

ColorBrowserAgent: Complex Long-Horizon Browser Agent with Adaptive Knowledge Evolution

Jihong Wang^{1,*}, Jiamu Zhou^{1,*}, Weiming Zhang^{1,2}, Teng Wang¹, Weiwen Liu²,
Zhuosheng Zhang², Xingyu Lou^{1,†}, Weinan Zhang², Huarong Deng¹, Jun Wang¹

¹OPPO Research Institute

²Shanghai Jiao Tong University

wangjihong@oppo.com, zhoujiamu@oppo.com

Abstract

With the advancement of vision-language models, web automation has made significant progress. However, deploying autonomous agents in real-world settings remains challenging, primarily due to site heterogeneity, where generalist models lack domain-specific priors for diverse interfaces, and long-horizon instability, characterized by the accumulation of decision drift over extended interactions. To address these challenges, we introduce ColorBrowserAgent (Complex Long-Horizon Browser Agent), a knowledge-evolving agent for robust web automation. Our approach addresses these challenges through two synergistic mechanisms: human-in-the-loop knowledge adaptation that transforms sparse human feedback into reusable domain knowledge, and knowledge-aligned progressive summarization that stabilizes long interactions through memory compression. Extensive experiments on WebArena, WebChoreArena and industrial deployment show that ColorBrowserAgent consistently outperforms strong baselines. It achieves a state-of-the-art success rate of 71.2% on WebArena and maintains 47.4% performance under zero-shot transfer setting on WebChoreArena. In commercial deployment, it improves user satisfaction by 19.3% relatively, verifying its robustness in real-world scenarios.

1 Introduction

Web browsers serve as the universal central interface for digital workflows across daily user activities and complex enterprise business operations. Recent advances in function-calling capabilities of large vision-language models (VLMs) (Yang et al., 2025a; Wang et al., 2025a; Achiam et al., 2023; Ma et al., 2025) have empowered web agents to interpret natural language instructions and directly operate on graphical user interfaces (GUIs). This capability enables the automation of complex, multi-step browser workflows that traditionally require manual intervention (Zhang et al., 2024).

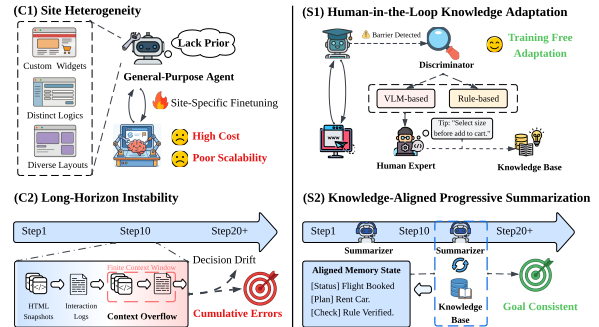


Figure 1: Challenges and Solutions in Complex Web Automation.

Despite recent progress, reliable deployment of autonomous web agents in production remains difficult (Xue et al., 2025; Kara et al., 2025). Web interfaces are inherently open and visually driven, introducing substantial variance compared to structured APIs. In practice, two core challenges impede scalability, as depicted in Figure 1. (1) **Site Heterogeneity**. General-purpose models lack domain-specific priors to handle diverse layout conventions and idiosyncratic business logic across websites. Existing approaches typically rely on post-training using interaction traces (Qi et al., 2024; Wei et al., 2025). This strategy is not only prohibitively expensive but also critically lacks generalization to out-of-distribution websites and robustness to interface evolution. (2) **Long-Horizon Instability**. As interaction trajectories extend, agents frequently suffer from decision drift, gradually deviating from the original task objectives (Liu et al., 2025a). This challenge is exacerbated by site heterogeneity: long-horizon interactions are more prone to lacking domain-specific priors, and minor interface ambiguities (e.g., transient pop-ups) that are negligible in short trajectories can accumulate, leading to severe task deviation.

To address site heterogeneity and long-horizon instability, we introduce **ColorBrowserAgent**

(Complex Long-Horizon Browser Agent), a web agent based on tuning-free adaptation that externalizes domain knowledge and stabilizes long-horizon interaction trajectories. As illustrated in Figure 1, our framework resolves the two challenges through two synergistic mechanisms: (1) **Human-in-the-loop Knowledge Adaptation**, which tackles site heterogeneity by capturing site-specific logic from sparse expert feedback upon failure. These insights are crystallized into an external knowledge base and retrieved contextually, enabling the agent to adapt to diverse interfaces without retraining; and (2) **Knowledge-aligned Progressive Summarization**, which mitigates long-horizon decision drift by maintaining a compressed belief state. This mechanism progressively summarizes interaction history and enforces consistency with retrieved expert priors, ensuring the agent remains focused on the objective throughout long-horizon trajectories.

In summary, our contributions are as follows:

- **Tuning-Free Adaptation.** We propose a human-in-the-loop knowledge adaptation framework that externalizes domain-specific priors into a persistent memory, decoupling adaptability from model parameters and eliminating the need for retraining.
- **Long-Horizon Stability.** We introduce a knowledge-aligned progressive summarization mechanism that mitigates decision drift by compressing a belief state across long-horizon interaction trajectories.
- **Empirical and Deployment Validation.** ColorBrowserAgent achieves 71.2% success rate on WebArena and strong zero-shot transfer to WebChoreArena. In a commercial deployment, it further improves user satisfaction by 19.3%, demonstrating practical effectiveness in dynamic environments.

2 Related Works

2.1 LLM-based Web Agents

Several recent works have introduced realistic browser environments and benchmarks to evaluate LLM-based web agents (Zhou et al., 2023; Ning et al., 2025; Li et al., 2025; Miyai et al., 2025).

Existing approaches fall into two paradigms: **training-based methods**, which update model parameters with interaction data, and **test-time search methods**, which improve performance via

inference-time exploration. Training-based methods like WebRL (Qi et al., 2024) and WebAgent-r1 (Wei et al., 2025) learn navigation policies from interaction trajectories via reinforcement learning. However, in real-world industrial settings, obtaining high-quality interaction trajectories is costly and labor-intensive, making frequent retraining infeasible. In contrast, search-based methods (best-first search (Koh et al., 2024; Wang et al., 2025c), Monte Carlo Tree Search (Zhang et al., 2025d; Wang et al., 2025b)) improve performance by exploring candidate action trajectories at inference time. While they boost task success rates, their extensive inference-time exploration incurs substantial latency and cost, limiting real-time deployment suitability. To address these limitations, ColorBrowserAgent bypasses expensive parameter updates and high-latency search via a collaborative, knowledge-driven mechanism enabling robust, efficient execution.

2.2 Human-in-the-Loop Strategy

Human-in-the-Loop (HITL) approaches are widely adopted to improve autonomous agent reliability in complex environments (Zanzotto, 2019; Wu et al., 2022; Wang et al., 2026). Existing research focuses on leveraging human data for policy learning. For example, demonstration-based methods like AppAgent (Zhang et al., 2025a; Li et al., 2024) collect expert trajectories to initialize agent skills before deployment. Similarly, feedback-driven approaches such as WebCoach (Liu et al., 2025b) use human corrections to iteratively refine the agent’s policy. These methods require significant data accumulation and training overhead to show improvements. In contrast, ColorBrowserAgent treats expert interventions as explicit knowledge injection: instead of model tuning, we encode expert guidance into a retrieval-based knowledge base, enabling immediate failure resolution and cross-task logic reuse without retraining.

3 Methodology

3.1 Preliminary

We view web automation as a goal-conditioned, sequential decision-making process. Given a user instruction g , the agent interacts with a browser environment over discrete time steps. At each time step t , the agent receives a multimodal observation o_t consisting of the DOM tree and a rendered screenshot of the current webpage. Based on the

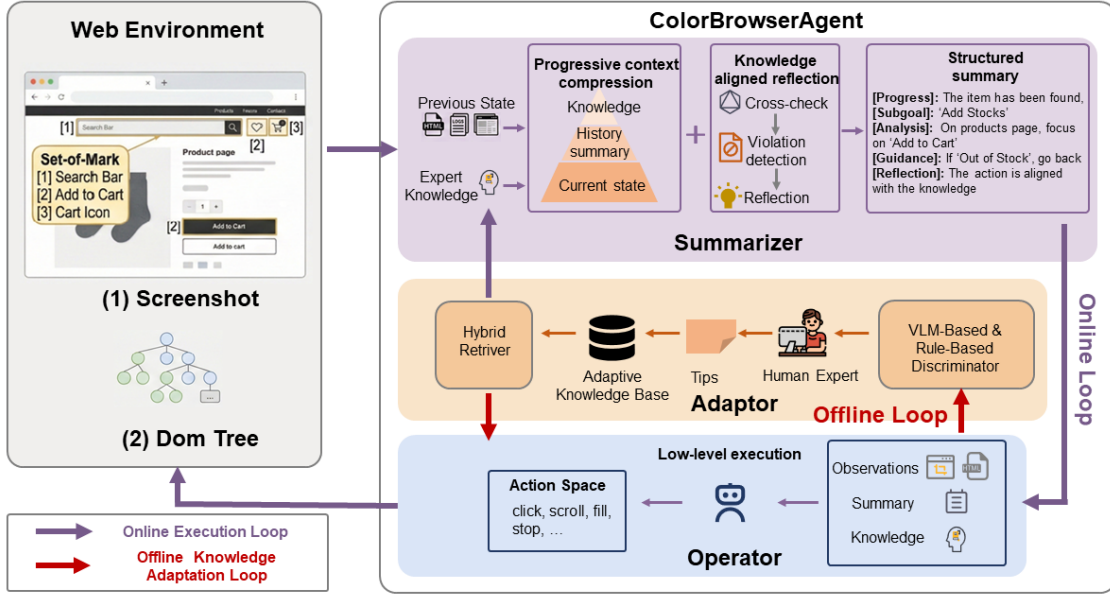


Figure 2: The framework of ColorBrowserAgent. It operates via two synergistic loops: (1) The Offline Knowledge Adaptation Loop, where the Adaptor crystallizes expert feedback into an Adaptive Knowledge Base; and (2) The Online Execution Loop, where the Summarizer maintains a compressed belief state and the Operator executes actions based on aligned intent.

interaction history $h_t = (o_0, a_0, \dots, o_t)$ and the instruction g , the agent selects an action a_t from a predefined action space (e.g., click, type, scroll, stop). A task is considered successful if the terminal environment state satisfies the goal constraints induced by the instruction g .

3.2 Overview of ColorBrowserAgent

Existing approaches (Wei et al., 2025; Qi et al., 2024) adapt to specific web environments by internalizing site-specific behaviors into model parameters through additional training. However, such parametric adaptation often requires retraining when web interfaces evolve and can struggle to transfer to previously unseen websites. In contrast, we propose **ColorBrowserAgent**, a tuning-free adaptation framework that externalizes domain-specific expert priors into an explicit knowledge base and retrieves them at runtime to guide action selection.

As shown in Figure 2, ColorBrowserAgent addresses the challenges of site heterogeneity and long-horizon instability through three core components: the Adaptor, the Summarizer, and the Operator. These components function within two synergistic loops:

- Offline Knowledge Adaptation Loop: To

tackle site heterogeneity, the Adaptor constructs the Adaptive Knowledge Base (AKB) through a human-in-the-loop process. During an exploratory phase, execution failures are monitored and analyzed by human experts, who distill site-specific logic into reusable expert priors and inject them into the AKB.

- Online Execution Loop: During inference, the system follows a sequential pipeline: Web Observation \rightarrow AKB Retrieval \rightarrow Summarization \rightarrow Execution. The Summarizer mitigates long-horizon instability by performing knowledge-aligned progressive summarization, compressing interaction history while ensuring consistency with constraints retrieved by the Adaptor. Finally, the Operator generates and executes the concrete browser action.

The detailed algorithmic procedure of our framework is provided in Appendix G.

3.3 Adaptor

The Adaptor fulfills two primary responsibilities: the Offline Knowledge Adaptation Loop for knowledge extraction and injection, and the Online Execution Loop for knowledge retrieval.

Offline Knowledge Adaptation Loop. To ensure reliability, the system employs an on-demand intervention strategy via a human-in-the-loop mechanism. It continuously monitors execution via a hybrid detector $f_{\text{trigger}}(o_t, h_t)$, which combines a rule-based detector for deterministic failures and a VLM-based evaluator for semantic inconsistencies between the UI state and the user’s intent. Intervention is triggered whenever either detector signals a failure. Upon failure detection, a human expert analyzes the failure context and formulates a corrective tip. Crucially, to ensure generalizability, these tips describe the site’s operational logic rather than the specific execution flow. This feedback is persistently crystallized into the AKB.

Online Execution Loop. During the processing of a new query, the system dynamically retrieves relevant knowledge from the knowledge base via a cascade of mechanisms: (1) URL Pattern Matching for site-specific lookups; (2) Keyword Search for content-aware constraints; and (3) Visual-Semantic Embedding for fuzzy UI matching.

3.4 Summarizer

To mitigate decision drift in long-horizon tasks, where the linearly growing history $O(T)$ exhausts the context window and weakens instruction adherence, we introduce a Summarizer agent. This supervisory component maintains a structured belief state to preserve temporal coherence and enforce alignment with expert priors. The update process iteratively refines this belief state through two core mechanisms:

Progressive Context Compression. To handle extended trajectories, the Summarizer employs a hierarchical retention strategy. It retains fine-grained interaction details only for the current active subgoal, while recursively collapsing completed history into high-level semantic summaries. This mechanism bounds the memory footprint to near-constant size ($O(1)$) regardless of task length, effectively preventing context overflow and the resulting hallucination or loss of goal focus.

Knowledge Aligned Reflection. Beyond compression, the Summarizer also serves as a knowledge-aware consistency monitor. It compares the agent’s planned actions with the retrieved expert knowledge. When a mismatch is detected, the Summarizer injects corrective guidance into the observation stream. This mechanism ensures

that the domain priors incorporated by the Adaptor are consistently applied during execution, reducing deviations from expert-informed strategies.

3.5 Operator

The Operator translates aligned intent into executable browser actions. At each time step, it conditions on the current multimodal observation (DOM tree and screenshot), the retrieved knowledge constraints from the AKB, and the structured summary produced by the Summarizer. Instead of relying on raw interaction history, the Operator utilizes the compressed belief state to focus on the current subgoal while respecting site-level priors. It first infers the immediate subgoal implied by the current state, then grounds this intent to specific UI elements by jointly reasoning over DOM attributes and visual cues, similar to recent multimodal GUI agents for web automation.

4 Experiments

4.1 Experimental Setup

We evaluate ColorBrowserAgent on two complementary benchmarks to assess both in-domain performance and zero-shot knowledge generalization:

- **WebArena** (Zhou et al., 2023): Our primary benchmark, consisting of 812 long-horizon tasks across multiple real-world domains (e.g., Shopping, GitLab, Reddit, Map).
- **WebChoreArena** (Miyai et al., 2025): An extended benchmark featuring more complex and memory-intensive tasks. We evaluate in a zero-shot transfer setting, using only knowledge derived from WebArena.

We use GPT-5 (OpenAI, 2025) as the underlying reasoning model. Our system is implemented on top of BrowserGym (Chezelles et al., 2025) with Playwright for programmatic browser control. Observations consist of multimodal inputs, combining Set-of-Marks (Yang et al., 2023) and accessibility tree representations (Zhou et al., 2023). We adopt an action space adapted from AgentOccam (Yang et al., 2024) to reduce execution errors. The maximum interaction length is set to 30 steps.

We construct the AKB through a one-off exploration phase, where human operators summarize failure cases into reusable site-level tips rather than query-specific traces. This process requires less than 1 person-day and yields a compact set of

Table 1: Comparison of success rates on the WebArena (Zhou et al., 2023) benchmark. Best results are highlighted in **bold**, and second-best results are underlined. Relative improvements are computed with respect to the strongest baseline. “-” means that the result is not reported in the original paper.

| Method | Overall | Reddit | GitLab | Shopping | Admin | Map | Multisite |
|--------------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| BrowserGym (Drouin et al., 2024) | 15.0 | 20.2 | 19.0 | 17.2 | 14.8 | 25.5 | - |
| LM-TS (Koh et al., 2024) | 19.2 | 11.3 | 13.9 | 27.8 | 16.5 | 26.6 | 16.7 |
| Go-Browse (Gandhi et al., 2025) | 22.6 | 30.7 | 15.3 | 22.4 | 25.3 | 17.9 | - |
| AWM (Wang et al., 2024a) | 35.5 | 50.9 | 31.8 | 30.8 | 29.1 | 43.3 | - |
| Branch-n-Browse (He et al., 2025) | 35.8 | 50.9 | 36.7 | 34.6 | 26.4 | 46.8 | 18.8 |
| WebPilot (Zhang et al., 2025c) | 37.2 | 65.1 | 39.4 | 36.9 | 24.7 | 33.9 | - |
| AgentOccam (Yang et al., 2025c) | 45.7 | 67.0 | 43.3 | 46.2 | 38.9 | 52.3 | 16.7 |
| AgentSymbiotic (Zhang et al., 2025b) | 52.1 | 66.0 | 51.0 | 48.0 | 49.0 | <u>60.0</u> | 29.0 |
| ScribeAgent (Shen et al., 2024) | 53.0 | 73.7 | 59.7 | 45.8 | 37.9 | 56.3 | - |
| WebOperator (Dihan et al., 2025) | 54.6 | <u>76.4</u> | 52.8 | 49.2 | 55.0 | 55.2 | 31.3 |
| CUGA (Shlomov et al., 2025) | <u>61.7</u> | 75.5 | <u>61.7</u> | <u>58.3</u> | <u>62.6</u> | 64.2 | <u>35.4</u> |
| ColorBrowserAgent | 71.2 | 87.4 | 65.7 | 72.9 | 76.4 | 55.9 | 64.8 |
| Relative Improvement | 15.38% | 14.38% | 6.48% | 25.04% | 22.04% | -14.19% | 83.05% |

Table 2: Zero-shot knowledge generalization on WebChoreArena (Miyai et al., 2025). Best results are highlighted in **bold**, and second-best results are underlined. Relative improvements are computed with respect to the strongest baseline. “-” indicates that the result is not reported in the original paper.

| Method | Overall | Shopping | Reddit | Admin | GitLab |
|--|-------------|-------------|-------------|-------------|-------------|
| SteP (Sodhi et al., 2023) | 3.1 | 2.6 | 4.4 | 0.7 | 4.7 |
| BrowserGym (Chezelles et al., 2024) | 21.2 | 15.4 | 15.4 | 26.5 | 27.6 |
| AWM (Wang et al., 2024b) | 22.4 | 18.0 | 14.3 | 30.3 | 26.8 |
| AgentOccam (Yang et al., 2024) | 21.5 | 21.3 | 11.0 | 30.8 | 22.8 |
| WEBDART (Yang et al., 2025b) | <u>31.1</u> | <u>35.0</u> | <u>26.4</u> | <u>33.8</u> | <u>29.1</u> |
| ColorBrowserAgent w/o knowledge | 34.4 | 38.5 | 27.5 | 52.2 | 33.9 |
| ColorBrowserAgent | 47.4 | 43.6 | 44.0 | 58.7 | 53.5 |
| Relative Improvement | 52.41% | 24.57% | 66.67% | 73.67% | 83.85% |

reusable site-level tips totaling 52 distilled rules across domains (GitLab: 13, Map: 7, Reddit: 5, Shopping: 9, Admin: 18). After construction, the knowledge base remains frozen throughout evaluation, with no test-time human intervention.

4.2 Performance on WebArena Benchmark

As detailed in Table 1, ColorBrowserAgent achieves a 71.2% overall task success rate on the WebArena benchmark, establishing a new state-of-the-art and substantially outperforming the strongest prior method, CUGA (61.7%), by a relative margin of 15.38%.

A granular analysis across individual domains reveals that our framework delivers consistent, double-digit relative improvements in the majority of categories. Specifically, we observe exceptional gains in domains characterized by diverse interface structures and complex, long-horizon workflows:

Multisite (+83.05%), Shopping (+25.04%), and Admin (+22.04%). Additionally, performance in the Reddit domain reaches an impressive 87.4% success rate. These substantial margins suggest that our synergistic approach—combining externalized domain knowledge via the Adaptor and memory stabilization via the Summarizer—effectively mitigates the core challenges of site heterogeneity and long-horizon instability. While performance in the Map domain is slightly lower (55.9%) compared to the best baseline (64.2%), ColorBrowserAgent still remains highly competitive across the board.

4.3 Knowledge Generalization on WebChoreArena

To evaluate whether crystallized site-level priors generalize beyond their original task distribution, we directly transfer the frozen knowledge base from WebArena to WebChoreArena without any

task-specific adaptation. Importantly, no query-level information from WebChoreArena is used during knowledge construction, ensuring a strict zero-shot transfer setting.

Despite the substantially increased task complexity in WebChoreArena—characterized by longer reasoning chains and heavier memory demands—ColorBrowserAgent achieves a 47.4% overall success rate, substantially outperforming both the strongest prior method (WEBDART, 31.1%) and our knowledge-free variant (34.4%). Crucially, the reported results for all baselines on WebChoreArena are directly cited from the WebDART paper (Yang et al., 2025b) under a GPT-5-based setting, making this comparison backbone-fair. The substantial gain over the strongest baseline demonstrates that the performance improvements cannot be attributed solely to the use of a more advanced LLM. Furthermore, the performance gap between the knowledge-enabled and knowledge-free variants indicates that reusable site-level priors contribute substantially beyond backbone reasoning alone.

The gains are observed across most domains, with particularly large margins in GitLab and Admin. This pattern suggests that crystallized knowledge effectively stabilizes long-horizon execution while remaining invariant to query complexity. Together, these results indicate that domain priors learned from simpler tasks can generalize to substantially more complex scenarios, decoupling operational logic from task difficulty.

4.4 Ablation Study

To assess the contribution of each component, we conduct ablation experiments on WebArena-Lite (165 tasks). Importantly, these experiments use a unified GPT-5 setup, guaranteeing that differences in performance are strictly attributable to the proposed architecture rather than the backbone. As shown in Figure 3, the full ColorBrowserAgent achieves a 72.6% success rate, improving over the vanilla baseline by 10.9%.

Removing Summarizer reduces performance to 68.8%, indicating that structured memory alignment helps stabilize long trajectories. Removing Adaptor results in a larger drop to 65.4%, suggesting that externalized site-level priors play a central role in handling heterogeneity.

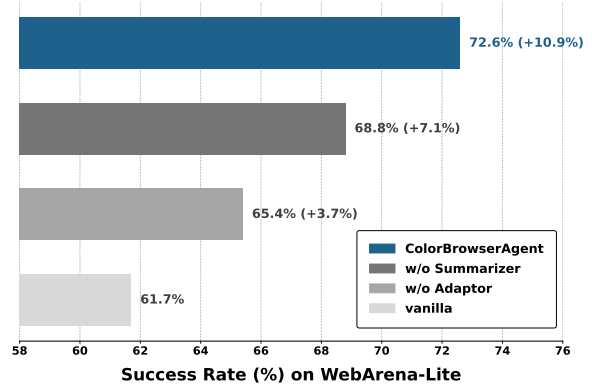


Figure 3: Ablation studies on WebArena-Lite.

4.5 Industrial Practice: Real-Time Service Verification

We deploy ColorBrowserAgent in a commercial intelligent travel planning system, where it operates as the Web Execution Unit within a multi-agent architecture. The agent autonomously navigates heterogeneous third-party service portals to retrieve real-time information (e.g., flight prices) in response to user requests. An online A/B test compares our system against a traditional search-engine-based pipeline. Note that specific token consumption and cost details are omitted for commercial confidentiality.

Engineering Adaptations. To ensure reliability under real-world conditions, we introduce several system-level optimizations. First, navigation prioritizes direct URL parameter construction when possible, reducing dependence on fragile UI interactions. Second, the action space is extended with lightweight primitives such as `take_note()` for persistent memory and `calculate()` for deterministic arithmetic, which reduces reasoning errors in multi-step workflows. Third, a session watchdog monitors execution and automatically restores state upon transient failures.

Deployment Results. In this production environment, the system achieves a 19.3% relative improvement in the User Satisfaction Rate compared to the baseline. Notably, in vertical domains with accumulated site-level priors, the agent exceeds 95% task success, demonstrating that crystallized knowledge remains effective under dynamic interface updates.

5 Conclusion

We presented ColorBrowserAgent, a knowledge-evolving web agent designed for reliable industrial deployment. By externalizing site-level expertise through Adaptor and stabilizing long interaction trajectories with Summarizer, Our framework addresses two practical bottlenecks in production web automation: interface heterogeneity and long-horizon instability. Extensive evaluations on WebArena and WebChoreArena demonstrate strong performance and zero-shot knowledge transfer without task-specific retraining. More importantly, deployment in a commercial real-time service system shows measurable improvements in user satisfaction in dynamic, frequently changing web environments. These results suggest that decoupling domain knowledge from model parameters through expert-guided persistent memory provides a scalable, low-maintenance paradigm for industrial web agents.

Limitations

Cold-Start Adaptation. When deployed to entirely new domains without prior knowledge coverage, the system requires increased human intervention to identify and crystallize site-specific logic. During this initial phase, performance may fluctuate until sufficient domain priors are accumulated.

Model Dependency. Our current implementation relies on state-of-the-art proprietary VLMs (e.g., GPT-5) to guarantee high-precision visual reasoning. Distilling these capabilities into smaller, cost-effective open-weights models for edge deployment is a critical next step.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. [arXiv preprint arXiv:2303.08774](#).
- De Chezelles, Thibault Le Sellier, Sahar Omid Shayegan, Lawrence Keunho Jang, Xing Han Lu, Ori Yoran, Dehan Kong, Frank F Xu, Siva Reddy, Quentin Cappart, and 1 others. 2024. The browsergym ecosystem for web agent research. [arXiv preprint arXiv:2412.05467](#).
- Thibault Le Sellier De Chezelles, Maxime Gasse, Alexandre Drouin, Massimo Caccia, Léo Boisvert, Megh Thakkar, Tom Marty, Rim Assouel, Sahar Omid Shayegan, Lawrence Keunho Jang, Xing Han Lù, Ori Yoran, Dehan Kong, Frank F. Xu, Siva Reddy, Quentin Cappart, Graham Neubig, Ruslan Salakhutdinov, Nicolas Chapados, and Alexandre Lacoste. 2025. [The browsergym ecosystem for web agent research](#). [arXiv preprint arXiv:2412.05467](#), arXiv:2412.05467.
- Mahir Labib Dihan, Tanzima Hashem, Mohammed Eunus Ali, and Md Rizwan Parvez. 2025. Weboperator: Action-aware tree search for autonomous agents in web environment. [arXiv preprint arXiv:2512.12692](#).
- A. Drouin, M. Gasse, M. Caccia, and 1 others. 2024. Workarena: How capable are web agents at solving common knowledge work tasks? In [Forty-first International Conference on Machine Learning](#).
- A. Gandhi, G. Neubig, and M. L. Dihan. 2025. Gobrowse: Training web agents with structured exploration. [arXiv preprint arXiv:2506.03533](#).
- S. He, Y. Cui, X. Ma, and 1 others. 2025. Branch-and-browse: Efficient and controllable web exploration with tree-structured reasoning and action memory. [arXiv preprint arXiv:2510.19838](#).
- Su Kara, Fazole Faisal, and Suman Nath. 2025. Waber: Evaluating reliability and efficiency of web agents with existing benchmarks. In [ICLR 2025 Workshop on Foundation Models in the Wild](#).
- J. Y. Koh, S. McAleer, D. Fried, and 1 others. 2024. Tree search for language model agents. [arXiv preprint arXiv:2407.01476](#).
- Ning Li, Qiqiang Lin, Zheng Wu, Xiaoyun Mo, Weiming Zhang, Yin Zhao, Xiangmou Qu, Jiamu Zhou, Jun Wang, Congmin Zheng, and 1 others. 2025. Col-oragent: Building a robust, personalized, and interactive os agent. [arXiv preprint arXiv:2510.19386](#).
- Yanda Li, Chi Zhang, Wenjia Jiang, Wanqi Yang, Bin Fu, Pei Cheng, Xin Chen, Ling Chen, and Yun-chao Wei. 2024. Appagent v2: Advanced agent for flexible mobile interactions. [arXiv preprint arXiv:2408.11824](#).
- Junteng Liu, Yunji Li, Chi Zhang, Jingyang Li, Aili Chen, Ke Ji, Weiyu Cheng, Zijia Wu, Chengyu Du, Qidi Xu, Jiayuan Song, Zhengmao Zhu, Wenhu Chen, Pengyu Zhao, and Junxian He. 2025a. Webexplorer: Explore and evolve for training long-horizon web agents. [arXiv preprint arXiv:2509.06501](#).
- Xiao Liu and 1 others. 2025b. Webcoach: Learning from human feedback via cross-session memory guidance. [Nature Machine Intelligence \(In Press\)](#).
- SHengjie Ma, Chenlong Deng, Jiaxin Mao, Jiadeng Huang, Teng Wang, Junjie Wu, Changwang Zhang, and 1 others. 2025. Pou: Proof-of-use to counter tool-call hacking in deepresearch agents. [arXiv preprint arXiv:2510.10931](#).

- Atsuyuki Miyai, Zaiying Zhao, Kazuki Egashira, Atsuki Sato, Tatsumi Sunada, Shota Onohara, Hiromasa Yamaniishi, Mashiro Toyooka, Kunato Nishina, Ryoma Maeda, and 1 others. 2025. Webchorearena: Evaluating web browsing agents on realistic tedious web tasks. [arXiv preprint arXiv:2506.01952](https://arxiv.org/abs/2506.01952).
- Liangbo Ning, Ziran Liang, Zhuohang Jiang, Haohao Qu, Yujuan Ding, Wenqi Fan, Xiao-yong Wei, Shanru Lin, Hui Liu, Philip S Yu, and 1 others. 2025. A survey of webagents: Towards next-generation ai agents for web automation with large foundation models. In [Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2](https://arxiv.org/abs/2506.01952), pages 6140–6150.
- OpenAI. 2025. [Gpt-5 is here](https://openai.com/gpt-5/). <https://openai.com/gpt-5/>.
- Zehan Qi, Xiao Liu, Iat Long Iong, Hanyu Lai, Xueqiao Sun, Wenyi Zhao, Yu Yang, Xinyue Yang, Jidai Sun, Shuntian Yao, and 1 others. 2024. Webrl: Training llm web agents via self-evolving online curriculum reinforcement learning. [arXiv preprint arXiv:2411.02337](https://arxiv.org/abs/2411.02337).
- J. Shen, A. Jain, Z. Xiao, and 1 others. 2024. Scribeagent: Towards specialized web agents using production-scale workflow data. [arXiv preprint arXiv:2411.15004](https://arxiv.org/abs/2411.15004).
- Segev Shlomov, Alon Oved, Sami Marreed, Ido Levy, Offer Akrabi, Avi Yaeli, Lukasz Strak, Elizabeth Koumpan, Yinon Goldshtein, Eilam Shapira, and 1 others. 2025. From benchmarks to business impact: Deploying ibm generalist agent in enterprise production. [arXiv preprint arXiv:2510.23856](https://arxiv.org/abs/2510.23856).
- Paloma Sodhi, SRK Branavan, Yoav Artzi, and Ryan McDonald. 2023. Step: Stacked llm policies for web actions. [arXiv preprint arXiv:2310.03720](https://arxiv.org/abs/2310.03720).
- Haoming Wang, Haoyang Zou, Huatong Song, Jiazhan Feng, Junjie Fang, Junting Lu, Longxiang Liu, Qinyu Luo, Shihao Liang, Shijue Huang, and 1 others. 2025a. Ui-tars-2 technical report: Advancing gui agent with multi-turn reinforcement learning. [arXiv preprint arXiv:2509.02544](https://arxiv.org/abs/2509.02544).
- Teng Wang, Zhangyi Jiang, Zhenqi He, Shenyang Tong, Wenhan Yang, Yanan Zheng, Zeyu Li, Zifan He, Hailei Gong, Zewen Ye, and 1 others. 2025b. Towards hierarchical multi-step reward models for enhanced reasoning in large language models. [arXiv preprint arXiv:2503.13551](https://arxiv.org/abs/2503.13551).
- Teng Wang, Rong Shan, Jianghao Lin, Junjie Wu, Tianyi Xu, Jianping Zhang, Wenteng Chen, Changwang Zhang, Zhaoxiang Wang, Weinan Zhang, and Jun Wang. 2026. Oscar: Optimization-steered agentic planning for composed image retrieval. [Preprint, arXiv:2602.08603](https://arxiv.org/abs/2602.08603).
- Teng Wang, Wing-Yin Yu, Zhenqi He, Zehua Liu, HaileiGong HaileiGong, Han Wu, Xiongwei Han, Wei Shi, Ruifeng She, Fangzhou Zhu, and 1 others. 2025c. Bpp-search: Enhancing tree of thought reasoning for mathematical modeling problem solving. In [Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](https://arxiv.org/abs/2506.01952), pages 821–838.
- Z. Z. Wang, J. Y. Mao, D. Fried, and 1 others. 2024a. Agent workflow memory. [arXiv preprint arXiv:2409.07429](https://arxiv.org/abs/2409.07429).
- Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. 2024b. Agent workflow memory. [arXiv preprint arXiv:2409.07429](https://arxiv.org/abs/2409.07429).
- Zhepei Wei, Wenlin Yao, Yao Liu, Weizhi Zhang, Qin Lu, Liang Qiu, Changlong Yu, Puyang Xu, Chao Zhang, Bing Yin, and 1 others. 2025. Webagent-r1: Training web agents via end-to-end multi-turn reinforcement learning. [arXiv preprint arXiv:2505.16421](https://arxiv.org/abs/2505.16421).
- Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. A survey of human-in-the-loop for machine learning. [Future Generation Computer Systems](https://doi.org/10.1007/978-1-4939-9832-1_135), 135:364–381.
- Tianci Xue, Weijian Qi, Tianneng Shi, Chan Hee Song, Boyu Gou, Dawn Song, Huan Sun, and Yu Su. 2025. An illusion of progress? assessing the current state of web agents. [arXiv preprint arXiv:2504.01382](https://arxiv.org/abs/2504.01382).
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. [arXiv preprint arXiv:2505.09388](https://arxiv.org/abs/2505.09388).
- Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. 2023. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. [arXiv preprint arXiv:2310.11441](https://arxiv.org/abs/2310.11441).
- Jingbo Yang, Bairu Hou, Wei Wei, Shiyu Chang, and Yujia Bao. 2025b. Webdart: Dynamic decomposition and re-planning for complex web tasks. [arXiv preprint arXiv:2510.06587](https://arxiv.org/abs/2510.06587).
- K. Yang, Y. Liu, S. Chaudhary, and 1 others. 2025c. Agentoccam: A simple yet strong baseline for llm-based web agents. In [The Thirteenth International Conference on Learning Representations](https://arxiv.org/abs/2506.01952).
- Ke Yang, Yao Liu, Sapana Chaudhary, Rasool Fakoore, Pratik Chaudhari, George Karypis, and Huzefa Rangwala. 2024. Agentoccam: A simple yet strong baseline for llm-based web agents. [arXiv preprint arXiv:2410.13825](https://arxiv.org/abs/2410.13825).
- Fabio Massimo Zanzotto. 2019. Human-in-the-loop artificial intelligence. [Journal of Artificial Intelligence Research](https://doi.org/10.1007/978-1-4939-9832-1_135), 64:243–252.
- Chaoyun Zhang, Shilin He, Jiayu Qian, Bowen Li, Liqun Li, Si Qin, Yu Kang, Minghua Ma, Guyue Liu, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. 2024. Large language model-brained gui agents: A survey. [arXiv preprint arXiv:2411.18279](https://arxiv.org/abs/2411.18279).

- Chi Zhang, Zhao Yang, Jiakuan Liu, Yanda Li, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2025a. Appagent: Multimodal agents as smartphone users. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, pages 1–20.
- R. Zhang, M. Qiu, Z. Tan, and 1 others. 2025b. Symbiotic cooperation for web agents: Harnessing complementary strengths of large and small llms. CoRR.
- Y. Zhang, Z. Ma, Y. Ma, and 1 others. 2025c. Webpilot: A versatile and autonomous multi-agent system for web task execution with strategic exploration. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 39, pages 23378–23386.
- Yao Zhang, Zijian Ma, Yunpu Ma, Zhen Han, Yu Wu, and Volker Tresp. 2025d. Webpilot: A versatile and autonomous multi-agent system for web task execution with strategic exploration. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 39, pages 23378–23386.
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, and Graham Neubig. 2023. Webarena: A realistic web environment for building autonomous agents. arXiv preprint arXiv:2307.13854.

A Benchmark Details

WebArena. WebArena is a reproducible, realistic benchmark for evaluating web-browsing agents that must execute high-level natural-language objectives through real browser interactions. The environment consists of four self-hosted websites covering distinct domains (online shopping, discussion forum, collaborative development, and an admin/CMS site), plus auxiliary tool sites (e.g., maps). It contains 812 long-horizon tasks specified as user objectives. Agents observe the environment through structured representations such as HTML/DOM or an accessibility-tree, and interact via a compound action space that simulates mouse/keyboard and multi-tab browsing. Concretely, the action space includes element-level actions (click/hover/type/press/scroll), tab management (new_tab/tab_focus/tab_close), navigation (goto/go_back/go_forward), and a terminal action stop[answer].

WebChoreArena. WebChoreArena extends WebArena with a focus on more realistic and tedious ‘web chores’ that require longer procedures and tighter state constraints, while keeping the same websites, interaction protocol, and action space so that agents are directly comparable. It contributes 532 newly curated tasks, including cross-site workflows. Both benchmarks emphasize functional correctness: success is determined by whether the final state (and/or produced answer) satisfies the task requirements, rather than matching a specific trajectory.

Queries and evaluation. Task queries are natural-language objectives written by human annotators to reflect multi-step, compositional workflows. Evaluation uses two complementary mechanisms: (i) answer-based checks for information-seeking tasks, using exact-match / must-include rules and optionally semantic (LLM-based) fuzzy matching; and (ii) programmatic checks for interaction tasks, where a validator inspects the final web/database state via site APIs, database queries, or DOM/JS selectors and verifies goal constraints. A simple unified view is:

$$\text{Success} = \begin{cases} \mathbb{I}[\text{Match}(\hat{a}, a^*)], & \text{answer-based} \\ \mathbb{I}[\text{Validate}(s_T)], & \text{programmatic} \end{cases}$$

where \hat{a} is the agent answer (if any), a^* is the reference, and s_T is the terminal environment state.

| Benchmark | #Tasks | Focus |
|---------------|--------|--------------------------------------|
| WebArena | 812 | General realistic web tasks |
| WebChoreArena | 532 | More tedious/longer chore-like tasks |

Table 3: Summary of the two benchmarks used in this work.

B Evaluation Protocol and Human Intervention Policy

B.1 AKB Construction and Refinement Workflow

The construction of the Adaptive Knowledge Base (AKB) adheres to a structured, two-phase iterative process designed to minimize human intervention while maximizing agent adaptability:

1. **Exploratory Deployment Phase:** An initial deployment of the base agent is conducted on WebArena tasks without external knowledge support. During this phase, Adaptor autonomously logs execution anomalies and task failures.
2. **Knowledge Distillation Phase:** For identified failures, human experts analyze the interaction traces to diagnose root causes (e.g., non-standard navigation logic, latent UI states). These insights are formalized into generalizable *heuristic entries* following the annotation guidelines (Appendix C) and integrated into the AKB. This process is sparse and efficient, requiring minimal expert effort (approximately one person-day for the entire benchmark).

B.2 Rigorous WebArena Evaluation Protocol

To guarantee evaluation strictness and reproducibility, we enforce the following constraints:

- **Static Knowledge Base:** The AKB is strictly **frozen** prior to the final evaluation. No dynamic knowledge injection or modifications are permitted during inference.
- **Autonomous Execution:** All 812 WebArena tasks are executed in a strictly sequential, fully autonomous manner.
- **Absence of Test-Time Intervention:** The agent must rely exclusively on the pre-compiled knowledge within the static AKB. No human-in-the-loop feedback or real-time correction is accessible during the evaluation phase.

B.3 WebChoreArena Transferability Protocol

For the WebChoreArena benchmark, we adopt a strict cross-task transfer learning protocol to assess generalization:

- **Zero-Shot Exposure:** The agent is strictly precluded from exposure to WebChoreArena tasks during the knowledge acquisition phase. The AKB comprises solely of heuristics derived from WebArena environments.
- **Direct Cross-Task Transfer:** We evaluate the agent’s capacity to generalize site-specific priors learned from WebArena to the more complex query formulations in WebChoreArena, without any task-specific fine-tuning or supplementary knowledge curation.

C AKB Tip Annotation Guideline

C.1 Tip Granularity and Scope

Tips must capture **site-level operational logic** rather than query-specific execution traces.

- **Bad (Query-Specific):** "Click the blue button to buy the red socks." (Only valid for one query)
- **Good (Site-Level):** "To add any item to the cart, you must first select a size option, even if only one size is available." (Valid for all purchase tasks on the site)

C.2 Recommended Tip Template

Annotators are encouraged to follow this structure:

- **Scope:** Define when the tip applies (e.g., "For all product pages...").
- **Action:** Describe the required interaction pattern (e.g., "Hover over the menu...").
- **Constraint:** Mention any critical preconditions (e.g., "...before clicking the link.").
- **Goal Alignment:** Explain why this is necessary (e.g., "...to reveal the submenu.").

D AKB Statistics and Example Tips

D.1 Tip Statistics from WebArena Failures

The final frozen AKB contains a compact set of tips distributed across the domains:

| Domain | Number of Tips |
|----------------|----------------|
| GitLab | 13 |
| Map | 7 |
| Reddit | 5 |
| Shopping | 9 |
| Shopping Admin | 18 |

Table 4: Count of crystallized tips per domain in the Adaptive Knowledge Base.

D.2 Representative Tips by Domain

Below are representative examples of crystallized knowledge:

GitLab:

- "When asked to make changes to a repository, commit changes to the main branch UNLESS the task specifically asks you to create a branch."
- "To add new members to the project, you can visit project information > members tab and click blue "invite members" button on top right."

Reddit:

- "if you need find a relevant subreddit or forum, you can find the name after clicking "alphabetical" in the "Forum" tab."
- "When creating a post/submission, put all content including URLs in the Body field. Leave the URL field empty."

Shopping Admin:

- "When you add a new product in the CATALOG > Products tab, you can click the downward arrow beside the "Add Product" button to select options like "Simple Product", "Configurable Product", etc. "
- "You can generate various reports by using menus in the REPORTS tab. Select REPORTS > "report type", select options, and click "Show Report" to view report."

Shopping:

- "If you need to find information about your previous purchases, you can go My Account > My Orders, and find order by date, order number, or any other available information"
- "An order is considered out of delivery if it is marked as "processing" in the order status."

E Prompt Templates

E.1 Operator Prompt

The following system prompt guides the low-level execution agent (Operator).

Role: You are a UI Assistant, your goal is to help the user perform tasks using a web browser. You can communicate with the user via a chat, to which the user gives you instructions and to which you can send back messages. You have access to a web browser that both you and the user can see, and with which only you can interact via specific commands.

Instructions: Review the instructions from the user, the current state of the page and all other information to find the best possible next action to accomplish your goal. Your answer will be interpreted and executed by a program, make sure to follow the formatting instructions.

Output Requirements: Your response must strictly follow this output format:

<think>

Explain your reasoning step-by-step here. Describe how you analyze the current page state and your plan to accomplish the user's goal.

</think>

<action>

Only a **SINGLE** action is allowed in this tag, which will be executed by the program.

</action>

Formatting Rules:

- The <think> tag should contain a clear and detailed explanation of your reasoning process, including how you interpret the task, analyze the page, and decide on your next step.
- Inside the <action> tag, you **MUST** provide exactly one valid action string formatted as ““ “. This represents the single action to be executed. The action must be one of the valid actions from the action space described below.

E.2 ColorSummarizer Prompt

The Summarizer Agent uses the following prompt to compress history, ensure alignment with re-

trieved knowledge, and provide guidance.

Role: You are a Progress Summarization and Knowledge Alignment Agent. Your role is to: 1. Analyze the current execution context to track progress. 2. Verify alignment between the agent's actions and the provided expert knowledge (tips). 3. Generate a concise, actionable summary to guide the browser automation agent.

Inputs Provided:

- goal: The ultimate task objective.
- relevant_knowledge: Key tips or site-specific constraints retrieved from the Adaptive Knowledge Base (AKB).(e.g., "On this site, use the search bar instead of category links.")
- axtree_txt: The accessibility tree of the current webpage.
- screenshot: A visual rendering of the current webpage.
- action_history: Sequence of actions already performed.
- previous_summary: The progress summary from the previous step.

Output Requirements: Generate a structured summary with exactly three sections:

- **Current Progress & Knowledge Check**
 - State clearly which sub-goals have been completed.
 - Explicitly evaluate if the recent actions align with relevant_knowledge.
 - If a specific tip was violated (e.g., agent clicked a link when the tip suggested searching), flag this deviation immediately.
- **Current State Analysis**
 - Assess whether the current page is on the critical path toward goal.
 - Identify relevant interactive elements (buttons, inputs) necessary for the next step.

- **Next-step Guidance** (Conditional Section)

- Include this section if the agent has deviated from the goal OR violated relevant_knowledge.
- Provide 1–2 concrete recommendations.
- If available, quote the specific tip from relevant_knowledge that dictates the correct next step.

Formatting Rules:

- Always include the Knowledge Check within the Progress section.
- If relevant_knowledge is empty, proceed with standard progress tracking.

F Practical Insights and Future Directions

The development and deployment of ColorBrowser-Agent have yielded valuable insights into the design of human-centred autonomous systems, shedding light on the pragmatic challenges of integrating agents into human workflows.

Engineering for Trustworthy Collaboration.

While novel architectures often capture the spotlight, our experience underscores that rigorous engineering optimization is the bedrock of user trust. Seemingly mundane enhancements—such as robust error handling for network timeouts, atomic action primitives, and defensive navigation strategies—are essential for transforming an agent from a fragile prototype into a reliable digital partner. In a human-centred context, reliability is not just a metric but a prerequisite for adoption; users will only delegate tasks to agents that demonstrate consistent stability in real-world conditions.

Symbiotic Human-AI Collaboration. Our results validate that Human-in-the-Loop integration is not merely a fallback mechanism but a powerful paradigm for **Collaborative Autonomy**. By allowing the agent to request help, we create a symbiotic loop: the AI scales execution for standard tasks, while the human provides high-level intuition for edge cases. A critical avenue for future research is optimizing this interaction cost—developing methods to dynamically identify the "minimal viable intervention" so that the system maximizes the

leverage of expert priors while minimizing human cognitive load. The ultimate goal is an agent that learns from its user’s guidance, progressively requiring less intervention over time.

Foundation Models as Reasoning Engines. Finally, our findings reaffirm the critical role of strong foundation models in interpreting human intent. The agent’s ability to generalize to novel sites relies heavily on the reasoning capabilities of the underlying LLM. We observe that a capable model (like GPT-5) serves as the bridge between abstract human instructions and concrete interface actions. As these models evolve, we envision a future where agents can infer not just "how" to click, but "why" a user prefers a certain workflow, leading to more personalized and intuitive automation.

G Algorithms

Algorithm 1 presents the unified process of ColorBrowserAgent, encompassing both the offline construction of the Adaptive Knowledge Base and the online execution phase.

Algorithm 1 ColorBrowserAgent: Knowledge Adaptation and Execution

Require: Training Tasks $\mathcal{T}_{\text{train}}$, Test Instruction g_{test} , Environment \mathcal{E} , Human Expert \mathcal{H}

Ensure: Adaptive Knowledge Base \mathcal{K} , Interaction Trajectory h

1: **Phase 1: Offline Knowledge Adaptation Loop**

2: Initialize $\mathcal{K} \leftarrow \emptyset$

3: **for all** task $g \in \mathcal{T}_{\text{train}}$ **do**

4: Run base agent on g , monitor with f_{trigger}

5: **if** failure detected **then**

6: Expert \mathcal{H} analyzes failure and extracts site-level logic τ

7: $\mathcal{K} \leftarrow \mathcal{K} \cup \{\tau\}$ {Inject into AKB}

8: **end if**

9: **end for**

10: **Phase 2: Online Execution Loop (Inference)**

11: Initialize history $h_0 \leftarrow \emptyset$, belief state $m_0 \leftarrow \emptyset$

12: Target goal $g \leftarrow g_{\text{test}}$

13: **for** $t = 0$ **to** T_{max} **do**

14: $o_t \leftarrow \mathcal{E}.\text{observe}()$ {Obtain observation}

15: **// Adaptor: Retrieve Knowledge**

16: $k_t \leftarrow \text{Retrieve}(\mathcal{K}, o_t)$

17: **// Summarizer: Update Belief State**

18: $m_t \leftarrow \text{Summarize}(m_{t-1}, o_t, k_t)$ {Align with knowledge}

19: **// Operator: Execute Action**

20: $a_t \leftarrow \text{Policy}(o_t, m_t, k_t)$

21: $\mathcal{E}.\text{step}(a_t)$

22: $h_{t+1} \leftarrow h_t \cup \{o_t, a_t\}$

23: **if** a_t is stop **then**

24: **return** h_{t+1}

25: **end if**

26: **end for**

H Case Study

To intuitively illustrate the execution of ColorBrowserAgent, we present a failure trajectory in the absence of prior knowledge and demonstrate how injecting priors resolves the issue.

Table 5: The failure trajectory of ColorBrowserAgent without prior knowledge.

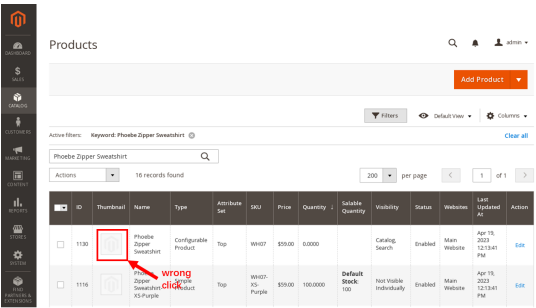
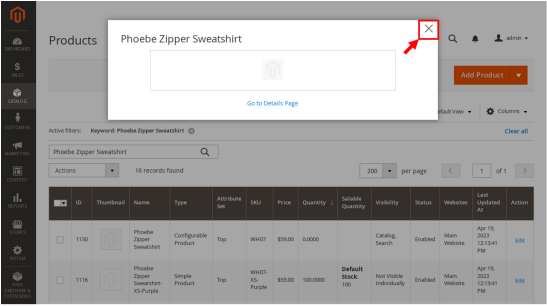
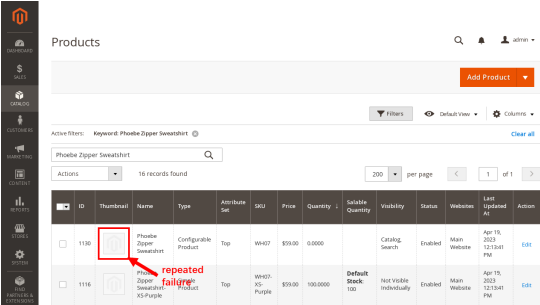
| Query: Add a new color option brown to the size S of Phoebe Zipper Sweatshirt | | |
|---|---|-------------------------------|
| Screenshot | Thought | Remark |
|  | <p><think> The user wants to add a new color option "brown" to size S for the product "Phoebe Zipper Sweatshirt." I'm on the Magento Admin Products grid page with a keyword filter applied for "Phoebe Zipper Sweatshirt," and the grid shows 16 results. To proceed, I need to open the specific product's edit page. The SOM context lists a grid cell for the Name "Phoebe Zipper Sweatshirt" with bid 1773, which is typically clickable and opens the product edit view. This is the logical next step before accessing configurations to add the new color option for size S. I will click the product name cell to open its edit page. </think></p> | <p>Navigating to Product.</p> |
|  | <p><think> The goal is to edit the configurable product "Phoebe Zipper Sweatshirt" to add a new color option brown for size S. I previously clicked the product name in the grid to open its edit page, but the current screenshot still shows the Products grid and a modal dialog overlay near the top (it appears to be a grid-related modal like "Reset Grid Columns"), which likely blocked navigation. To proceed, I need to dismiss this modal. Since I don't have the specific bid for the modal's Cancel/Close button, the most reliable way to close it is to press Escape, which commonly dismisses modals in Magento's admin UI. After the modal is closed, I will be able to click into the product and continue with the configuration changes. </think></p> | <p>Missed Button.</p> |
|  | <p><think> The task is to add a new color option "brown" to size S for the product "Phoebe Zipper Sweatshirt" in Magento. I'm on the Products grid page with the keyword filter applied, and earlier a modal overlay prevented navigation when I tried to open the product. I dismissed the modal by pressing Escape, so the next logical step is to open the product's edit page. The product name cell "Phoebe Zipper Sweatshirt" (bid 1773) is typically clickable to navigate to the edit view for a configurable product. Opening this page is necessary before accessing the Configurations section to add the new color-size variant. I will click the product name cell now. </think></p> | <p>Trapped in Loop.</p> |

Table 6: The success trajectory of ColorBrowserAgent with injected knowledge.

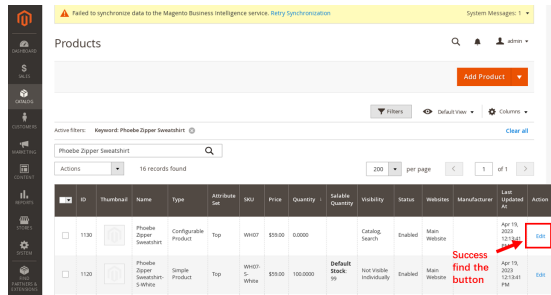
Query: Add a new color option brown to the size S of Phoebe Zipper Sweatshirt

Knowledge: If you need to add new attribute values (e.g. size, color, etc) to a product, you can find the product at CATALOG > Products, search for the product, edit product with “Configurable Product” type, and use “**Edit Configurations**” to add the product with new attribute values. If the value that you want does not exist, you may need to add new values to the attribute.

Screenshot

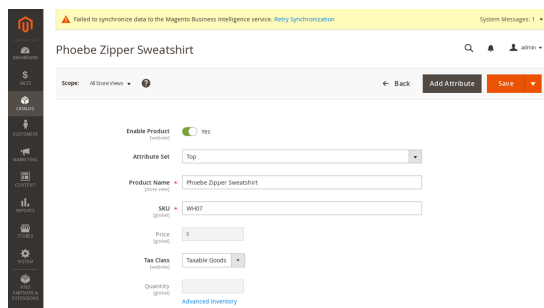
Thought

Remark



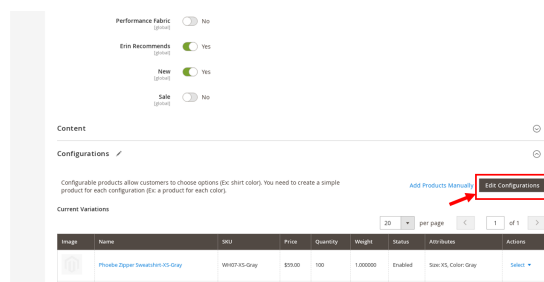
<think> The task is to add a new color option “brown” specifically for size S of the “Phoebe Zipper Sweatshirt,” which is a configurable product. On the current Products page, the grid lists the parent configurable product and all its child simples. To modify configurations (add the S-Brown variant), I need to open the parent configurable product’s edit page. In the table, I see the row for “Phoebe Zipper Sweatshirt” with Type “Configurable Product” and an Action cell containing the link “**Edit Phoebe Zipper Sweatshirt**”. The correct next step is to click this edit link to access the product editing interface, where I can use “**Edit Configurations**” to add the new color option. If “brown” doesn’t exist as a Color attribute value, I may need to add it later under STORES > Attributes > Product, but first I must open the product. </think>

Navigating to Product.



<think> The goal is to add a new color option Brown to size S for the configurable product “Phoebe Zipper Sweatshirt.” I am currently on the product search page and have located the item. The next step is to enter the product edit page. However, the specific configuration button “Edit Configurations” is not immediately visible in the current viewport. I need to scroll down the page to locate the “Configurations” section where the attribute modification options are typically found. Therefore, the immediate action is to scroll down to reveal the configuration interface. </think>

Locating Config.



<think> The goal is to add a new color option Brown to size S for the configurable product “Phoebe Zipper Sweatshirt.” I am already on the product’s edit page, and I can see the Configurations section with the “Edit Configurations” button. Since Size and Color are the configurable attributes (their comboboxes are disabled, indicating they’re tied to configurations), the correct workflow is to open the configurations wizard via “**Edit Configurations**.” In that wizard, I will select Size=S and Color=Brown, then proceed through the steps to generate the new simple product variant. The Brown option already exists in the Color attribute list (observed in the disabled combobox options), so I do not need to add a new attribute value in STORES. The next step is to click “**Edit Configurations**” to start the wizard. </think>

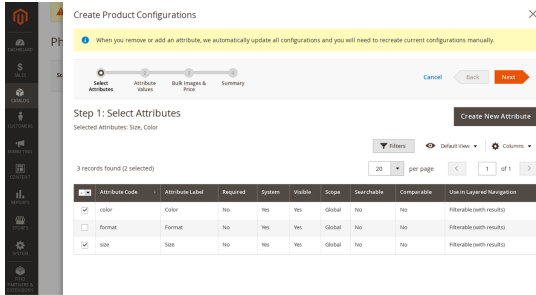
Starting Wizard.

Continued on next page

Table 6: The success trajectory of ColorBrowserAgent with injected knowledge (continued).

Query: Add a new color option brown to the size S of Phoebe Zipper Sweatshirt

Screenshot



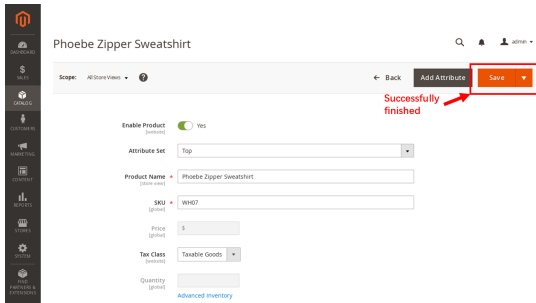
Thought

<think> The user wants to add a new Color option Brown specifically to Size S for the configurable product. I have the “Create Product Configurations” modal open on Step 1: Select Attributes, with Size and Color already selected, which matches the attributes we need. To proceed to choosing specific attribute values (where I can select Size S and Color Brown only), I need to go to Step 2: Attribute Values. **The “Next” button is present. Clicking it should take me to the values selection screen.** </think>

Remark

Wizard Step 1.

...



<think> The user’s goal is to add a new color option "Brown" to Size "S" for the product "Phoebe Zipper Sweatshirt." I’m on the product edit page for the configurable product, and the Configurations section shows the "Current Variations" grid with a row "Phoebe Zipper Sweatshirt-S-Brown," SKU "WH07-S-Brown," Status "Enabled," and Attributes "Size: S, Color: Brown." This indicates the S–Brown variant exists and is enabled. I already executed a Save action earlier, and the variant appears in the grid, which implies the configuration was successfully generated and saved. **The best next step is to confirm completion to the user, including the variant name and SKU, aligning with the requirement to provide a complete answer that stands alone.** </think>

Task Completed.