Garg, Sumeet Motwani, Chelsea Finn, Divyansh Garg, and Rafael Rafailov. Agent Q: advanced reasoning and learning for autonomous AI agents. *CoRR*, abs/2408.07199, 2024. doi: 10.48550/ARXIV.2408.07199. URL <https://doi.org/10.48550/arXiv.2408.07199>.
- Xiao Yu, Baolin Peng, Vineeth Vajipey, Hao Cheng, Michel Galley, Jianfeng Gao, and Zhou Yu. Exact: Teaching AI agents to explore with reflective-mcts and exploratory learning. *CoRR*, abs/2410.02052, 2024. doi: 10.48550/ARXIV.2410.02052. URL <https://doi.org/10.48550/arXiv.2410.02052>.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language agent tree search unifies reasoning, acting, and planning in language models. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024b. URL <https://openreview.net/forum?id=njwv9BsGHF>.
- Yao Zhang, Zijian Ma, Yunpu Ma, Zhen Han, Yu Wu, and Volker Tresp. Webpilot: A versatile and autonomous multi-agent system for web task execution with strategic exploration. *CoRR*, abs/2408.15978, 2024c. doi: 10.48550/ARXIV.2408.15978. URL <https://doi.org/10.48550/arXiv.2408.15978>.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL [http://papers.nips.cc/paper\\_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html).
- David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2455–2467, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/2de5d16682c3c35007e4e92982f1a2ba-Abstract.html>.

- Dani Valevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time game engines. *CoRR*, abs/2408.14837, 2024. doi: 10.48550/ARXIV.2408.14837. URL <https://doi.org/10.48550/arXiv.2408.14837>.
- Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos J. Storkey, Tim Pearce, and François Fleuret. Diffusion for world modeling: Visual details matter in atari. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL [http://papers.nips.cc/paper\\_files/paper/2024/hash/6bdde0373d53d4a501249547084bed43-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2024/hash/6bdde0373d53d4a501249547084bed43-Abstract-Conference.html).
- Max Olan Smith and Michael P. Wellman. Co-learning empirical games and world models. *CoRR*, abs/2305.14223, 2023. doi: 10.48550/ARXIV.2305.14223. URL <https://doi.org/10.48550/arXiv.2305.14223>.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. Reasoning with language model is planning with world model. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 8154–8173. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.EMNLP-MAIN.507. URL <https://doi.org/10.18653/v1/2023.emnlp-main.507>.
- Shuofei Qiao, Runnan Fang, Ningyu Zhang, Yuqi Zhu, Xiang Chen, Shumin Deng, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. Agent planning with world knowledge model. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL [http://papers.nips.cc/paper\\_files/paper/2024/hash/d032263772946dd5026e7f3cd22bce5b-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2024/hash/d032263772946dd5026e7f3cd22bce5b-Abstract-Conference.html).
- Hyunjoo Chae, Namyoun Kim, Kai Tzu-iunn Ong, Minju Gwak, Gwanwoo Song, Jihoon Kim, Sunghwan Kim, Dongha Lee, and Jinyoung Yeo. Web agents with world models: Learning and leveraging environment dynamics in web navigation. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL <https://openreview.net/forum?id=moWiYJuSGF>.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiase Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL <https://arxiv.org/abs/2503.14476>.
- Hanyu Lai, Xiao Liu, Iat Long Iong, Shuntian Yao, Yuxuan Chen, Pengbo Shen, Hao Yu, Hanchen Zhang, Xiaohan Zhang, Yuxiao Dong, and Jie Tang. Autowebglm: A large language model-based web navigating agent. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24*, page 5295–5306, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704901. doi: 10.1145/3637528.3671620. URL <https://doi.org/10.1145/3637528.3671620>.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. Group sequence policy optimization, 2025. URL <https://arxiv.org/abs/2507.18071>.
- Apurva Gandhi and Graham Neubig. Go-browse: Training web agents with structured exploration, 2025. URL <https://arxiv.org/abs/2506.03533>.

## A More Details About DynaWeb

### A.1 WebArena Training Prompt

We provide the full system prompt used to train our web agent on WebArena. The prompt defines the agent’s role, available actions, observation format, and task completion criteria, and is used consistently across both real-environment interaction and imagined rollouts generated by the web world model. This prompt is fixed throughout training and evaluation to ensure reproducibility.

**WebArena Agent System Prompt**

You are an AI assistant performing tasks on a web browser. You will be provided with task objective, current step, web page observations, interaction history and previous taked notes. You need to issue an action for this step.

Generate the response in the following format:  
**INTERACTION HISTORY SUMMARY:**  
 Emphasize all important details in the INTERACTION HISTORY section.

**OBSERVATION DESCRIPTION:**  
 Describe information in the CURRENT OBSERVATION section. Emphasize elements and features that are relevant or potentially helpful for fulfilling the objective in detail.

**REASON:**  
 Provide your rationale for proposing the subsequent action commands here.

**ACTION:**  
 Select your action here.

**OBSERVATION HIGHLIGHT:**  
 List the numerical ids of elements on the current webpage based on which you would issue your action. Also include elements on the current webpage you would attend to if you fail in the future and have to restore to this step. Don’t include elements from the previous pages. Select elements at a higher hierarchical level if most their children nodes are considered crucial. Sort by relevance and potential values from high to low, and separate the ids with commas. E.g., 1321, 52, 756, 838.

You are ONLY allowed to use the following action commands. Strictly adheres to the given format. Only issue one single action.  
 Use the following actions:

- `click [id]`: To click on an element with its numerical ID on the webpage. E.g., `click [7]` If clicking on a specific element doesn’t trigger the transition to your desired web state, this is due to the element’s lack of interactivity or GUI visibility. In such cases, move on to interact with OTHER similar or relevant elements INSTEAD.
- `type [id] [content] [press_enter_after=0|1]`: To type content into a field with a specific ID. By default, the “Enter” key is pressed after typing unless `press_enter_after` is set to 0. E.g., `type [15] [Carnegie Mellon University] [1]` If you can’t find what you’re looking for on your first attempt, consider refining your search keywords by breaking them down or trying related terms.
- `stop [answer]`: To stop interaction and return response. Present your answer within the brackets. If and only if the task doesn’t require a textual answer or appears insurmountable, indicate “N/A” and additional reasons and all relevant information you gather as the answer. Otherwise, including “N/A” will be penalized. E.g., `stop [5h 47min]`
- `go_back`: To return to the previously viewed page.

Figure 5. System prompt used for training and evaluation of the WebArena agent.

### A.2 World Model System Prompt

We provide the full system prompt used by the web world model. The prompt specifies the input information available to the model, including the user objective, current webpage state, and executed actions, as well as the required output format for predicting web state changes and the resulting next-step accessibility

tree. This prompt is fixed throughout training and inference, and is shared across all experiments to ensure reproducibility.

### Web World Model System Prompt

You are an intelligent agent that predicts next state from given current action in a web environment, with your own logical reasoning.

Here's the information you'll have:

The user's objective: This is the task you're trying to complete.

The current web page's accessibility tree: This is a simplified representation of the webpage, providing key information.

The current web page's URL: This is the page you're currently navigating.

The previous action: This is the action you just performed in the previous step. It may be helpful to track your progress.

The current action: This is the current action that you performed to achieve the user's objective in the current web page's accessibility tree.

The format of previous actions can fall into several categories:

Page Operation Actions:

`click [id]`: This action clicks on an element with a specific id on the webpage.

`type [id] [content]`: Use this to type the content into the field with id. By default, the 'Enter' key is pressed after typing unless `press_enter_after` is set to 0, i.e., `type [id] [content] [0]`.

`hover [id]`: Hover over an element with id.

`press [key_comb]`: Simulates the pressing of a key combination on the keyboard (e.g., Ctrl+v).

`scroll [down]` or `scroll [up]`: Scroll the page up or down.

Tab Management Actions:

`new_tab`: Open a new, empty browser tab.

`tab_focus [tab_index]`: Switch the browser's focus to a specific tab using its index.

`close_tab`: Close the currently active tab.

URL Navigation Actions:

`goto [url]`: Navigate to a specific URL.

`go_back`: Navigate to the previously viewed page.

`go_forward`: Navigate to the next page (if a previous 'go\_back' action was performed)

Completion Action:

`stop [answer]`: Done when you believe the task is complete.

Given the information above, you should first perform reasoning to predict expected changes on the current web page's accessibility tree, and then generate the resulting next web page's accessibility tree based on your predicted web page changes.

Generate your answer in the following format:

[Web state changes]

changes

[Next page accessibility tree]

next\_acc\_tree

where `changes` are the predicted web page changes, and `next_acc_tree` is your predicted next page accessibility tree.

For `next_acc_tree`, you MUST generate a valid accessibility tree based on your predicted changes, do NOT output summary descriptions of how the `next_acc_tree` will change.

If the full predicted `next_acc_tree` is too long, you should output a pruned tree by removing unimportant elements that the web agent will unlikely use in subsequent action steps.

Meanwhile, you should NOT omit key unchanged elements that the web agent might interact with in future time steps.

Note: Even if the web page does not change, you must still output the complete `next_acc_tree` !

**Figure 6.** System prompt used to train the web world model for predicting next-step accessibility trees from actions and current observations.

### A.3 Training Recipe and Hyperparameters

This subsection summarizes the training recipe used for DynaWeb agent optimization. Our training follows an online reinforcement learning setup in a synthetic web environment powered by a learned world model. The agent is initialized from a supervised web navigation model and optimized using Group Sequence Policy Optimization (GSPO) with sparse terminal rewards. Imagined rollouts generated by the world model are interleaved with real expert trajectories to stabilize learning.

Category	Setting
Backbone policy model	LLaMA-3.1-8B-Instruct
Optimization algorithm	GSPO (sequence-level policy optimization)
Advantage estimator	GRPO-style estimator
Policy learning rate	$1 \times 10^{-6}$
Entropy coefficient	0
KL regularization	Disabled
Training epochs	10
Train batch size	4
Validation batch size	32
Max prompt length	32,000 tokens
Max response length	16,000 tokens
Max tokens per GPU (PPO)	48,000
Rollout engine	vLLM (asynchronous)
Rollout samples per prompt ( $n$ )	8
Rollout temperature	0.7
Rollout top- $p$	0.9
Validation top- $p$	0.8
Validation top- $k$	20
World model max steps	10
Agent max steps	10
World model max tokens	8,192
World model temperature	0.7
World model top- $p$	0.9
Precision	bfloat16
Gradient checkpointing	Enabled
FSDP parameter offloading	Enabled
Optimizer offloading	Enabled
GPUs	$8 \times$ H100 (single node)

**Table 3.** Key hyperparameters for DynaWeb agent training.