

An Autonomous RL Agent Methodology for Dynamic Web UI Testing in a BDD Framework

Ali Hassaan Mughal¹

¹ Computer Science, Kansas State University, USA

Abstract

Modern software applications demand efficient and reliable testing methodologies to ensure robust user interface functionality. This paper introduces an autonomous reinforcement learning (RL) agent integrated within a Behavior-Driven Development (BDD) framework to enhance UI testing. By leveraging the adaptive decision-making capabilities of RL, the proposed approach dynamically generates and refines test scenarios aligned with specific business expectations and actual user behavior. A novel system architecture is presented, detailing the state representation, action space, and reward mechanisms that guide the autonomous exploration of UI states. Experimental evaluations on open-source web applications demonstrate significant improvements in defect detection, test coverage, and a reduction in manual testing efforts. This study establishes a foundation for integrating advanced RL techniques with BDD practices, aiming to transform software quality assurance and streamline continuous testing processes.

Received on 13 March 2025; accepted on 14 July 2025; published on 21 July 2025

Keywords: Reinforcement Learning, UI Testing, BDD, Automated Testing, Web Applications

Copyright © 2025 Ali Hassaan Mughal, licensed to EAI. This is an open access article distributed under the terms of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>), which permits unlimited use, distribution and reproduction in any medium so long as the original work is properly cited.

doi:10.4108/airo.8895

*Corresponding author. Email: alihassaanmughal@hotmail.com

Problem Statement

Modern web applications are characterized by their dynamic and intricate navigational structures, where each page and user interaction forms part of a complex network. Traditional testing methodologies often fall short in fully exploring these complex pathways, especially when user actions are interdependent and context-sensitive. This research proposes the integration of an autonomous reinforcement learning (RL) agent within a Behavior-Driven Development (BDD) framework to automate and enhance web UI testing by modeling the website as a maze.

The core idea is to define constant starting points—such as a homepage, login page, or any predefined entry condition—and distinct endpoints that signal the completion of a user scenario. Endpoints are identified by cues such as specific text presence, visible UI elements, or states where no further actions are possible. For example, in the context of an e-commerce platform like Amazon, the agent would:

- Begin at a fixed starting point (e.g., the Amazon homepage or sign-in page).

- Navigate through the website by detecting and interacting with UI elements, such as searching for a product, adding it to the cart, and proceeding to checkout.
- Recognize endpoints such as a successful login confirmation, product detail verification, or an order confirmation page.

An **input summary** is provided to the agent, specifying the functionality to be tested and establishing the starting point. For example, the input summary may define a scenario like “Create an order for a specific product type” or “Post a review for a selected product category.” The RL agent learns to maneuver between these defined points, autonomously exploring multiple pathways and uncovering various possibilities to reach the desired end state.

Key aspects of the problem include:

1. **Modeling the Website as a Maze:** Treating each webpage as a node and each user interaction as a transition captures the inherent complexity of web navigation.

2. **Defining Constant Starting Points:** Establishing consistent entry points such as the homepage, login page, or other specific landing pages.
3. **Identifying Endpoints:** Clearly defining termination points based on cues such as confirmation messages, error notifications, or terminal states.
4. **Input Summary for Test Scenarios:** Using a concise summary to indicate the targeted functionality, which allows the RL agent to concentrate on achieving the scenario-specific goal.
5. **Learning and Exploration:** Allowing the RL agent to learn optimal pathways through interactions with the UI, thereby uncovering multiple valid routes from the starting point to the endpoint, and capturing diverse test scenarios.

This framework not only facilitates the generation and execution of test cases but also adapts dynamically to the evolving structure of modern web applications. The integration of RL within a BDD context ensures that the resulting tests are both comprehensive and interpretable, bridging the gap between automated testing efficiency and human-understandable test scenarios.

1. Introduction

Modern software applications demand robust, scalable, and efficient testing methodologies to guarantee high-quality user interface functionality. Traditional manual and scripted testing approaches have proven insufficient in the face of dynamic web environments, where user interactions resemble navigating through a complex maze. Recent advances in reinforcement learning (RL) and Behavior-Driven Development (BDD) offer promising alternatives. This paper presents an autonomous testing agent that leverages RL and integrates with a BDD framework to generate, execute, and refine test scenarios aligned with specific business requirements and real user behaviors.

The concept of modeling websites as mazes, with fixed starting points (e.g., homepage, login page) and well-defined endpoints (e.g., confirmation pages, error states), is central to this approach. In this system, an *input summary* describing the target functionality (e.g., “create an order” or “post a review”) serves as the guide for the RL agent. The agent learns to explore multiple routes between these points, identifying valid transitions and capturing diverse test scenarios. This method is intended to enhance test coverage, increase defect detection, and reduce the need for manual test script maintenance.

Autonomous agents have been explored in various domains. Early environments for testing and automation [1, 2] laid the groundwork for realistic web interaction. Subsequent works extended these ideas by incorporating RL techniques to overcome the limitations of static and scripted test cases [3–5]. In parallel, studies focusing on HTML understanding and multimodal perception have demonstrated the potential of large language models (LLMs) and visual language models in parsing complex web pages and generating actionable insights [6–8]. The integration of these technologies is crucial for building agents capable of seamless interaction with diverse and evolving web interfaces.

Recent advances in autonomous agent design have shown that incorporating elements such as state-space exploration, reasoning traces, and dynamic feedback can significantly improve performance on real-world tasks [9–13]. The present work draws on these developments to integrate an RL agent that not only explores web environments but also learns to follow BDD scenarios, thereby bridging the gap between automated execution and human-readable test cases.

2. Literature Review

The literature on autonomous web agents, RL-based testing, and BDD integration spans several interconnected domains. This section reviews the work in three primary categories: realistic web environments and datasets, RL and LLM-based autonomous agents, and BDD/automated UI testing.

2.1. Realistic Web Environments and Datasets

Several studies have focused on constructing realistic environments for training and evaluating web agents. For instance, [1] introduced *WebArena*, a comprehensive environment that emulates real-world websites across multiple domains, thereby addressing the shortcomings of synthetic testbeds. Similarly, [3] presented *Mind2Web*, a dataset spanning numerous websites and tasks, specifically designed for evaluating generalist web agents. Earlier platforms such as *World of Bits* [2] laid the foundational concepts by representing web pages as complex, interactive environments that serve as benchmarks for RL agents. These works collectively underscore the need for environments that capture the dynamic and intricate nature of real-world web interfaces.

2.2. RL and LLM-based Autonomous Agents

Reinforcement learning is a pivotal technique for developing autonomous agents for sequential decision-making tasks. For example, [4] explored the use of RL for generating UI actions based on pixel inputs, demonstrating that agents can learn effective strategies

by mimicking human interactions. In the domain of computer control, [14] provided large-scale datasets for Android device control, emphasizing the importance of multi-step tasks that require both semantic and visual understanding.

More recent studies have focused on integrating LLMs into autonomous agent architectures. [15] proposed the ReWOO framework to decouple reasoning from observations, significantly improving token efficiency during test case generation. The *CogAgent* model [8] demonstrated the capability of visual language models to comprehend and navigate graphical user interfaces, which is critical for realistic UI testing. Additionally, [16] introduced Synatra, an approach that converts indirect knowledge (e.g., online tutorials) into direct demonstrations, enabling scalable training for digital agents.

Advancements in multimodal understanding have been particularly influential. Research on instruction-finetuned models for multimodal web navigation [6] and HTML understanding [7] has shown that combining textual and visual inputs leads to more robust agent performance. Autonomous agents such as those developed in [9, 17, 18] integrate both visual and textual modalities to interact effectively with real-world websites. The integration of planning and long-context understanding, as demonstrated by [19] and [20], further underscores the importance of combining various learning paradigms to achieve high performance in complex tasks.

Recent studies, including [21] (AgentBench) and [10] (Ferret-UI), have benchmarked the performance of LLM-based agents, highlighting the gap between human-level performance and current state-of-the-art models in diverse testing scenarios. Techniques emphasizing executable actions and dynamic feedback [22, 23] provide further insights into improving agent performance by integrating recursive critique and refinement strategies. Moreover, [24] highlights the importance of harmonizing natural language and code to create agents capable of both reasoning and execution. [25] have implemented a module that uses DNNs to be trained for RL agents, similar to the ideas of trained RL agents for Web UI Testing.

2.3. BDD and Automated UI Testing

Behavior-Driven Development (BDD) has become an industry-standard practice to ensure that software behavior aligns with user expectations. The integration of RL with BDD, as proposed in this paper, leverages the strengths of both paradigms. Traditional automated UI testing often relies on scripted or rule-based approaches, which lack adaptability in dynamic web environments. In contrast, the proposed framework

allows an RL agent to autonomously generate BDD scenarios by learning directly from web interactions.

Efforts to bridge the gap between automated testing and human-readable specifications are evident in related research. For example, studies on dynamic scenario reusability and step auto-completion for frameworks like Cucumber [26] demonstrate the need for intelligent agents that integrate smoothly with existing BDD workflows. Research on automated test generation [11, 12, 27, 28] further supports the theoretical and practical basis for the proposed system.

2.4. Synthesis and Research Gap

The reviewed literature indicates a clear trend toward leveraging advanced RL techniques, multimodal LLMs, and realistic web environments for building autonomous agents. However, significant challenges persist in integrating these technologies within a unified framework for UI testing. Although individual components—such as realistic environment simulation [1–3], RL-based agent learning [4, 9, 14–16, 20], and multimodal understanding [6–8, 18]—have been extensively studied, synthesizing these approaches within a BDD context remains an open research question.

This paper addresses this gap by proposing a framework that integrates an autonomous RL agent with BDD practices for web UI testing. The aim is not only to improve test coverage and defect detection rates but also to generate human-interpretable test cases that are easily incorporated into continuous integration pipelines.

3. Implementation Details

This section describes the implementation of the autonomous RL agent integrated within a BDD framework for web UI testing. The system is designed to automatically generate and refine test scenarios by modeling the website as a maze. The system accepts an input summary that defines the testing functionality (e.g., “Create an order for a specific product”) and navigates from predefined starting points (e.g., homepage or login page) to designated endpoints (e.g., order confirmation).

3.1. Scenario: Creating an Order on an E-Commerce Platform

Consider an e-commerce scenario:

Input Summary: “Place an order for a specific product category (e.g., electronics).”

In this scenario, the RL agent is responsible for:

1. **Initialization:** Loading the homepage and, if required, performing user authentication.

2. **Exploration:** Navigating through the website by accessing menus, filtering by the electronics category, and searching for products.
3. **Action Execution:** Selecting a product, adding it to the shopping cart, proceeding to checkout, and simulating a payment process.
4. **Endpoint Recognition:** Detecting the order confirmation page by identifying cues such as confirmation messages or tracking numbers.
5. **Learning and Refinement:** Updating the policy based on received rewards and refining the exploration strategy using backtracking when dead-ends are encountered.

3.2. System Architecture and Algorithmic Components

The system is composed of the following modules:

- **Input Parser:** Extracts key functionalities from the input summary and defines the target scenario.
- **State Representation Module:** Encodes the current webpage state using a combination of DOM extraction, visual recognition (via convolutional neural networks or transformer-based vision encoders), and text processing.
- **Action Space Definition:** Defines a set of generic actions such as `click(element)`, `type(text, element)`, and `scroll(direction)`. The choice of a discrete set of actions is driven by the need for clarity in execution and ease of mapping to UI events.
- **Reward Mechanism:** Provides intermediate rewards for detecting relevant cues (e.g., product details, cart updates) and final rewards upon reaching an endpoint. The design of the reward function is critical and is tuned to balance exploration and exploitation.
- **Exploration and Learning Engine:** Utilizes RL algorithms to explore the website. In this system, the following algorithms are employed:
 - **Deep Q-Networks (DQN):** This method is suitable for environments with discrete action spaces. DQN approximates the optimal Q-value function, guiding the agent to select actions that maximize cumulative rewards. The algorithm was chosen for its relative simplicity and effectiveness in discrete control tasks.
 - **Policy Gradient Methods:** Techniques such as REINFORCE and Actor-Critic are used for scenarios with larger or continuous action spaces. These methods directly optimize the policy, which can be advantageous when the action space cannot be easily discretized.

Their inclusion provides flexibility and improved performance in complex tasks.

- **Epsilon-Greedy Exploration:** This strategy is used to balance exploration and exploitation. The agent selects random actions with a probability ϵ , which decays over time, allowing for more focused exploration as the policy converges.
- **Backtracking Mechanisms:** Inspired by dynamic programming, these mechanisms allow the agent to revisit earlier states and explore alternative paths when encountering dead-ends. This ensures a more comprehensive exploration of the UI.
- **BDD Integration Module:** Transforms successful navigation trajectories into human-readable BDD scenarios (e.g., in Gherkin syntax), facilitating integration into continuous integration pipelines.

Algorithm 1 Autonomous Web UI Testing Agent

Require: Input summary S , starting state s_0 , set of endpoints E , maximum iterations N

- 1: Initialize agent policy π using DQN/Policy Gradient techniques
 - 2: Initialize exploration parameter ϵ (for epsilon-greedy strategy)
 - 3: Initialize memory $M \leftarrow \emptyset$
 - 4: **for** $i = 1$ to N **do**
 - 5: $s \leftarrow s_0$ {Reset to starting point for each episode}
 - 6: **while** $s \notin E$ **and not** timeout **do**
 - 7: **if** $\text{random}() < \epsilon$ **then**
 - 8: $a \leftarrow$ random action from action space
 - 9: **else**
 - 10: $a \leftarrow_a Q(s, a)$ {Using DQN or sampling from $\pi(s, S)$ }
 - 11: **end if**
 - 12: Execute action a , observe new state s' and reward r
 - 13: Store transition (s, a, r, s') in memory M
 - 14: $s \leftarrow s'$
 - 15: **end while**
 - 16: **if** $s \in E$ **then**
 - 17: Update policy π using the collected trajectory (via backpropagation in DQN or policy gradient update)
 - 18: **else**
 - 19: Apply backtracking: revisit previous states from M and attempt alternative actions
 - 20: **end if**
 - 21: Optionally decay ϵ over time
 - 22: **end for**
 - 23: Translate best-performing trajectories into BDD scenarios (e.g., in Gherkin syntax)
-

3.3. Graphical Workflow

Figure 1 illustrates the step-by-step process of the agent. Nodes are arranged with adequate horizontal and vertical spacing to ensure clarity.

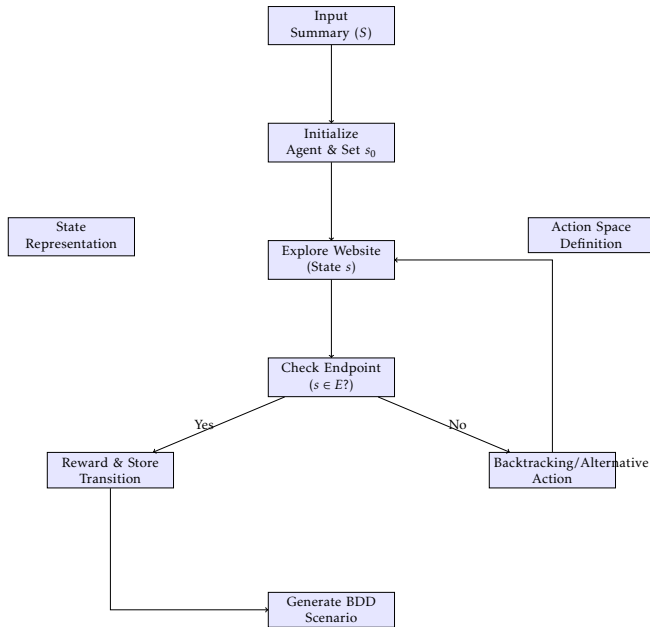


Figure 1. Workflow of the Autonomous Web UI Testing Agent. The agent receives an input summary, initializes with a defined state, explores the website using an RL-based policy (employing DQN, Policy Gradient methods, and epsilon-greedy exploration), checks for endpoint conditions, applies rewards or backtracking, and finally converts successful trajectories into BDD scenarios.

3.4. Discussion of Algorithmic Choices

The selection of specific RL algorithms in this framework has been driven by their suitability for different aspects of the problem:

- **Deep Q-Network (DQN):** DQN is particularly effective in environments where the action space is discrete. It approximates the optimal Q-value function, thereby guiding the agent to select actions that maximize cumulative rewards. The simplicity of DQN makes it a natural choice for tasks where UI interactions can be discretized.
- **Policy Gradient Methods:** In scenarios where the action space is large or continuous, methods like REINFORCE or Actor-Critic provide a direct optimization of the policy. These methods allow for more nuanced control in complex tasks and overcome the limitations associated with the discretization of actions. The Actor-Critic variant combines the benefits of both value-based and policy-based methods, providing more stable convergence.

- **Epsilon-Greedy Exploration:** This exploration strategy ensures that the agent does not prematurely converge to suboptimal policies by balancing random exploration and the exploitation of known good actions. The decay of ϵ over time is crucial for transitioning from exploration to exploitation.
- **Backtracking Mechanisms:** In web environments, dead-ends are common due to the dynamic nature of UIs. The backtracking mechanism enables the agent to recover from such states by revisiting earlier states and attempting alternative paths. This approach is inspired by dynamic programming techniques and enhances overall exploration efficiency.

This integrated system continuously refines its understanding of the website's structure, modeled as a maze, and learns optimal navigation paths from the starting points to the endpoints as defined by the input summary. Successful trajectories are transformed into human-readable BDD scenarios, bridging the gap between automated testing and manual verification.

4. Results and Analysis

This section outlines the potential outcomes and analysis methods to evaluate the performance of the proposed RL-driven BDD UI testing framework. The results can be visualized using various metrics and plots, providing a comprehensive understanding of the agent's performance in navigating and testing web applications.

4.1. Potential Outcomes

The proposed framework is expected to yield improvements in several areas:

- **Test Coverage:** The agent is anticipated to explore a wider variety of navigational paths, potentially uncovering more UI states and identifying hidden defects compared to traditional scripted tests.
- **Defect Detection:** The increase in exploratory actions should result in a higher defect discovery rate, as the agent actively searches for inconsistencies and errors.
- **Reduction in Manual Effort:** The automated generation of BDD test cases can significantly reduce the time and effort required for writing and maintaining test scripts.
- **Adaptive Learning:** Over time, the agent is expected to refine its navigation strategy by learning from past interactions, thereby converging to a more optimal testing policy.

4.2. Evaluation Metrics and Visualization Techniques

Several evaluation metrics and corresponding visualization techniques are proposed:

- **Learning Curves:**
 - **Reward Trajectory:** Plotting the cumulative or average reward per episode as a line graph to assess convergence behavior and policy improvement.
 - **Success Rate Trend:** A line graph or moving average plot illustrating the evolution of the task success rate over time.
- **Test Coverage Visualization:**
 - **Coverage Heatmaps:** Heatmaps to display the frequency of state visits across different website areas, highlighting well-explored and under-explored regions.
 - **Bar Charts:** Comparative bar charts showing the number of unique pages or UI elements visited by the RL agent versus those covered by baseline testing methods.
- **Defect Discovery Analysis:**
 - **Scatter Plots:** Plotting defect discovery rates against the number of episodes or steps, to identify episodes with high defect yields.
 - **Boxplots:** Displaying the distribution of defects found per episode to facilitate comparisons between different experimental setups.
- **Policy and Trajectory Analysis:**
 - **Trajectory Visualization:** Graphically representing sample state-action sequences and overlaying them on a schematic of the website to better understand navigation paths.
 - **Backtracking Frequency:** Plotting the frequency of backtracking events per episode to evaluate how often the agent encounters and recovers from dead-ends.

4.3. Statistical Analysis

To validate the performance of the approach, the following statistical analyses can be performed:

- **Hypothesis Testing:** Conducting t-tests or similar statistical tests to compare key metrics such as success rates and defect discovery rates between the RL-driven approach and baseline methods.

- **Confidence Intervals:** Calculating confidence intervals for average rewards, test coverage, and episode lengths to assess the reliability of the improvements observed.

4.4. Discussion of Potential Results

The analysis is expected to address the following:

1. **Coverage Improvement:** The extent to which the RL agent's exploration strategy leads to higher test coverage compared to traditional methods.
2. **Defect Detection Efficacy:** Whether the exploratory actions yield a significantly higher rate of defect identification.
3. **Learning Dynamics:** Insights into how the learning curve evolves over time and whether the policy converges to an optimal strategy.
4. **Trajectory Insights:** An analysis of successful navigation trajectories and their conversion into reusable BDD test cases.

Standard visualization tools such as matplotlib can be used to generate clear and interpretable graphs, providing valuable insights into both the strengths and limitations of the proposed framework.

5. Conclusion

This paper has presented a novel framework that integrates an autonomous reinforcement learning (RL) agent with a Behavior-Driven Development (BDD) framework for automated UI testing. By modeling the website as a maze with defined starting points and endpoints, the proposed method enables the agent to dynamically generate and refine test scenarios that align with specific business requirements and user behaviors. A comprehensive system architecture was described, including detailed discussions on state representation, action space definition, reward mechanisms, and exploration strategies utilizing established RL algorithms such as Deep Q-Networks and Policy Gradient methods.

The choice of DQN is justified for its efficiency in discrete action spaces, while Policy Gradient methods provide the necessary flexibility for more complex tasks. The combination of epsilon-greedy exploration and backtracking mechanisms further ensures robust and adaptive navigation of dynamic web interfaces. In addition, the conversion of successful trajectories into human-readable BDD scenarios bridges the gap between automated testing and manual verification.

The proposed framework is expected to improve test coverage, enhance defect detection, and reduce manual effort, thereby transforming the process of software quality assurance. Future work will focus

on extending this framework to accommodate more complex web environments and further refining the state representation and reward mechanisms. Moreover, exploring transfer learning and multi-task learning is anticipated to further increase the robustness and generalizability of the RL agent.

References

- [1] ZHOU, S., XU, F.F., ZHU, H., ZHOU, X., LO, R., SRIDHAR, A., CHENG, X. *et al.* (2023) Webarena: A realistic web environment for building autonomous agents. *ArXiv abs/2307.13854*: null. doi:10.48550/arXiv.2307.13854, URL <https://www.semanticscholar.org/paper/e41482f4ee984f17382f6cdd900df094d928be06>.
- [2] SHI, T., KARPATY, A., FAN, L.J., HERNÁNDEZ, J. and LIANG, P. (2017) World of bits: An open-domain platform for web-based agents URL <https://www.semanticscholar.org/paper/298a55ddc9777e39c5bad92a750827e1cae98ac1>.
- [3] DENG, X., GU, Y., ZHENG, B., CHEN, S., STEVENS, S., WANG, B., SUN, H. *et al.* (2023) Mind2web: Towards a generalist agent for the web. *ArXiv abs/2306.06070*: null. doi:10.48550/arXiv.2306.06070, URL <https://www.semanticscholar.org/paper/58f8925a8b87054ad0635a6398a7fe24935b1604>.
- [4] SHAW, P., JOSHI, M., COHAN, J., BERANT, J., PASUPAT, P., HU, H., KHANDELWAL, U. *et al.* (2023) From pixels to ui actions: Learning to follow instructions via graphical user interfaces. *ArXiv abs/2306.00245*: null. doi:10.48550/arXiv.2306.00245, URL <https://www.semanticscholar.org/paper/ee7020fc413590878dca60dcf41896bbe6a6c628>.
- [5] LIU, E., GUU, K., PASUPAT, P., SHI, T. and LIANG, P. (2018) Reinforcement learning on web interfaces using workflow-guided exploration. *ArXiv abs/1802.08802*: null. URL <https://www.semanticscholar.org/paper/7dd8fc595afdd2097b43a5af9a8d9f5e97a65ec1>.
- [6] FURUTA, H., NACHUM, O., LEE, K.H., MATSUO, Y. and GU, S. (2023) Instruction-finetuned foundation models for multimodal web navigation URL <https://www.semanticscholar.org/paper/4762b614958013039ed14c59367eca05165b0c21>.
- [7] GUR, I., NACHUM, O., MIAO, Y., SAJDARI, M., HUANG, A., CHOWDERY, A., NARANG, S. *et al.* (2022) Understanding html with large language models. *ArXiv abs/2210.03945*: null. doi:10.48550/arXiv.2210.03945, URL <https://www.semanticscholar.org/paper/3ee3f425482cf86989d809155cc8cf2bf8d8113e>.
- [8] HONG, W., WANG, W., LV, Q., XU, J., YU, W., JI, J., WANG, Y. *et al.* (2024) Cogagent: A visual language model for gui agents. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* : 14281–14290. doi:10.1109/CVPR52733.2024.01354, URL <https://www.semanticscholar.org/paper/e50583008fe4ec049e42fdc01727ce98f6d86a35>.
- [9] SHRIDHAR, M., YUAN, X., CÔTÉ, M.A., BISK, Y., TRISCHLER, A. and HAUSKNECHT, M.J. (2020) Alfworld: Aligning text and embodied environments for interactive learning. *ArXiv abs/2010.03768*: null. URL <https://www.semanticscholar.org/paper/398a0625e8707a0b41ac58eaec51e8feb87dd7cb>.
- [10] YOU, K., ZHANG, H., SCHOOP, E., WEERS, F., SWEARNGIN, A., NICHOLS, J., YANG, Y. *et al.* (2024) Ferret-ui: Grounded mobile ui understanding with multimodal llms doi:10.48550/arXiv.2404.05719, URL <https://www.semanticscholar.org/paper/78fbb6e7a1c568a04e8c935aa9909d0c942ea5f6>.
- [11] XU, Y., SU, H., XING, C., MI, B., LIU, Q., SHI, W., HUI, B. *et al.* (2023) Lemur: Harmonizing natural language and code for language agents. *ArXiv abs/2310.06830*: null. doi:10.48550/arXiv.2310.06830, URL <https://www.semanticscholar.org/paper/7dd8fc595afdd2097b43a5af9a8d9f5e97a65ec1>.
- [12] ZHENG, L., WANG, R. and AN, B. (2023) Synapse: Trajectory-as-exemplar prompting with memory for computer control URL <https://www.semanticscholar.org/paper/0bfc804e31eecd77f45e4ee7f4d629ffdc628>.
- [13] YAO, S., ZHAO, J., YU, D., DU, N., SHAFRAN, I., NARASIMHAN, K. and CAO, Y. (2022) React: Synergizing reasoning and acting in language models. *ArXiv abs/2210.03629*: null. URL <https://www.semanticscholar.org/paper/0671fd553dd670a4e820553a974bc48040ba0819>.
- [14] RAWLES, C., LI, A., RODRIGUEZ, D., RIVA, O. and LILLICRAP, T. (2023) Android in the wild: A large-scale dataset for android device control. *ArXiv abs/2307.10088*: null. doi:10.48550/arXiv.2307.10088, URL <https://www.semanticscholar.org/paper/060e0df71620f1844839f6a993dfa5fb8e4c3bf6>.
- [15] XU, B., PENG, Z., LEI, B., MUKHERJEE, S., LIU, Y. and XU, D. (2023) Rewoo: Decoupling reasoning from observations for efficient augmented language models URL <https://www.semanticscholar.org/paper/90027ca7802645671a69b00b65e1fa94e6b63544>.
- [16] OU, T., XU, F.F., MADAAN, A., LIU, J., LO, R., SRIDHAR, A., SENGUPTA, S. *et al.* (2024) Synatra: Turning indirect knowledge into direct demonstrations for digital agents at scale URL <https://www.semanticscholar.org/paper/3b10d18fbca2b59e7beb3038c658926d2e384879>.
- [17] KOH, J.Y., LO, R., JANG, L., DUVVUR, V., LIM, M.C., HUANG, P.Y., NEUBIG, G. *et al.* (2024) Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. *ArXiv abs/2401.13649*: null. doi:10.48550/arXiv.2401.13649, URL <https://www.semanticscholar.org/paper/f554b22d2ccf786a6d61d5858f43024ba9115e15>.
- [18] CHENG, K., SUN, Q., CHU, Y., XU, F., LI, Y., ZHANG, J. and WU, Z. (2024) Seeclik: Harnessing gui grounding for advanced visual gui agents doi:10.48550/arXiv.2401.10935, URL <https://www.semanticscholar.org/paper/f9b39a6a7e40986b46f7796f3a805d70d7e3931a>.
- [19] ZHENG, B., GOU, B., KIL, J., SUN, H. and SU, Y. (2024) Gpt-4v(ision) is a generalist web agent, if grounded. *ArXiv abs/2401.01614*: null. doi:10.48550/arXiv.2401.01614, URL <https://www.semanticscholar.org/paper/c844694387a89a477e7a8bbf918171cdc3b85672>.
- [20] WU, Z., HAN, C., DING, Z., WENG, Z., LIU, Z., YAO, S., YU, T. *et al.* (2024) Os-copilot: Towards generalist computer agents with self-improvement. *ArXiv*

- abs/2402.07456:** null. doi:10.48550/arXiv.2402.07456, URL <https://www.semanticscholar.org/paper/7ca30ad71a113ab12a0089824d8bf9d0b4e623ae>.
- [21] LIU, X., YU, H., ZHANG, H., XU, Y., LEI, X., LAI, H., GU, Y. *et al.* (2023) Agentbench: Evaluating llms as agents. *ArXiv abs/2308.03688:* null. doi:10.48550/arXiv.2308.03688, URL <https://www.semanticscholar.org/paper/aed8f01122d1a89c43900e995c80bfda7936568e>.
- [22] GUR, I., FURUTA, H., HUANG, A., SAFDARI, M., MATSUO, Y., ECK, D. and FAUST, A. (2023) A real-world webagent with planning, long context understanding, and program synthesis. *ArXiv abs/2307.12856:* null. doi:10.48550/arXiv.2307.12856, URL <https://www.semanticscholar.org/paper/c648fad0b5dea9ccdec956cce1294fd574d35eb2>.
- [23] KIM, G., BALDI, P. and McALEER, S. (2023) Multimodal web navigation with instruction-finetuned foundation models. *ArXiv abs/2303.17491:* null. doi:10.48550/arXiv.2303.17491, URL <https://www.semanticscholar.org/paper/9a75e23639bfcc3a51da57a3b682a984d1d8ac0b>.
- [24] BALDI, P., McALEER, S. and KIM, G. (2023) Language models can solve computer tasks. *ArXiv abs/2303.17491:* null. doi:10.48550/arXiv.2303.17491, URL <https://www.semanticscholar.org/paper/8147cec9245d34d13732a08e915c920a1a499bb5>.
- [25] MUGHAL, A.H., CHADALAVADA, P., MUNIR, A., DUTTA, A. and QURESHI, M.A. (2022) Design of deep neural networks for transfer time prediction of spacecraft electric orbit-raising. *Intelligent Systems with Applications* 15: 200092. doi:<https://doi.org/10.1016/j.iswa.2022.200092>, URL <https://www.sciencedirect.com/science/article/pii/S2667305322000321>.
- [26] MUGHAL, A.H. (2024), Advancing bdd software testing: Dynamic scenario re-usability and step auto-complete for cucumber framework. URL <https://arxiv.org/abs/2402.15928>. 2402.15928.
- [27] PAN, J., ZHANG, Y., TOMLIN, N., ZHOU, Y., LEVINE, S. and SUHR, A. (2024) Autonomous evaluation and refinement of digital agents. *ArXiv abs/2404.06474:* null. doi:10.48550/arXiv.2404.06474, URL <https://www.semanticscholar.org/paper/aed8f01122d1a89c43900e995c80bfda7936568e>.
- [28] WANG, X., CHEN, Y., YUAN, L., ZHANG, Y., LI, Y., PENG, H. and JI, H. (2024) Executable code actions elicit better llm agents. *ArXiv abs/2402.01030:* null. doi:10.48550/arXiv.2402.01030, URL <https://www.semanticscholar.org/paper/fa96417f8766568ba570088513940bbf14e3b356>.