GPT-4V(ision) is a Generalist Web Agent, if Grounded

Boyuan Zheng [1] **Boyu Gou** [1] **Jihyung Kil** [1] **Huan Sun** [1] **Yu Su** [1]
https://osu-nlp-group.github.io/SeeAct

Abstract

The recent development on large multimodal models (LMMs), especially GPT-4V(ision) and Gemini, has been quickly expanding the capability boundaries of multimodal models beyond traditional tasks like image captioning and visual question answering. In this work, we explore the potential of LMMs like GPT-4V as a generalist web agent that can follow natural language instructions to complete tasks on any given website. We propose SeeAct, a generalist web agent that harnesses the power of LMMs for integrated visual understanding and acting on the web. We evaluate on the recent Mind2Web benchmark. In addition to standard offline evaluation on cached websites, we enable a new online evaluation setting by developing a tool that allows running web agents on live websites. We show that GPT-4V presents a great potential for web agents—it can successfully complete 51.1% of the tasks on live websites if we manually ground its textual plans into actions on the websites. This substantially outperforms text-only LLMs like GPT-4 or smaller models (FLAN-T5 and BLIP-2) specifically fine-tuned for web agents. However, grounding still remains a major challenge. Existing LMM grounding strategies like set-of-mark prompting turns out to be not effective for web agents, and the best grounding strategy we develop in this paper leverages both the HTML structure and visuals. Yet, there is still a substantial gap with oracle grounding, leaving ample room for further improvement. All code, data, and evaluation tools are available at `https://github.com/OSU-NLP-Group/SeeAct`.

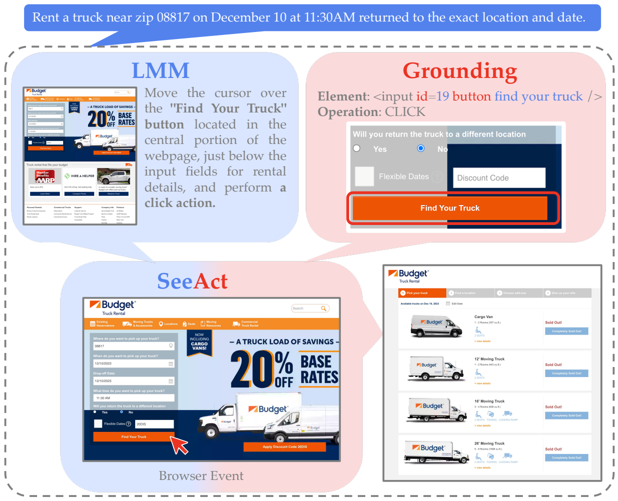

Rent a truck near zip 08817 on December 10 at 11:30AM returned to the exact location and date.

Move the cursor over the "Find Your Truck" button located in the central portion of the webpage, just below the input fields for rental details, and perform a click action.

Element: <input id=19 button find your truck />
Operation: CLICK

Figure 1: SeeAct leverages an LMM like GPT-4V to visually perceive websites and generate plans in textual forms. The textual plans are then grounded onto the HTML elements and operations to act on the website.

1. Introduction

Large multimodal models (LMMs; Li et al. (2023); Alayrac et al. (2022); Liu et al. (2023b)), especially recent ones such as GPT-4V(ision) (OpenAI, 2023) and Gemini (Anil et al., 2023), have shown a remarkable capability on standard vision-and-language understanding and reasoning benchmarks (Kazemzadeh et al., 2014; Goyal et al., 2016; Hendrycks et al., 2020; Saikh et al., 2022; Lu et al., 2022; Zhong et al., 2023; Yue et al., 2023). While web content has been a primary source of training data, a largely overlooked part of the web is the websites themselves—every website is designed to be rendered visually for easy consumption by human users. This poses a new challenge and a new opportunity for LMMs. On the one hand, screenshots of rendered websites, which could contain thousands of elements with rich relations, are more complex than the images in most existing benchmarks, which are usually object- or scene-centric. On the other hand, if LMMs can accurately comprehend websites, it will open the door for numerous applications on the web.

In this work, we aim to investigate the potential of LMMs

---

[1] The Ohio State University, Columbus, OH. Correspondence to: Boyuan Zheng <zheng.2372@osu.edu>, Huan Sun <sun.397@osu.edu>, Yu Su <su.809@osu.edu>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

as generalist web agents (Deng et al., 2023). A generalist web agent, as defined in MIND2WEB (Deng et al., 2023), is expected to follow natural language instructions and complete tasks on any given real-world website (e.g., Figure 1). The tasks can be fairly diverse and complex, with one task possibly taking 10+ actions across multiple dynamically rendered webpages. Existing work (Deng et al., 2023; Liu et al., 2023d) primarily uses large language models (LLMs) such as GPT-4 (OpenAI, 2023) on the raw HTML input. However, HTML code is noisier than the rendered visuals and has a lower information density. For example, the screenshot in Figure 1 contains 423 HTML elements that would require 186,490 textual tokens with the GPT-2 Tokenizer, while requiring only 1,445 visual tokens using GPT-4V's visual tokenizer. Furthermore, HTML alone provides incomplete information and misses critical semantics from, e.g., embedded images.

To this end, we propose SEEACT, a generalist web agent that harnesses the power of LMMs for integrated visual understanding and acting on the web. We will focus on GPT-4V, the most advanced LMM publicly available to date, and compare it with smaller LMMs such as BLIP-2 (Li et al., 2023) and LLaVA-1.5 (Liu et al., 2023a;c). We find that GPT-4V exhibits a strong capability in visually understanding rendered webpages and generate the right plans in textual forms across a wide range of websites and tasks. However, *grounding* (Chandu et al., 2021; Gu et al., 2023), *i.e.*, converting the textual plan into precise actions on the website, remains a major challenge. It involves selecting the right HTML element to interact with as well as the right operation (e.g., `Click`, `Type`, or `Select`). We propose multiple grounding methods, including superpositioning bounding boxes and index labels onto the image, similar to set-of-mark prompting (Yang et al., 2023a) that has been shown effective on object- or scene-centric images. However, we find that on complex images with rich semantic and spatial relationships like webpage screenshots, severe hallucination is observed from GPT-4V. The most effective grounding strategy leverages the known correspondence between HTML element and their visual rendering, a unique property for websites compared to natural images.

We evaluate SEEACT on the MIND2WEB dataset (Deng et al., 2023) and compare it with text-only large language models (LLMs) like GPT-4 (OpenAI, 2023) as well as smaller models (FLAN-T5 (Chung et al., 2022) and BLIP-2 (Li et al., 2023) for supervised fine-tuning, LLaVA-1.5 (Liu et al., 2023a;c) and CogAgent (Hong et al., 2023) for in-context learning). In addition to the standard offline evaluation setting on cached websites, we further establish a new *online evaluation* setting by developing a tool that allows for running web agents on live websites. The major findings from our exploration are summarized below:

- SEEACT with GPT-4V is a strong generalist web agent, if oracle grounding is provided. In online evaluation, it can successfully complete 51.1% of tasks on different websites, substantially outperforming existing methods like GPT-4 (13.3%) or FLAN-T5 (8.9%). This strongly demonstrates the potential of LMMs like GPT-4V for web agents.
- However, grounding is still a major challenge. The best grounding strategy still has a 20-32% gap with oracle grounding. Among the various grounding strategies, the best one organically leverages both HTML text and visuals, substantially outperforming image annotation strategies (Yang et al., 2023a) by up to 10%.
- In-context learning with large models (both LMMs and LLMs) shows better generalization to unseen websites, while supervised fine-tuning still has an edge on websites seen during training.
- There is a non-negligible discrepancy between online and offline evaluation because there can often be multiple viable plans for completing the same task. Online evaluation is more indicative of a model's true performance.

2. SeeAct

In this section, we first explain the problem formulation of web agents and then introduce SEEACT, a generalist web agent based on LMMs. Specifically, given a web-based task (e.g., "Rent a truck with the lowest rate" in the car rental website), we examine two essential capabilities of LMMs as a generalist web agent: (i) **Action Generation** to produce an action description at each step (e.g., "Move the cursor over the 'Find Your Truck' button and perform a click") towards completing the task, and (ii) **Element Grounding** to identify an HTML element (e.g., "[button] Find Your Truck") at the current step on the webpage.

2.1. Formulation

Given a website $\mathcal{S}$ (*e.g.,* a car rental website) and a task T (*e.g., "Rent a truck with the lowest rate"*), the web agent should generate a sequence of executable actions $A = [\boldsymbol{a}_1, \boldsymbol{a}_2, ..., \boldsymbol{a}_n]$ to complete the task. Specifically, at time step t, the agent should generate an action $\boldsymbol{a}_t$ based on the current environment observation s_t, the previous actions $\{\boldsymbol{a}_1, \boldsymbol{a}_2, ..., \boldsymbol{a}_{t-1}\}$, and the task T:

$$\boldsymbol{a}_t = \pi(s_t, T, \{\boldsymbol{a}_1, \boldsymbol{a}_2, ..., \boldsymbol{a}_{t-1}\})$$

The environment observation s_t comprises an HTML document h_t and a screenshot image i_t. LLMs can only be grounded on the HTML document, while LMMs can be grounded on both the HTML document and the screenshot image. The website status is updated accordingly after each

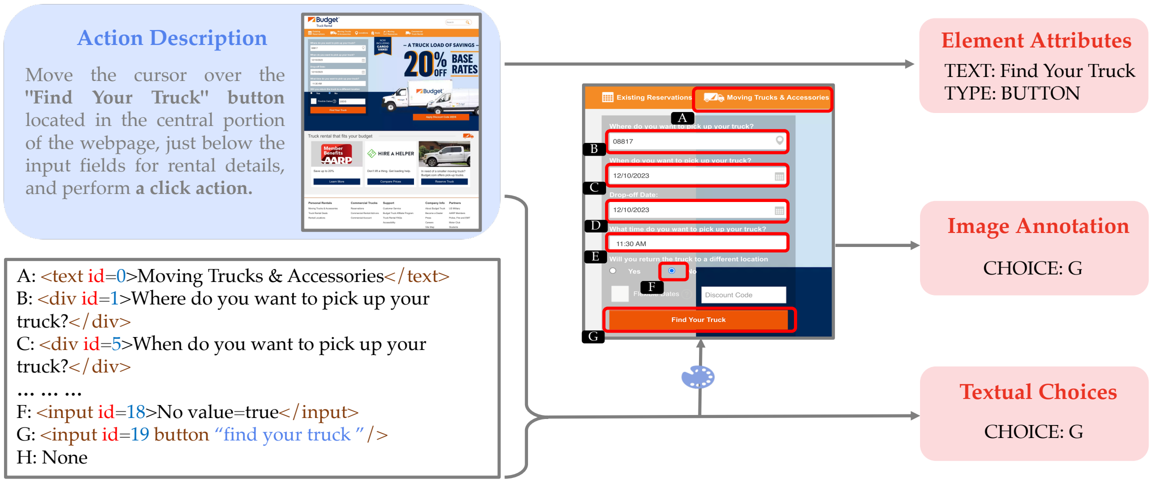

Move the cursor over the "Find Your Truck" button located in the central portion of the webpage, just below the input fields for rental details, and perform **a click action.**

TEXT: Find Your Truck
TYPE: BUTTON

CHOICE: G

CHOICE: G

```
A: <text id=0>Moving Trucks & Accessories</text>
B: <div id=1>Where do you want to pick up your
truck?</div>
C: <div id=5>When do you want to pick up your
truck?</div>
... ... ...
F: <input id=18>No value=true</input>
G: <input id=19 button “find your truck ”/>
H: None
```

Figure 2: An example of the element grounding process for a single action during completing the given task with three different methods. In this action step, the model needs to click the "Find Your Truck" button to perform a search. For grounding with textual choices, some element candidates represented with HTML text are given, the model is required to generate the choice index of the target element. For image annotation, bounding boxes and index labels are added to the image. The model is required to generate the label on the bottom-left of the target element. For grounding with element attributes, the model needs to predict the text and type of the target element.

action:

$$s_{t+1} = \mathcal{S}(\boldsymbol{a}_t) = \{h_{t+1}, i_{t+1}\}$$

For simplicity, in subsequent step-wise formulations, the time step notation t is omitted.

An action $\boldsymbol{a}$ corresponds to a browser event provided by the website environment. Therefore, we formulate an action as a triplet of three necessary variables for a browser event (e, o, v). $e \in \mathcal{E}$ identifies the target webpage element to operate on, such as the "Find Your Truck" button in Figure 2. $\mathcal{E}$ represents the set of webpage elements within the environment $\mathcal{S}$. The operation $o \in \mathcal{O}$ is the action to be performed on the target element, with $\mathcal{O}$ encompassing all possible operations in $\mathcal{S}$ (e.g., `Click`, `Type`). The variable v denotes the additional value needed for a certain operation (e.g., the date *12/10/2023* for a `Type` operation).

However, agents based on LLMs or LMMs are typically unable to directly generate the three variables (e, o, v) required for a browser event. Instead, they generate a textual description of the intended action $\tilde{\boldsymbol{a}}$, containing information about these variables as $(\tilde{e}, \tilde{o}, \tilde{v})$. This process is referred to as **Action Generation**. To interact with the environment, a further step is required to convert $\tilde{\boldsymbol{a}}$ into $\boldsymbol{a}$, which we refer to as **Action Grounding**.

2.2. Action Generation

We explicitly instruct the LMM to imitate humans browsing a webpage and analyze the task, webpage, and previous actions. It is asked to generate an action description $\tilde{a}$ based on its analysis and reasoning. We take the screenshot image i as the visual context without utilizing the HTML document h for action generation.

2.3. Action Grounding

Despite the capability of LMMs in identifying and describing the next action to complete the given task in natural language, it is still challenging to convert the action description $\tilde{\boldsymbol{a}}$ into an executable action $\boldsymbol{a}$ within the environment. Deriving operation type o and value v from the action description $\tilde{\boldsymbol{a}}$ can be solved through string parsing reasonably well. The key challenge is to identify the target element e from the generated $\tilde{e}$, which we refer to as **Element Grounding**.

To address this challenge, we explore three approaches using different types of information: Grounding via Element Attributes, Grounding via Textual Choices, and Grounding via Image Annotation, as depicted in Figure 2. The prompting details of action generation and grounding are included in Appendix D.

Grounding via Element Attributes. This approach involves prompting the model to generate as detailed attributes of the target element as possible, thereby providing more information to precisely match with the target HTML element. Specifically, we prompt the model to not only describe the element e, but also specify the target element's type and the textual content in $\tilde{e}$. For example, as illustrated in Figure 2, the model would generate element text as "Find Your Truck" and identify its type as a "BUTTON." Following this, a heuristic search is performed across the DOM elements, using the element text and type to locate matching elements.

In cases where a single match is found, it is automatically selected. Otherwise, when multiple matches arise, the model is further prompted to select the final selection.

Grounding via Textual Choices. The above approach demands precise and sufficient attribute descriptions from the LMM and accurate matching by the heuristic search, which can be highly demanding. For instance, many elements may have no textual content or have textual information in a nearby element instead of itself.

Alternatively, we provide the model with textual representations of elements as choices to facilitate grounding, which has already been proven effective in MindAct (Deng et al., 2023). Specifically, MindAct utilizes a ranking model to select top-k candidate elements $(e_1, e_2, ..., e_k)$ with a pretrained cross-encoder. Each candidate element is represented as a choice in a multi-choice question with its HTML text, as illustrated in Figure 2. After generating the action description $\tilde{\boldsymbol{a}}$, the model is further asked a multi-choice question to choose its intended element from the given multiple choices (including a ‘none’ option).

Grounding via Image Annotation. Textual representations alone are sometimes insufficient to distinguish similar or identical elements, as illustrated in Appendix H. Therefore, in this approach, we propose to overlay a bounding box for each candidate element e selected by the ranker as well as a label around the bounding box[1] with a label assignment method to avoid overlapping between markups. The model is expected to generate the label corresponding to the target element.

Oracle Action Grounding. Ideally, the action description $\tilde{a}$ must encompass all necessary details to precisely identify each variable (e, o, v) of the action triplet. To assess the performance of action generation, an oracle grounding method, which ensures the variables be identified as long as they are mentioned in the action description, is desired. Here we approximate the oracle grounding method by asking human annotators to identify the model’s intended actions.

3. Experiments

3.1. Dataset

We evaluate our methods on MIND2WEB (Deng et al., 2023), a comprehensive dataset encompassing over 2,000 complex web tasks with annotated actions. This dataset spans 137 websites across 31 low-level domains, categorized into 12 high-level domains. It supports three primary operations: `Click`, `Type`, and `Select`, with `Hover` and `Press Enter` operations integrated into `Click` to avoid ambiguity.

The dataset’s test sets aim to measure the generalization of web agents across different tasks, websites, and domains. Specifically, the Cross-Task setting focuses on evaluating agents on tasks that are new to the training data but within included domains and websites. The Cross-Website setting evaluates agents with tasks across 10 new websites for each of the top-level domains in the training data. The Cross-Domain setting assesses agent performance on tasks in two top-level domains held out from the training data.

We align each HTML document in the dataset with its corresponding webpage screenshot image from the MIND2WEB raw dump, which undergoes human verification to confirm element visibility and correct rendering for action prediction. This cleaned version of the dataset is called *Multimodal Mind2Web*, with the statistics in Table 1.[2]

3.2. Methods

SeeAct. In grounding via image annotation and textual choices, we first employ the DeBERTa-base cross-encoder from MindAct (Deng et al., 2023) to rank the top 50 elements for better comparison with its text-only counterparts. Then, we cluster elements into groups of 17 options for inference. In grounding via element attributes, no candidate element is provided. We experiment all three grounding methods with GPT-4V API, and use the best-performing grounding method for Gemini Pro Vision (Anil et al., 2023), and LLaVA-1.5 (Liu et al., 2023a;c).

MindAct. To compare with SEEACT, we also implement methods based on text-only LLMs and BLIP-2 (Li et al., 2023) following the two-stage strategy of MindAct (Deng et al., 2023). Firstly, we employ the ranker above to pick the top 50 elements. Subsequently, the action generation problem is formulated as a multi-choice question answering problem, with the candidate elements as options, including a "None" option if the target element is absent. During inference, elements are clustered into groups of 5 elements, with iterative refinement, until a single choice is made or all options are discarded. We evaluate supervised fine-tuning (SFT) methods using FLAN-T5 (Chung et al., 2022) and BLIP-2-T5 and in-context learning (ICL) methods using GPT-3.5 and GPT-4.

Pixel-Level Grounding. LMMs can generate target element coordinates in the image via training on datasets augmented with object coordinates, especially for open-sourced models (Hong et al., 2023; Cheng et al., 2024; You et al., 2023). We choose CogAgent (Hong et al., 2023) as a representative model for this experiment. Details of each method can be found in Appendix A.

---

[1] We use the Supervision library for image annotation: https://supervision.roboflow.com/

[2] The dataset is released at https://huggingface.co/datasets/osunlp/Multimodal-Mind2Web.

Table 1: Statistics of the cleaned Multimodal Mind2Web dataset. The average # of visual tokens is based on OpenAI visual token calculator.

Split	# Tasks	# Domains	# Websites	Avg # Actions	Avg # Visual Tokens	Avg # HTML Elements	Avg # HTML Tokens
Train	1,009	17	73	7.7	4,240	602	128,827
Cross-Domain	694	13	53	5.9	4,314	494	91,163
Cross-Task	177	17	64	7.6	4,172	607	123,274
Cross-Website	142	9	10	7.2	4,653	612	114,358

3.3. Offline Evaluation

We adopt the evaluation metrics utilized in MIND2WEB. **Element Accuracy** (Ele. Acc) compares the predicted element with the ground-truth elements. **Operation F1** (Op. F1) calculates the token-level F1 score for the predicted operation comprised of action and input value. **Step Success Rate** (Step SR) measures the success of each action step. A step is successful only if the selected element and the predicted operation are correct. We report macro averages across tasks for these step-wise metrics. **Success Rate** (SR) measures the success of an entire task. A task is regarded successful only if all steps have succeeded. This metric is stringent without allowing the model any space for exploration and error correction. Therefore, for offline evaluation, we focus on the first three metrics. However, we also conduct online evaluation on live websites for better evaluation on the whole task success rate, as detailed below.

3.4. Online Evaluation

We develop a new online evaluation tool using Playwright[3] to evaluate web agents on live websites (instead of cached websites in offline evaluation). Our tool can efficiently tunnel multimodal inputs from the browser to the agent and convert the predicted action (e, o, v) into a browser event for execution. To adhere to ethical standards, our experiments are restricted to non-login tasks in compliance with user agreements, and we closely monitor agent activities during online evaluation to prevent any actions that have potentially harmful impacts, like placing an order or modifying the user profile.

4. Results and Analysis

4.1. Offline Evaluation Results

GPT-4V can be a Generalist Web Agent with Oracle Action Grounding. Given an effective action grounding method, GPT-4V has the potential to serve as a generalist web agent. Specifically, as described in subsection 2.3, we provide GPT-4V with an oracle action grounding method (SEEACT$_{\text{Oracle}}$) through human annotation, the model achieves a step success rate of 61.9%, 65.0%, and 62.1% across three test splits, respectively. As shown in Table 2, this method substantially outperforms other models under all metrics across three test splits. Specifically, it achieves a 8.4% step success rate improvement over the second-best method in the Cross Task setting. The performance advantage is more pronounced under the Cross-Website and Cross-Domain settings, where it leads by 23.9% and 23.2% step success rates, demonstrating its generality compared with supervised fine-tuning. This observation is further corroborated within the online evaluation (Table 4).

Element Grounding Method Comparison. However, there is a noticeable gap between oracle grounding and all three proposed grounding methods, as shown in Table 3. This demonstrates that grounding, especially element grounding, is a major bottleneck. Element grounding via textual choice (SEEACT$_{\text{Choice}}$) demonstrates the best performance under all metrics across all settings, comparable to supervised fine-tuning and showing a substantial improvement over text-only LLMs.

Grounding via image annotation (SEEACT$_{\text{Annotation}}$) offers an intuitive approach and shows promising results in recent work that focuses on object- or scene-centric images (Yang et al., 2023a). However, we find that on complex images with rich semantic and spatial relationships like webpage screenshots, severe hallucination is observed from GPT-4V. Specifically, it often fails to correctly map its generated element description (which is often correct according to oracle grounding) to the right bounding box and index label in the image, leading to a low element accuracy. This limitation primarily arises from GPT-4V's weakness in understanding image details and relative spatial location, a topic that we will further delve into in Appendix E. We leverage bottom-left number labels around bounding boxes as it is identified as the optimal markups in the ablation study in Appendix B.

Grounding via element attributes (SEEACT$_{\text{Attribute}}$) also demonstrates inferior performance. This method's effectiveness is primarily limited by its heuristic-based element localization strategy, which depends on textual and locality characteristics. This becomes problematic as not all webpage elements contain text, and sometimes the relevant text is associated with a nearby but distinct element.

LMMs vs. LLMs. The SEEACT$_{\text{Choice}}$ with GPT-4V demon-

---

[3] `https://playwright.dev/`

Table 2: Performance of different models. All models under SEEACT utilize "Choices" for grounding. Methods with * mark are conducted on a subset with 30 tasks for each task split.

Model	Cross-Task			Cross-Website			Cross-Domain		
	Ele. Acc	Op. F1	Step SR	Ele. Acc	Op. F1	Step SR	Ele. Acc	Op. F1	Step SR
Supervised Fine-Tuning									
FLAN-T5-XL	**57.1**	75.7	**53.5**	**43.8**	67.7	**41.1**	**41.4**	65.9	**38.9**
BLIP-2-T5-XL	50.1	**77.0**	47.0	39.4	**69.3**	37.0	41.2	**69.3**	38.9
In-Context Learning									
GPT-3.5*	19.4	59.2	16.8	14.9	56.5	14.1	25.2	57.9	24.1
GPT-4*	40.8	63.1	32.3	30.2	61.0	27.0	35.4	61.9	29.7
COGAGENT	22.4	53.0	17.6	18.4	42.2	13.4	20.6	42.0	15.5
SEEACT									
– LLAVA-1.5	9.7	65.6	8.1	9.1	60.8	7.5	10.9	63.9	8.5
– GEMINI PRO VISION	21.5	67.7	19.6	17.1	61.3	15.4	20.7	64.3	18.0
– GPT-4V	**46.4**	**73.4**	**40.2**	**38.0**	**67.8**	**32.4**	**42.4**	**69.3**	**36.8**
– GPT-4V-Oracle*	66.4	79.2	61.9	69.5	78.9	65.0	72.8	73.6	62.1

Table 3: Step success rate (%) of GPT-4V on a subset of 30 tasks for each task split with different grounding methods. "Attributes", "Choices", "Annotation", and "Oracle" refer to element grounding via Element Attributes, Textual Choices, Image Annotation, and Human Annotation, respectively, as described in subsection 2.3.

Grounding	Cross-Task	Cross-Website	Cross-Domain
Attribute	16.1	12.1	19.0
Annotation	20.3	13.9	23.7
Choice	39.1	32.7	42.0
Oracle	61.9	65.0	62.1

strates a substantial performance advantage over the text-only GPT-4 under all three metrics across all three test splits. Specifically, it outperforms GPT-4 in step success rate of 7.9%, 5.4%, and 7.1% on three settings, respectively. Interestingly, fine-tuned BLIP-2-T5 does not show a noticeable gain over FLAN-T5, despite having additional visual input. Several factors may contribute to this. First, the CLIP model used as the image encoder may not be sufficiently adept at image details, as explored by Shen et al. (2021). This limitation is particularly relevant for our web navigation task, which demands a high level of image detail comprehension. Second, BLIP-2-T5 utilizes an off-the-shelf CLIP model that may not be optimal for webpage screenshots. Finally, although the screenshots in the test splits are error-free, some of the examples in the training set might contain issues such as rendering failures or inaccuracies when annotators capture the screenshot.

SFT vs. ICL. We compare SFT and ICL methods to offer insights for developing web agents in different scenarios. ICL (with SEEACT) demonstrates consistent and robust performance across three test splits. ICL is particularly advantageous in scenarios lacking annotations or requiring strong generalization capabilities for new domains and websites. As grounding methods improve towards oracle grounding, ICL is poised to show even stronger performance. On the other hand, SFT methods show better generalization across tasks on websites already seen during training. Considering the high cost of data annotation for web agents and the billions of websites on the Internet, ICL offers a more compelling solution for generalist web agents. However, if one only needs to develop a strong web agent for a certain website, SFT is still a competitive solution.

Table 4: Whole task success rate (%) under both offline and online evaluation. Offline_0 and Offline_1 refer to no tolerance for error at any step and allowing for error at one step, respectively.

	Offline_0	Offline_1	Online
FLAN-T5-XL	4.4	24.4	8.9
GPT-4	1.1	12.2	13.3
$\text{SEEACT}_{\text{Choice}}$	3.3	12.2	37.8
$\text{SEEACT}_{\text{Oracle}}$	13.3	27.8	51.1

4.2. Online Evaluation Results

In online evaluation, we pair a web agent with a human annotator, where the human was tasked to monitor agent actions that may change real-world states and determine whether each task was successfully completed. For comparative analysis, we include success rates from offline evaluation, denoted as Offline_0 (allowing zero wrong action) and Offline_1 (allowing one wrong action). For a fair comparison between offline and online evaluations, we only re-write time-sensitive tasks to ensure they are still valid when the evaluation is conducted. For instance, we update the dates for flight-related tasks. Finally, we conduct the online evalu-

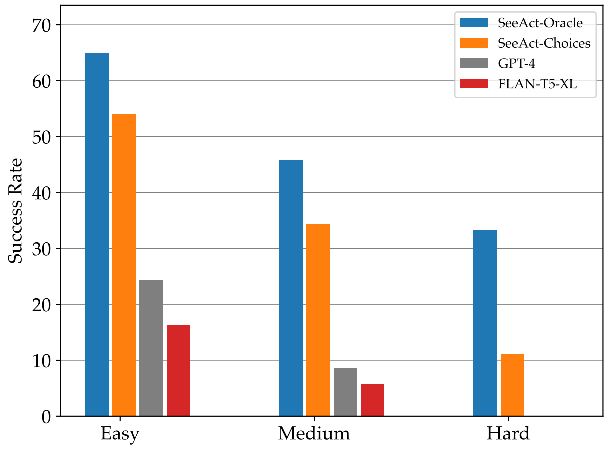

Figure 3: Whole task success rate across task difficulty levels. We categorize tasks based on the number of actions to complete, i.e., Easy: 1-4, Medium: 5-9, and Hard: 10-18, with 37, 35, and 18 tasks in each group, respectively.

ation on the same subset of total 90 tasks from the three test splits and report the mean performance across these tasks.

Table 4 shows that the whole task success rate in online evaluation substantially exceeds that of offline evaluation (Offline$_0$). This finding suggests that the whole task success rate is likely underestimated in the offline evaluation due to the variability in actions and plans. In other words, there may be multiple viable plans for a task, but the reference plan in offline evaluation only captures one of them.

Across all three settings, SEEACT$_{\text{Choice}}$ outperforms both GPT-4 and FLAN-T5-XL by a large margin of over 20% whole task success rate. Using oracle grounding further improves the performance substantially, reaching a remarkable whole task success rate of 51.1%. Although GPT-4 shows much worse performance than FLAN-T5-XL in step success rate under offline evaluation (Table 2), it outperforms FLAN-T5-XL by 4.4% whole task success rate in the online evaluation. These results further confirm the potential of large models for generalist web agents compared with fine-tuned small models.

4.3. Analysis

Online Success Rate by Task Difficulty. We investigate the performance of web agents on tasks across different difficulty levels. We estimate the task difficulty based on the number of actions taken by annotators during action trace annotation. As shown in Figure 3, the whole task success rate is negatively correlated with the number of actions—it decreases as the number of actions increases across all four methods. SEEACT$_{\text{Oracle}}$ consistently outperforms other methods across all difficulty levels. Interestingly, the gap between SEEACT$_{\text{Oracle}}$ and SEEACT$_{\text{Choice}}$ enlarges on longer-horizon tasks. This is understandable because grounding errors cascade to later steps; nonetheless, it further shows the challenge of grounding for GPT-4V and the need for better grounding methods.

Error Analysis in Grounding via Image Annotation. Set-of-mark prompting (Yang et al., 2023a) uses a similar method as grounding via image annotation and has been shown effective on object- or scene-centric images (Lin et al., 2014; Plummer et al., 2015; Zhou et al., 2017). However, this grounding method is suboptimal on webpage screenshot images that are complex and contain rich semantic and spatial relationships. To analyze the reasons behind the failures, we randomly sample 100 action predictions with correct action generation but wrong grounding results. We observes major types of errors as : (1) Making up bounding box & label; (2) Failure to link bounding boxes with the correct labels. Illustrative examples are included in Appendix E.

Our analysis reveals that 54% of the errors can be attributed to GPT-4V's tendency of visual illusion (Guan et al., 2023), where the model misinterprets and fabricates content over the image. Specifically, the target element described in action generation does not have a bounding box or a label on the bottom-left, where the model is supposed to generate "NA". However, the model falsely assumes the presence of a bounding box and makes up a label as the answer. Another 46% of errors are caused by GPT-4V's limitation in recognizing the relative position within an image. Specifically, the model is capable of identifying the target element within the bounding box. However, it struggles to correctly link the bounding box with its corresponding label.

4.4. Case Study

GPT-4V exhibits promising capabilities, ranging from speculative planning, webpage content reasoning, and error correction to surpassing the limitations of superficial textual similarity matching inherent in fine-tuned, text-only models.

World Knowledge. GPT-4V demonstrates substantial advantages in tasks requiring certain knowledge over fine-tuned models at a smaller scale. As shown in Appendix I, GPT-4V is able to identify the IATA code of the airport in Los Cabos as SJD. In contrast, smaller models are typically weaker at knowledge-intensive tasks and are also likely to lose knowledge during the fine-tuning process due to catastrophic forgetting.

World Model (for Websites). GPT-4V exhibits the potential of a "world model" for websites. As shown in Appendix G, GPT-4V can predict the state transitions on a website (e.g., what would happen if I clicked this button). Based on its awareness of website state transitions, GPT-4V can conduct speculative planning involving a sequence of subsequent actions in the future to complete the given task.

Error Correction Awareness. GPT-4V also exhibits the awareness of error correction in the previous actions. In the example in Appendix J, it realizes that the mobile phone number is invalid due to the wrong format and generates the description of the action to correct this error. This highlights the model's potential for adaptation in online settings, where actions may not always follow pre-defined, ideal paths as in offline evaluations. This capability paves the way for adding robustness and reasonable dynamic planning.

5. Related Work

Web Agent. Many works have focused on improving web agents relying on the HTML document (Deng et al., 2023; Gur et al., 2023; 2022; 2023; Kim et al., 2023; Sridhar et al., 2023). However, a raw HTML document is often massive making it infeasible or cost-prohibitively to feed into LLMs directly. MindAct (Deng et al., 2023) instead employs a small language model to rank each HTML element and selectively consider top elements as the context. WebAgent (Gur et al., 2023) proposes an enhanced planning strategy by summarizing the HTML documents and decomposing the instruction into multiple sub-instructions. Another stream considers visual information for web agents (Shaw et al., 2023; Furuta et al., 2023; Hong et al., 2023). Pix2Act (Shaw et al., 2023) leverages Pix2Struct (Lee et al., 2022) to parse screenshot images into simplified HTML to complete GUI-based tasks (Shaw et al., 2023; Liu et al., 2018; Shi et al., 2017; Mazumder & Riva, 2020; Yao et al., 2022). WebGUM (Furuta et al., 2023) and CogAgent (Hong et al., 2023) pre-train an LMM with massive screenshot-HTML data to enhance its decision-making on real-world web navigation like Mind2Web. While all these prior works show promise, generalizing to various web environments remains a challenge for existing models. Thus, SEEACT explores recently released, more powerful LMMs such as GPT-4V and Gemini, to demonstrate their potential as generalist web agents with comprehensive online and offline evaluation and analysis. In a concurrent work (Yan et al., 2023), GPT-4V exhibits strong performance on mobile UI understanding, which is less complex than the desktop websites we study.

Large Multimodal Models. GPT-4V (OpenAI, 2023) and Gemini (Anil et al., 2023) represent significant progress in LMMs. Several studies (Akter et al., 2023; OpenAI, 2023; Yang et al., 2023c; Zhang et al., 2023; Yang et al., 2023a; Yan et al., 2023) have highlighted their remarkable multimodal capabilities, emphasizing the advanced and versatile integration of visual and language reasoning abilities. Their performance on a series of benchmarks (Kazemzadeh et al., 2014; Goyal et al., 2016; Hendrycks et al., 2020; Saikh et al., 2022; Lu et al., 2022; Zhong et al., 2023; Yue et al., 2023) also showcases remarkable capabilities on vision-and-language understanding and reasoning. Although open-sourced models still exhibit a performance gap with GPT-4V, they have the advantages of controllability and ease of fine-tuning for various applications. For example, in CogAgent (Hong et al., 2023), LMMs are fine-tuned on HTML and screenshot image pairs to enhance webpage understanding ability and further enhanced with an image encoder for high-resolution image details. Ferret (You et al., 2023) is finetuned to allow visual referring and grounding.

Visual Grounding. Despite LMMs having achieved remarkable vision-language understanding capabilities, they still face challenges in fine-grained visual grounding. Various visual prompting (Shtedritski et al., 2023; Yang et al., 2023b;c; Yan et al., 2023) methods have been proposed to augment GPT-4V's image detail grounding ability by overlaying visual marks onto the image. SoM (Yang et al., 2023a) involves segmenting the image into semantically meaningful regions and overlaying an array of visual marks like numbers, alphabets, masks, or bounding boxes. Fine-tuning vision-language models with image-annotated data is effective. Kosmos-2 (Peng et al., 2023) represents bounding box locations through textual location tokens. BuboGPT (Zhao et al., 2023) extract entities and find corresponding masks for objects in the image. Shikra (Chen et al., 2023) handles image detail referring and grounding by applying spatial coordinates as text tokens in inputs and outputs, respectively. Ferret (You et al., 2023) represents regions with both discrete coordinates and continuous features along with a spatial-aware visual sampler to handle diverse spatial characteristics across various shapes.

6. Conclusion

In this work, we developed SEEACT, a generalist web agent that harnesses the power of large multimodal models (LMMs) like GPT-4V to integrate visual understanding and acting on the web. We showed that LMMs present a great promise for generalist web agents, with a success rate of 50% on live websites given an oracle grounding method. GPT-4V also exhibits impressive capabilities, such as error correction and speculative planning. However, fine-grained visual grounding is still a major challenge. The most effective grounding strategies we explored in this paper still exhibit a 20-25% performance gap compared to oracle grounding. Future work should better leverage the unique properties of the Web, *e.g.*, the known correspondence between HTML and visual elements, for improving grounding and reducing hallucinations from LMMs. Furthermore, we show a significant discrepancy between online and offline evaluations, emphasizing the importance of online evaluation for an accurate assessment of a model's capabilities.

This discrepancy is largely due to the variability in potential plans for completing the same task, pointing to the dynamic nature of web interactions.

Impact Statement

Generalist web agents hold the potential to automate routine web tasks, enhance user experiences, and promote web accessibility, safety concerns related to their real-world deployment are also critical. These concerns span privacy issues, such as access to users' personal profiles, and sensitive operations, such as financial transactions or application form submissions. During the online evaluation, we noticed the possibility for these web agents to generate harmful actions on the web, and we manually validated the safety of all the actions before execution. It is critical for further research to thoroughly assess and mitigate the safety risks associated with web agents, ensuring they are safeguarded against producing and executing harmful actions. The code will also be released solely for research purposes, with the goal of making the web more accessible via language technologies under an OPEN-RAIL License. We are strongly against any potentially harmful use of the data or technology by any party.

Acknowledgments

The authors would like to thank colleagues from the OSU NLP group for their thoughtful comments. This research was supported in part by ARL W911NF2220144 and Cisco.

References

Akter, S. N., Yu, Z., Muhamed, A., Ou, T., Bauerle, A., Cabrera, Á. A., Dholakia, K., Xiong, C., and Neubig, G. An in-depth look at gemini's language abilities. 2023. URL `https://api.semanticscholar.org/CorpusID:266359502`.

Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., Ring, R., Rutherford, E., Cabi, S., Han, T., Gong, Z., Samangooei, S., Monteiro, M., Menick, J., Borgeaud, S., Brock, A., Nematzadeh, A., Sharifzadeh, S., Binkowski, M., Barreira, R., Vinyals, O., Zisserman, A., and Simonyan, K. Flamingo: a visual language model for few-shot learning. *ArXiv*, abs/2204.14198, 2022. URL `https://api.semanticscholar.org/CorpusID:248476411`.

Anil, G. T. G. R., Borgeaud, S., and et al., Y. W. Gemini: A family of highly capable multimodal models. 2023. URL `https://api.semanticscholar.org/CorpusID:266361876`.

Chandu, K. R., Bisk, Y., and Black, A. W. Grounding 'grounding' in nlp. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 4283–4305, 2021.

Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., and Zhao, R. Shikra: Unleashing multimodal llm's referential dialogue magic. *ArXiv*, abs/2306.15195, 2023. URL `https://api.semanticscholar.org/CorpusID:259262082`.

Cheng, K., Sun, Q., Chu, Y., Xu, F., Li, Y., Zhang, J., and Wu, Z. Seeclick: Harnessing gui grounding for advanced visual gui agents. 2024. URL `https://api.semanticscholar.org/CorpusID:267069082`.

Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, E., Wang, X., Dehghani, M., Brahma, S., Webson, A., Gu, S. S., Dai, Z., Suzgun, M., Chen, X., Chowdhery, A., Valter, D., Narang, S., Mishra, G., Yu, A. W., Zhao, V., Huang, Y., Dai, A. M., Yu, H., Petrov, S., hsin Chi, E. H., Dean, J., Devlin, J., Roberts, A., Zhou, D., Le, Q. V., and Wei, J. Scaling instruction-finetuned language models. *ArXiv*, abs/2210.11416, 2022. URL `https://api.semanticscholar.org/CorpusID:253018554`.

Deng, X., Gu, Y., Zheng, B., Chen, S., Stevens, S., Wang, B., Sun, H., and Su, Y. Mind2web: Towards a generalist agent for the web. *arXiv preprint arXiv:2306.06070*, 2023.

Furuta, H., Nachum, O., Lee, K.-H., Matsuo, Y., Gu, S. S., and Gur, I. Multimodal web navigation with instruction-finetuned foundation models. *ArXiv*, abs/2305.11854, 2023. URL `https://api.semanticscholar.org/CorpusID:258823350`.

Goyal, Y., Khot, T., Summers-Stay, D., Batra, D., and Parikh, D. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. *International Journal of Computer Vision*, 127:398 – 414, 2016. URL `https://api.semanticscholar.org/CorpusID:8081284`.

Gu, Y., Deng, X., and Su, Y. Don't generate, discriminate: A proposal for grounding language models to real-world environments. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4928–4949, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.270. URL `https://aclanthology.org/2023.acl-long.270`.

Guan, T., Liu, F., Wu, X., Xian, R., Li, Z., Liu, X., Wang, X., Chen, L., Huang, F., Yacoob, Y., Manocha, D., and Zhou, T. Hallusionbench: An advanced diagnostic suite for entangled language hallucination&visual illusion in large vision-language models. 2023. URL `https://api.semanticscholar.org/CorpusID:265499116`.

Gur, I., Nachum, O., Miao, Y., Safdari, M., Huang, A., Chowdhery, A., Narang, S., Fiedel, N., and Faust, A. Understanding html with large language models. In *Conference on Empirical Methods in Natural Language Processing*, 2022. URL `https://api.semanticscholar.org/CorpusID:252780086`.

Gur, I., Furuta, H., Huang, A., Safdari, M., Matsuo, Y., Eck, D., and Faust, A. A real-world webagent with planning, long context understanding, and program synthesis. *ArXiv*, abs/2307.12856, 2023. URL `https://api.semanticscholar.org/CorpusID:260126067`.

Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D. X., and Steinhardt, J. Measuring massive multitask language understanding. *ArXiv*, abs/2009.03300, 2020. URL `https://api.semanticscholar.org/CorpusID:221516475`.

Hong, W., Wang, W., Lv, Q., Xu, J., Yu, W., Ji, J., Wang, Y., Wang, Z., Dong, Y., Ding, M., and Tang, J. Cogagent: A visual language model for gui agents. 2023. URL `https://api.semanticscholar.org/CorpusID:266210390`.

Kazemzadeh, S., Ordonez, V., andre Matten, M., and Berg, T. L. Referitgame: Referring to objects in photographs of natural scenes. In *Conference on Empirical Methods in Natural Language Processing*, 2014. URL `https://api.semanticscholar.org/CorpusID:6308361`.

Kim, G., Baldi, P., and McAleer, S. M. Language models can solve computer tasks. *ArXiv*, abs/2303.17491, 2023. URL `https://api.semanticscholar.org/CorpusID:257834038`.

Koh, J. Y., Lo, R., Jang, L., Duvvur, V., Lim, M. C., Huang, P.-Y., Neubig, G., Zhou, S., Salakhutdinov, R., and Fried, D. Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. 2024. URL `https://api.semanticscholar.org/CorpusID:267199749`.

Lee, K., Joshi, M., Turc, I., Hu, H., Liu, F., Eisenschlos, J. M., Khandelwal, U., Shaw, P., Chang, M.-W., and Toutanova, K. Pix2struct: Screenshot parsing as pretraining for visual language understanding. *ArXiv*, abs/2210.03347, 2022. URL `https://api.semanticscholar.org/CorpusID:252762394`.

Li, J., Li, D., Savarese, S., and Hoi, S. C. H. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *ArXiv*, abs/2301.12597, 2023. URL `https://api.semanticscholar.org/CorpusID:256390509`.

Lin, T.-Y., Maire, M., Belongie, S. J., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, 2014. URL `https://api.semanticscholar.org/CorpusID:14113767`.

Liu, E. Z., Guu, K., Pasupat, P., Shi, T., and Liang, P. Reinforcement learning on web interfaces using workflow-guided exploration. In *International Conference on Learning Representations (ICLR)*, 2018. URL `https://arxiv.org/abs/1802.08802`.

Liu, H., Li, C., Li, Y., and Lee, Y. J. Improved baselines with visual instruction tuning, 2023a.

Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning. *ArXiv*, abs/2304.08485, 2023b. URL `https://api.semanticscholar.org/CorpusID:258179774`.

Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning, 2023c.

Liu, X., Yu, H., Zhang, H., Xu, Y., Lei, X., Lai, H., Gu, Y., Gu, Y., Ding, H., Men, K., Yang, K., Zhang, S., Deng, X., Zeng, A., Du, Z., Zhang, C., Shen, S., Zhang, T., Su, Y., Sun, H., Huang, M., Dong, Y., and Tang, J. Agentbench: Evaluating llms as agents. *ArXiv*, abs/2308.03688, 2023d. URL `https://api.semanticscholar.org/CorpusID:260682249`.

Lu, P., Mishra, S., Xia, T., Qiu, L., Chang, K.-W., Zhu, S.-C., Tafjord, O., Clark, P., and Kalyan, A. Learn to explain: Multimodal reasoning via thought chains for science question answering. *ArXiv*, abs/2209.09513, 2022. URL `https://api.semanticscholar.org/CorpusID:252383606`.

Mazumder, S. and Riva, O. Flin: A flexible natural language interface for web navigation. *ArXiv*, abs/2010.12844, 2020. URL `https://api.semanticscholar.org/CorpusID:225067907`.

OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023. URL `https://api.semanticscholar.org/CorpusID:257532815`.

Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., and Wei, F. Kosmos-2: Grounding multimodal large language models to the world. *ArXiv*, abs/2306.14824, 2023. URL `https://api.semanticscholar.org/CorpusID:259262263`.

Plummer, B. A., Wang, L., Cervantes, C. M., Caicedo, J. C., Hockenmaier, J., and Lazebnik, S. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *International Journal of Computer Vision*, 123:74 – 93, 2015. URL `https://api.semanticscholar.org/CorpusID:6941275`.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.

Saikh, T., Ghosal, T., Mittal, A., Ekbal, A., and Bhattacharyya, P. Scienceqa: a novel resource for question answering on scholarly articles. *International Journal on Digital Libraries*, 23:289 – 301, 2022. URL `https://api.semanticscholar.org/CorpusID:250729995`.

Shaw, P., Joshi, M., Cohan, J., Berant, J., Pasupat, P., Hu, H., Khandelwal, U., Lee, K., and Toutanova, K. From pixels to ui actions: Learning to follow instructions via graphical user interfaces. *ArXiv*, abs/2306.00245, 2023. URL `https://api.semanticscholar.org/CorpusID:258999511`.

Shen, S., Li, L. H., Tan, H., Bansal, M., Rohrbach, A., Chang, K.-W., Yao, Z., and Keutzer, K. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*, 2021.

Shi, T., Karpathy, A., Fan, L. J., Hernández, J. Z., and Liang, P. World of bits: An open-domain platform for web-based agents. In *International Conference on Machine Learning*, 2017. URL `https://api.semanticscholar.org/CorpusID:34953552`.

Shtedritski, A., Rupprecht, C., and Vedaldi, A. What does clip know about a red circle? visual prompt engineering for vlms. *ArXiv*, abs/2304.06712, 2023. URL `https://api.semanticscholar.org/CorpusID:258108138`.

Sridhar, A., Lo, R., Xu, F. F., Zhu, H., and Zhou, S. Hierarchical prompting assists large language model on web navigation. *ArXiv*, abs/2305.14257, 2023. URL `https://api.semanticscholar.org/CorpusID:258841249`.

Yan, A., Yang, Z., Zhu, W., Lin, K., Li, L., Wang, J., Yang, J., Zhong, Y., McAuley, J., Gao, J., Liu, Z., and Wang, L. Gpt-4v in wonderland: Large multimodal models for zero-shot smartphone gui navigation. 2023. URL `https://api.semanticscholar.org/CorpusID:265149992`.

Yang, J., Zhang, H., Li, F., Zou, X., yue Li, C., and Gao, J. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *ArXiv*, abs/2310.11441, 2023a. URL `https://api.semanticscholar.org/CorpusID:264172201`.

Yang, L., Wang, Y., Li, X., Wang, X., and Yang, J. Fine-grained visual prompting. *ArXiv*, abs/2306.04356, 2023b. URL `https://api.semanticscholar.org/CorpusID:259096008`.

Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.-C., Liu, Z., and Wang, L. The dawn of lmms: Preliminary explorations with gpt-4v(ision). *ArXiv*, abs/2309.17421, 2023c. URL `https://api.semanticscholar.org/CorpusID:263310951`.

Yao, S., Chen, H., Yang, J., and Narasimhan, K. Webshop: Towards scalable real-world web interaction with grounded language agents. *ArXiv*, abs/2207.01206, 2022. URL `https://api.semanticscholar.org/CorpusID:250264533`.

You, H., Zhang, H., Gan, Z., Du, X., Zhang, B., Wang, Z., Cao, L., Chang, S.-F., and Yang, Y. Ferret: Refer and ground anything anywhere at any granularity. *ArXiv*, abs/2310.07704, 2023. URL `https://api.semanticscholar.org/CorpusID:263834718`.

Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., Wei, C., Yu, B., Yuan, R., Sun, R., Yin, M., Zheng, B., Yang, Z., Liu, Y., Huang, W., Sun, H., Su, Y., and Chen, W. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. *ArXiv*, abs/2311.16502, 2023. URL `https://api.semanticscholar.org/CorpusID:265466525`.

Zhang, X., Lu, Y., Wang, W., Yan, A., Yan, J., Qin, L., Wang, H., Yan, X., Wang, W. Y., and Petzold, L. R. Gpt-4v(ision) as a generalist evaluator for vision-language tasks. *ArXiv*, abs/2311.01361, 2023. URL `https://api.semanticscholar.org/CorpusID:264935635`.

Zhao, Y., Lin, Z., Zhou, D., Huang, Z., Feng, J., and Kang, B. Bubogpt: Enabling visual grounding in multi-modal llms. *ArXiv*, abs/2307.08581, 2023. URL `https://api.semanticscholar.org/CorpusID:259937702`.

Zhong, W., Cui, R., Guo, Y., Liang, Y., Lu, S., Wang, Y., Saied, A. S. S., Chen, W., and Duan, N. Agieval: A human-centric benchmark for evaluating foundation models. *ArXiv*, abs/2304.06364, 2023. URL `https://api.semanticscholar.org/CorpusID:258108259`.

Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., and Torralba, A. Scene parsing through ade20k dataset. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5122–5130, 2017. URL `https://api.semanticscholar.org/CorpusID:5636055`.

Table of Content:

1. Appendix A: Offline Experiments Method Details
2. Appendix B: Markup Type Ablation Study
3. Appendix C: Online Experiment Details
4. Appendix D: Offline Experiment Prompts
5. Appendix E: Error Examples for Grounding via Image Annotation
6. Appendix F: Error Example for Action Generation
7. Appendix G: Strong Capability of Planning
8. Appendix H: Challenges in Grounding via Textual Choices
9. Appendix I: Knowledge and Reasoning Requirements
10. Appendix J: Path Variation and Awareness of Error Correction

A. Offline Experiments Method Details

FLAN-T5. We fine-tune FLAN-T5 using a left-to-right language modeling objective with the target sequence of ground-truth actions in the Mind2Web training data. The fine-tuned FLAN-T5 then serves as the backbone for inference, enabling action generation in the target format for parsing.

BLIP-2-T5. The BLIP-2 model combines a vision encoder and an LLM with a bridging component for modality connection. We jointly fine-tune the LLM and the bridge module on Mind2Web training data while keeping the vision encoder frozen. For the vision encoder, we leverage the ViT-L/14 pre-trained from CLIP (Radford et al., 2021) with an image resolution of 2,048. To ensure a fair comparison with the FLAN-T5-based text-only model, we choose FLAN-T5 as the language model and initialize it with the parameters fine-tuned on Mind2Web.

GPT-3.5 and GPT-4. We also conduct experiments with text-only LLMs, specifically GPT-3.5-turbo-0613 and GPT-4-turbo-1106-preview, using in-context learning in 3-shot settings. We use the same multiple-choice formulation and include three demonstration examples for in-context learning as specified in MindAct.

SeeAct We experiment with GPT-4-vision-preview, Gemini Pro Vision, LLaVA-1.5. Gemini Pro Vision supports only single-turn conversations; therefore, we merge the two turns used in other models for compatibility.

CogAgent We utilize the cogagent-chat-hf checkpoint that hasn't been fine-tuned on Mind2Web for experiments.

B. Markup Type Ablation Study

The markup types might influence model performance as shown in Yang et al. (2023a); Yan et al. (2023); Koh et al. (2024). We first tested grounding via image annotation through different types of text labels of numerical value, single-digit characters, and two-digit characters, at two different positions to choose a relatively better markup type. The results are shown in Table 5.

Table 5: Grounding via Image Annotation with different markup types and locations. Method with * mark means the annotated image is used in action generation.

Label	Location	Ele. Acc	Op. F1	Step SR
Number	Bottom-Left	**27.0**	73.7	**24.3**
	Bottom-Center	23.0	76.4	21.8
Single Letter	Bottom-Left	19.4	**81.0**	17.2
	Bottom-Center	19.7	78.8	19.7
Double Letter	Bottom-Left	19.8	68.3	18.3
	Bottom-Center	22.4	74.6	22.4
Number*	Bottom-Left*	26.6	73.9	22.3

C. Online Experiment Details

We develop an online evaluation tool using Playwright to load webpages, acquire textual representation of interactive elements, and perform operations generated by web agents. We manually monitor each step of the model and assess whether it finishes the tasks. Attempts to log in, make final submissions, or perform other potentially harmful actions are prohibited to avoid negative consequences.

MindAct. We adhere strictly to the original settings in MindAct-FLAN-T5 and MindAct-GPT-4, as in offline experiment. For consistency with the MindAct framework and models, we use the scripts from Mind2Web for processing webpages, elements, and generating options, employing the same action space: `Click`, `Type`, and `Select`.

SeeAct$_{\text{Oracle}}$. In the oracle setting, we manually implement the model's intended actions. The action history is automatically generated by the model, with an added requirement to summarize actions in the "Element", "Operation", "Value" format in another turn conversation. To avoid overly long screenshots, we use screenshots of current views, and hence allow intentions of scrolling.

SeeAct$_{\text{Choice}}$. For a relatively fair comparison, we still adopt the ranker and top-50 candidate setting, batching them into three option groups as described in offline experiments. We enable `PRESS ENTER` and `TERMINATE` for the model to make confirmation or stop the process.

During our tests, pop-up ads on webpages were manually closed. The MindAct model, not trained on handling pop-up ads, lacks the feature to automatically manage them, potentially causing stalls. In contrast, SEEACT models can proactively suggest closing ads through visual analysis and reasoning.

D. Offline Experiment Prompts

The prompt for action generation is shown in Table 6. For grounding via textual choices, image annotation, and element attributes, the prompts are shown in Tables 7 to 9, along with specific tasks and examples in Figures 4 to 8.

Table 6: Prompt for SEEACT Action Generation with LMMs.

System Role	Imagine that you are imitating humans doing web navigation for a task step by step. At each stage, you can see the webpage like humans by a screenshot and know the previous actions before the current step decided by yourself through recorded history. You need to decide on the first following action to take. You can click an element with the mouse, select an option, or type text with the keyboard. (For your understanding, they are like the click(), select_option() and type() functions in playwright respectively) One next step means one operation within the three.
Action Generation	You are asked to complete the following task: {TASK} Previous Actions: {PREVIOUS ACTIONS} The screenshot below shows the webpage you see. Follow the following guidance to think step by step before outlining the next action step at the current stage: (Current Webpage Identification) Firstly, think about what the current webpage is. (Previous Action Analysis) Secondly, combined with the screenshot, analyze each step of the previous action history and their intention one by one. Particularly, pay more attention to the last step, which may be more related to what you should do now as the next step. (Screenshot Details Analysis) Closely examine the screenshot to check the status of every part of the webpage to understand what you can operate with and what has been set or completed. You should closely examine the screenshot details to see what steps have been completed by previous actions even though you are given the textual previous actions. Because the textual history may not clearly and sufficiently record some effects of previous actions, you should closely evaluate the status of every part of the webpage to understand what you have done. (Next Action Based on Webpage and Analysis) Then, based on your analysis, in conjunction with human web browsing habits and the logic of web design, decide on the following action. And clearly outline which element in the webpage users will operate with as the first next target element, its detailed location, and the corresponding operation. To be successful, it is important to follow the following rules: 1. You should only issue a valid action given the current observation. 2. You should only issue one action at a time.

Table 7: Prompt for SEEACT grounding via element attributes. We make a slight modification to enhance action generation and only show the modified part here to save space, as well as the prompts in Table 8 and Table 9.

System Role	Same as Table 6
Action Generation	Slightly modified from Table 6 ... (Next Action Based on Webpage and Analysis) Then, based on your analysis, in conjunction with human web browsing habits and the logic of web design, decide on the following action. And clearly outline which element in the webpage users will operate with as the first next target element, its detailed location, and the corresponding operation. Please also closely examine the screenshot to adequately describe its position relative to nearby elements and its textual or visual content (if it has). If you find multiple elements similar to your target element, use a more precise description to ensure people can distinguish your target element from them through your answer. ...
Format Answer	(Final Answer) Finally, conclude your answer using the format below. Ensure your answer is strictly adhering to the format provided below. Please do not leave any explanation in your answers of the final standardized format part, and this final part should be clear and certain. The element, element type, element text, action and value should be in five separate lines. Format: ELEMENT: Please describe which element you need to operate with. Describe it as detailed as possible, including what it is and where it is. ELEMENT TYPE: Please specify its type from these options: BUTTON, TEXTBOX, SELECTBOX, or LINK. ELEMENT TEXT: Please provide the exact text displayed on the element. Do not invent or modify the text; reproduce it as-is from the screenshot. ACTION: Choose an action from {CLICK, TYPE, SELECT}. VALUE: Provide additional input based on ACTION. The VALUE means: If ACTION == TYPE, specify the text to be typed. If ACTION == SELECT, specify the option to be chosen. If ACTION == CLICK, write "None".

Table 8: Prompt for SEEACT grounding via textual choices.

System Role	Same as Table 6
Action Generation	Same as Table 6
Referring Description	(Reiteration) First, reiterate your next target element, its detailed location, and the corresponding operation. (Multichoice Question) Below is a multi-choice question, where the choices are elements in the webpage. From the screenshot, find out where and what each one is on the webpage. Then, determine whether one matches your target element. Please examine the choices one by one. Choose the matching one. If multiple options match your answer, choose the most likely one by re-examining the screenshot, the choices, and your further reasoning. If none of these elements match your target element, please select [None of the other options match the correct element]. A. [CHOICE A] B. [CHOICE B] ...
Format Answer	(Final Answer) Finally, conclude your answer using the format below. Ensure your answer is strictly adhering to the format provided below. Please do not leave any explanation in your answers of the final standardized format part, and this final part should be clear and certain. The element choice, action, and value should be in three separate lines. Format: ELEMENT: The uppercase letter of your choice. ACTION: Choose an action from {CLICK, TYPE, SELECT}. VALUE: Provide additional input based on ACTION. The VALUE means: If ACTION == TYPE, specify the text to be typed. If ACTION == SELECT, specify the option to be chosen. If ACTION == CLICK, write "None".

Table 9: Prompt for SEEACT grounding via image annotation.

System Role	Same as Table 6
Action Generation	Same as Table 6
Referring Description	(Reiteration) First, reiterate your next target element, its detailed location, and the corresponding operation. (Verification with the Screenshot) Then, please closely re-examine the screenshot to find whether your target element is marked by a red bounding box and has a white number on a black background at the bottom left corner of the bounding box, which is positioned closely next to the bounding box. If yes, use that number for your final answer. If not, please do not make them up. If it is not marked, please output "NA" as your target element in the following final answer part.
Format Answer	(Final Answer) Finally, conclude your answer using the format below. Ensure your answer is strictly adhering to the format provided below. Please do not leave any explanation in your answers of the final standardized format part, and this final part should be clear and certain. The element choice, action, and value should be in three separate lines. Format: ELEMENT: The number of your choice. ACTION: Choose an action from {CLICK, TYPE, SELECT}. VALUE: Provide additional input based on ACTION. The VALUE means: If ACTION == TYPE, specify the text to be typed. If ACTION == SELECT, specify the option to be chosen. If ACTION == CLICK, write "None".

E. Error Examples for Grounding via Image Annotation

In the method of grounding via image annotation, we observe significant hallucination errors that can be classified into the following categories:

Making up bounding box & label. In our grounding method, if the correct element is absent from the set of candidate elements, the model is anticipated to generate "NA" as the answer. However, as depicted in Figure 10 and Figure 11, the model erroneously claims the element is included within a red bounding box and makes up a wrong index label as the answer.

Failure to link bounding boxes with the correct labels. Another challenge arises in accurately linking bounding boxes to their corresponding index labels. This challenge can be attributed to both LMMs' limitations in understanding relative spatial positions and the complex, dense layout of webpage elements. The model often mistakenly associates the labels of adjacent elements (as illustrated in Figure 12 and Figure 13), rather than accurately predicting the intended index label for the targeted element.

F. Error Examples for Action Generation

Although GPT-4V demonstrates remarkable capability in generating action descriptions, it does not always generate the optimal next step actions, as shown in Figure 14.

G. Strong Capability of Planning

GPT-4V shows remarkable understanding and planning capabilities during our experiments. As depicted in Figure 15, the model is capable of understanding the website and generating a full plan for the given task involving multiple low-level tasks. Specifically, GPT-4V could understand reasonably well about the process and the remaining work of the task by its careful examination of the webpage, as shown in Figure 16.

H. Challenges in Grounding via Textual Choices

Although textual choices achieved the best results among the three grounding approaches, it still suffers from challenges of similar or identical elements which are common in webpages. The model tends to choose the first text choice that seemingly corresponds to its intention. Moreover, this is inevitable, as web pages indeed contain many elements that may even have exact identical HTML information, as the "Schedule" button shown in Figure 17.

I. Knowledge and Reasoning Requirements

Some tasks require a certain degree of reasoning and knowledge, which may be challenging for fine-tuned models like MindAct. For instance, the task in Figure 18 necessitates the model to know the specific district of Dublin in Virginia. In the task of Figure 19, the model correctly provided the IATA airport code of airports in Indira Gandhi and Los Cabos.

J. Path Variation and Awareness of Error Correction

On webpages, multiple paths often exist to accomplish a given task. For instance, varying the execution order of actions within an interchangeable sequence can result in diverse routes to task completion. Additionally, the agent can navigate to different webpages but still accomplish the give tasks. Figure 20 presents a straightforward example where the model chose a more direct route that differs from the ground truth annotated in the dataset.

When running on live website, the agent's previous action histories is likely to be filled with redundant, unnecessary, erroneous, failed operations generated, or merely exploratory attempts by the model, resulting in a final path that deviates significantly from the ground truth. Despite these circumstances, the model can still accomplish the task amidst numerous incorrect explorations. The process of exploration and correction requires the model to possess a sense of self-correction. As shown in Figure 21, GPT-4V demonstrates this awareness of correcting errors caused by previous steps.

Example of Grounding via Element Attributes – 1st Turn

Input – 1st turn: Action Generation Prompts
Task: View the profile of a Wedding Photographer near 10203 for a 4 hour wedding on april 13
Previous Actions:
[textbox] Search on Thumbtack -> TYPE: wedding photographer
[textbox] Zip code -> TYPE: 10203

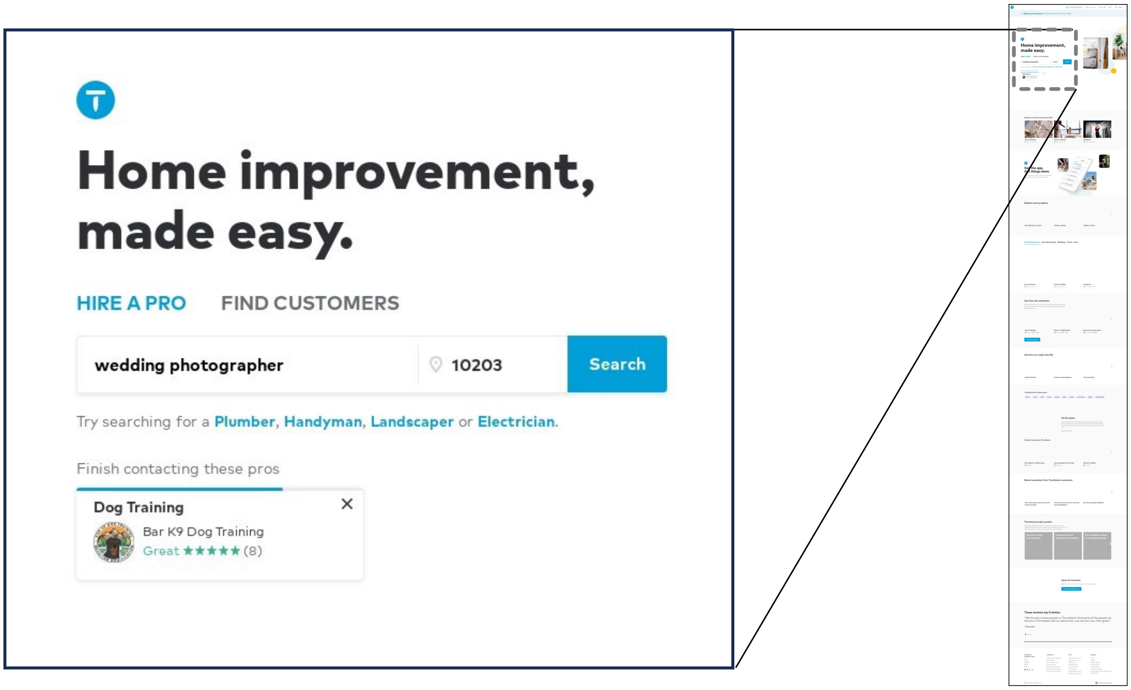

Home improvement, made easy.

HIRE A PRO FIND CUSTOMERS

wedding photographer | 10203 | Search

Try searching for a Plumber, Handyman, Landscaper or Electrician.

Finish contacting these pros

Dog Training
Bar K9 Dog Training
Great ★★★★★ (8)

GPT-4V: (Current Webpage Identification)
The current webpage appears to be the homepage of Thumbtack, which is a local services marketplace where users can find professionals for various services, including wedding photography.

(Previous Action Analysis)
The previous actions indicate that the user has typed "wedding photographer" into the search box and entered the zip code "10203" to find a wedding photographer near that location. The intention behind these actions is to search for a wedding photographer in the specified area.

(Screenshot Details Analysis)
The screenshot shows the Thumbtack homepage with the search bar at the top, containing the previously entered search term "wedding photographer" and the zip code "10203." Below the search bar, there are suggestions for different services and popular articles, but these are not relevant to the task at hand. The search button next to the zip code field is visible, indicating that the search has not yet been initiated.

(Next Action Based on Webpage and Analysis)
Based on the analysis, the next logical step is to initiate the search for wedding photographers in the specified zip code area. To do this, the user should click on the "Search" button.

The "Search" button is located to the right of the zip code field at the top of the webpage. It is a rectangular button with rounded corners, and the text "Search" is visible on it. This button is directly adjacent to the zip code field where "10203" has been entered.

The first next action to take is to click the "Search" button to proceed with finding a wedding photographer for the specified requirements.

Figure 4: An example of action generation in grounding via element attributes.

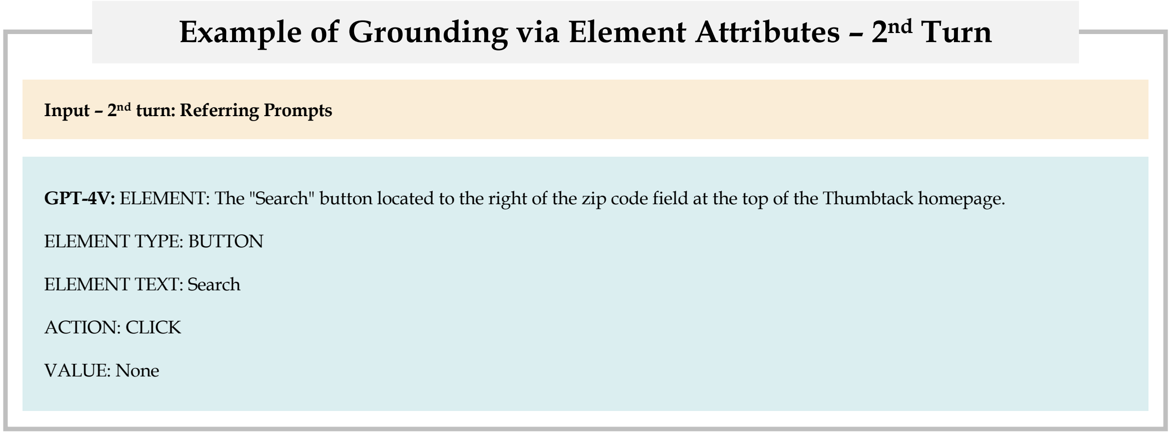

Example of Grounding via Element Attributes – 2nd Turn

Input – 2nd turn: Referring Prompts

GPT-4V: ELEMENT: The "Search" button located to the right of the zip code field at the top of the Thumbtack homepage.

ELEMENT TYPE: BUTTON

ELEMENT TEXT: Search

ACTION: CLICK

VALUE: None

Figure 5: An example of grounding via element attributes after action generation.

Example of Grounding via Textual Choices – 1st Turn

Input – 1st turn: Action Generation Prompts
Task: View the profile of a Wedding Photographer near 10203 for a 4 hour wedding on april 13
Previous Actions:
[textbox] Search on Thumbtack -> TYPE: wedding photographer
[textbox] Zip code -> TYPE: 10203

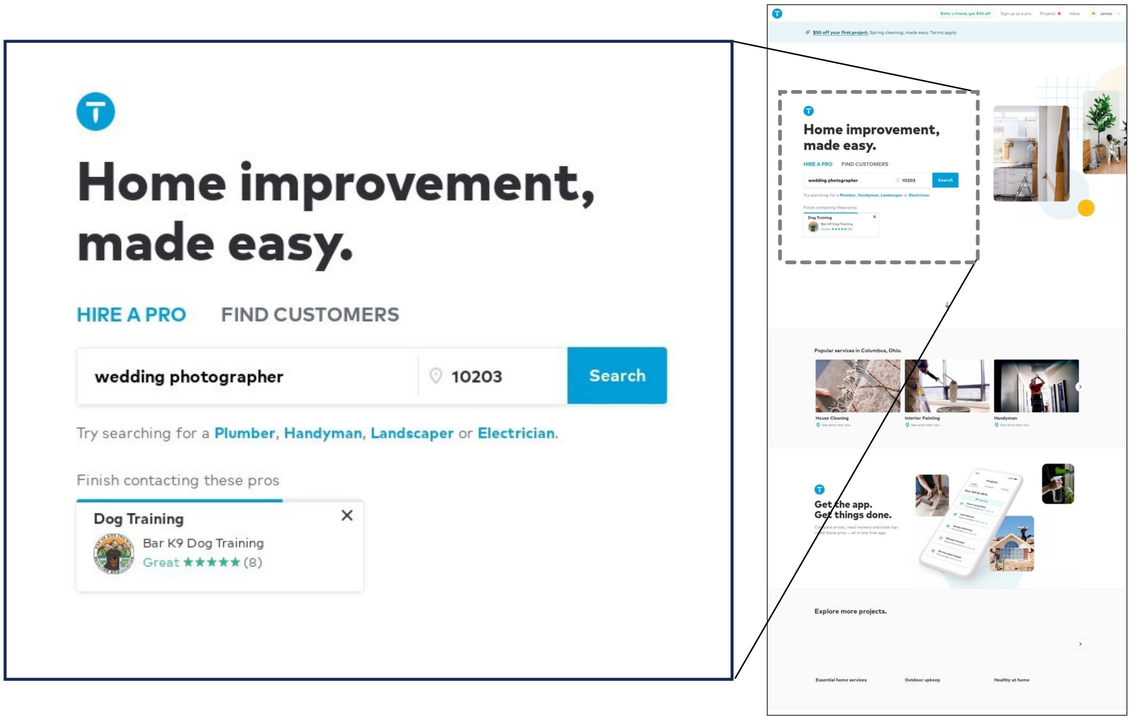

GPT-4V: (Current Webpage Identification)
The current webpage appears to be the homepage of Thumbtack, a platform for hiring local professionals for various services.

(Previous Action Analysis)
The previous actions indicate that the user is looking for a wedding photographer. The user has typed "wedding photographer" into a search box and entered the zip code "10203" into another box, presumably to localize the search to that area.

(Screenshot Details Analysis)
The screenshot shows the Thumbtack homepage with the search fields already populated with "wedding photographer" and "10203." The user has not yet initiated the search, as the search results are not displayed on the screen. The "Search" button is visible next to the zip code field, indicating that the next logical step is to execute the search.

(Next Action Based on Webpage and Analysis)
Based on the analysis, the next action should be to initiate the search for wedding photographers in the specified zip code area. The first next target element is the "Search" button located to the right of the zip code field.

The corresponding operation is to click the "Search" button to proceed with finding a list of wedding photographers near the zip code 10203.

Figure 6: An example of action generation in grounding via textual choices.

Example of Grounding via Textual Choices – 2nd Turn

Input – 2nd turn: Referring Prompts
Choices:
A. <button id=0 button header dropdown button> <div> <span> JS </span>
B. <a id=1> <span> Refer a friend, get $50 off </span>
C. <a id=2> <div> Profile </div> </a>
D. <button id=3 button> HIRE A PRO </button>
E. <button id=4 button> FIND CUSTOMERS </button>
F. <input id=5 search on thumbtack wedding photographer whats on your
G. <input id=6 zip code zip_code 10203 />
H. <button id=7 submit search> Search </button>
I. <span id=8> , </span>
J. <span id=9> . </span>
K. <button id=10 button navigation arrow down />
L. <a id=11> <source image/webp /> <div furniture moving and heavy
M. <input id=12 search on thumbtack whats on your to-do list?
N. <input id=13 zip code zip_code 43235 />
O. <button id=14 button> Cancel </button>
P. <button id=15 submit> Search </button>
Q. None of the other options match the correct element

GPT-4V: (Reiteration)
The next target element is the "Search" button located to the right of the zip code field.

(Multichoice Question)
H. <button id=7 submit search> Search </button>

(Final Answer)

ELEMENT: H

ACTION: CLICK

VALUE: None

Figure 7: An example of grounding via textual choices after action generation.

Example of Grounding via Image Annotation – 1st Turn

Input – 1st turn: Action Generation Prompts
Task: View the profile of a Wedding Photographer near 10203 for a 4 hour wedding on april 13
Previous Actions:
[textbox] Search on Thumbtack -> TYPE: wedding photographer
[textbox] Zip code -> TYPE: 10203

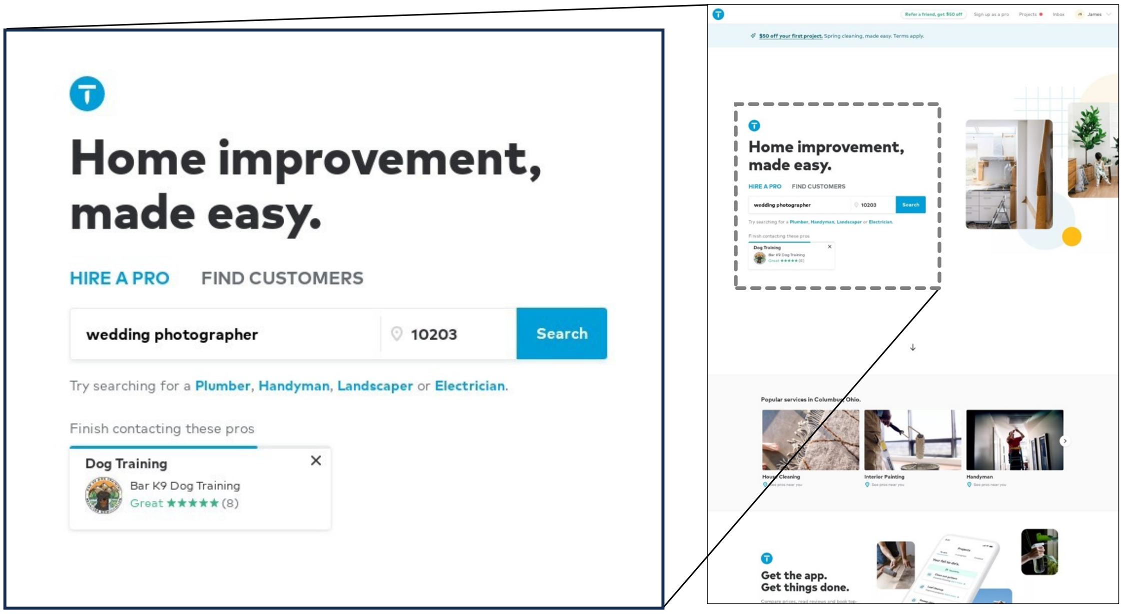

GPT-4V: (Current Webpage Identification)
The current webpage appears to be the homepage of Thumbtack, a local services marketplace where users can find professionals for various services, including home improvement, events, and more.

(Previous Action Analysis)
The first action was typing "wedding photographer" into a search box, indicating the user's intent to find a photographer for a wedding. The second action involved typing "10203" into a zip code field, specifying the location where the service is needed.

(Screenshot Details Analysis)
The screenshot shows the Thumbtack homepage with the search fields at the top. The "wedding photographer" text is visible in the search box, and "10203" is in the zip code field. There is a "Search" button next to these fields, which has not been clicked yet, as indicated by the presence of the search terms in the fields and the lack of search results on the page.

The first next target element is the "Search" button located to the right of the zip code field at the top of the webpage. The corresponding operation is to click this button to execute the search for wedding photographers near the zip code 10203.

Figure 8: An example of grounding via image annotation.

Example of Grounding via Image Annotation – 2nd Turn

Input – 2nd turn: Referring Prompts

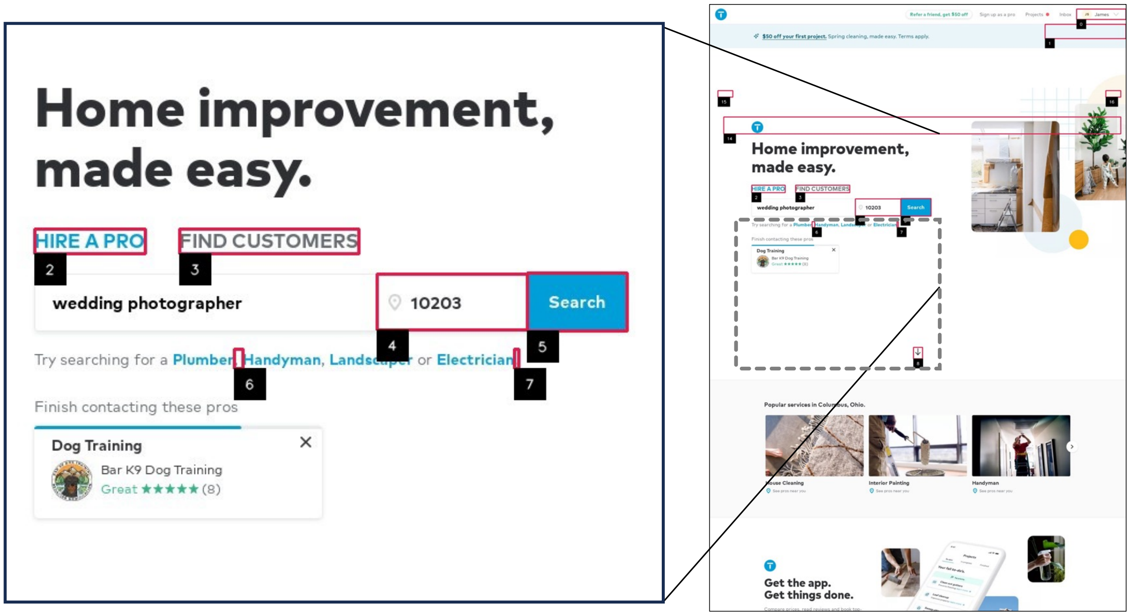

Home improvement, made easy.

wedding photographer

10203

Try searching for a Plumber, Handyman, Landscaper or Electrician

Finish contacting these pros

Dog Training

Bar K9 Dog Training

Great ★★★★★ (8)

Popular services in Columbus, Ohio.

Get the app.
Get things done.

GPT-4V: ELEMENT: 4

ACTION: CLICK

VALUE: None

Figure 9: An example of grounding via image annotation.

Makeup Image Annotation – Example 1

Task: As a Verizon user, finance a new blue iPhone 13 with 256gb along with monthly apple care.
Previous Actions:
...
[checkbox] AppleCare+ Theft and Loss for iPhone 13 - Monthly ... -> CLICK
[button] Add to Cart -> CLICK
[button] Let's Get Started -> CLICK
[label] Yes -> CLICK

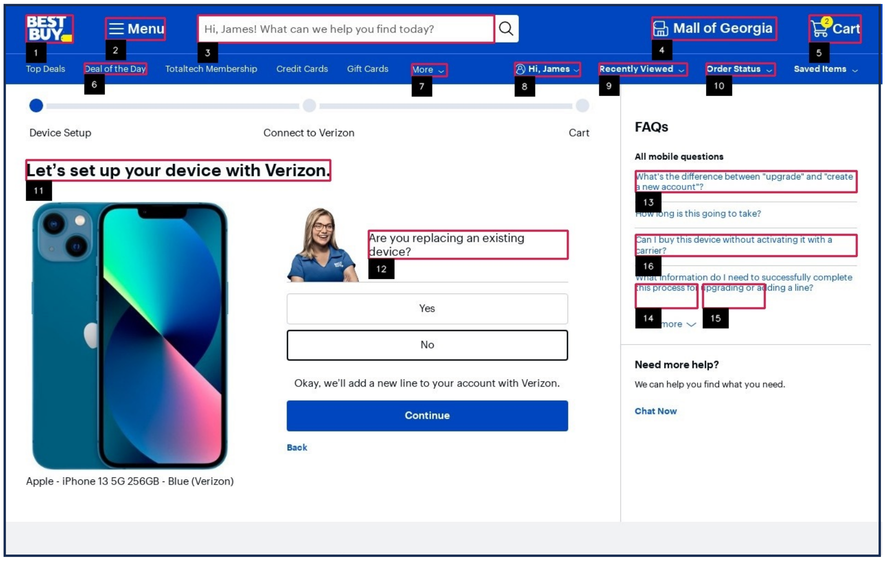

Let's set up your device with Verizon.

Are you replacing an existing device?

Okay, we'll add a new line to your account with Verizon.

Apple - iPhone 13 5G 256GB - Blue (Verizon)

What's the difference between "upgrade" and "create a new account"?

How long is this going to take?

Can I buy this device without activating it with a carrier?

What information do I need to successfully complete this process for upgrading or adding a line?

Need more help?

We can help you find what you need.

Action Generation: ...The first next target element is the "Continue" button located in the center of the page, just below the message about adding a new line to the Verizon account. The corresponding operation is to click this button.

Action Grounding: ELEMENT: **12**

ACTION: CLICK

VALUE: None

Figure 10: In this example, the model describes a correct element in action generation. However, the identified element is absent from the set of candidate elements. Despite this, the model erroneously assigns it the index number "12".

Makeup Image Annotation - Example 2

Task: Find me the cheapest Dodge Ram 1500 within 50 miles of 21122
Previous Actions:
[combobox] Make -> SELECT: Dodge
[combobox] Model -> SELECT: Ram 1500

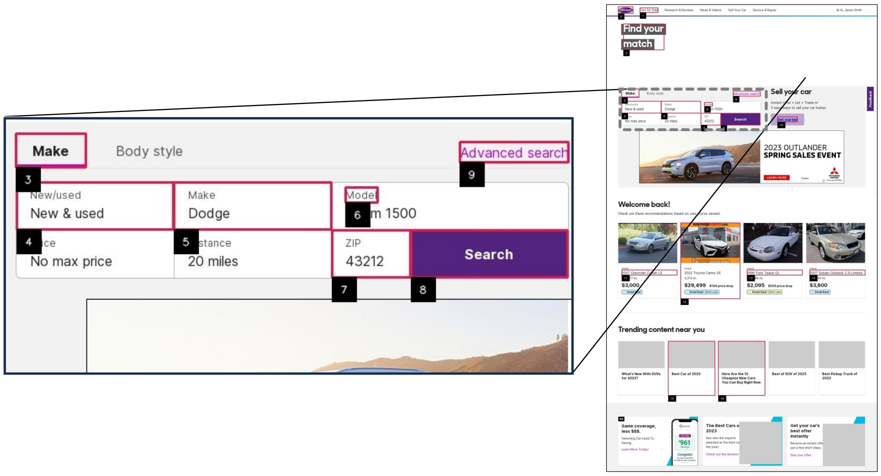

New/used: New & used
Make: Dodge
Model: …m 1500
…ce: No max price
…stance: 20 miles
ZIP: 43212

Action Generation: ... Next Action:
- Target Element: 'Distance' dropdown menu
- Location: Under the 'Model' dropdown, to the right of where '20 miles' is currently displayed
- Operation: Click on the 'Distance' dropdown menu to expand it and then select '50 miles'.

Grounding: ELEMENT: 5

ACTION: CLICK

VALUE: None

Figure 11: In this example, the model predicts the appropriate element. Nevertheless, the identified element is not present in the provided image options. Despite this, the model erroneously assigns it with the label ’5’ nearby.

Failure of Linking - Example 1

Task: Book a cheapest bundle and save option for 2 adults from Ahmedabad to Dubai on April 5 with free cancellation options, hotel should be a 3 star near Burj Khalifa with guest rating above 4, one night
Previous Actions:
[span] Bundle & Save -> CLICK

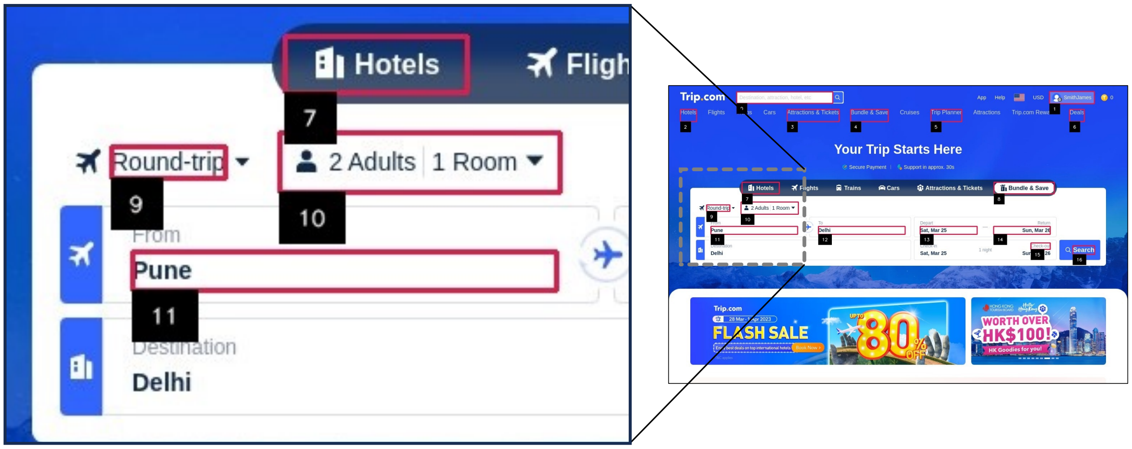

Action Generation: The first next target element is the "From" field, located in the search form at the top of the page, which currently has "Pune" as the input. The corresponding operation is to clear the current input and type "Ahmedabad."

Next action: Click on the "From" field and type "Ahmedabad."

Action Grounding:
ELEMENT: **10**

ACTION: TYPE

VALUE: Ahmedabad

Figure 12: In this case, while the model predicts the appropriate element, it incorrectly associates the element with the nearby label ’10’ instead of the correct label ’11’.

Failure of Linking - Example 2

Task: Find a psychiatrist who offers virtual appointmentsthat has experience with treating neurodevelopment disorders and accepts new patients.
Previous Actions:
[searchbox] Search doctors, conditions, or procedures -> TYPE: psychiatrist
[em] Psychiatrist -> CLICK

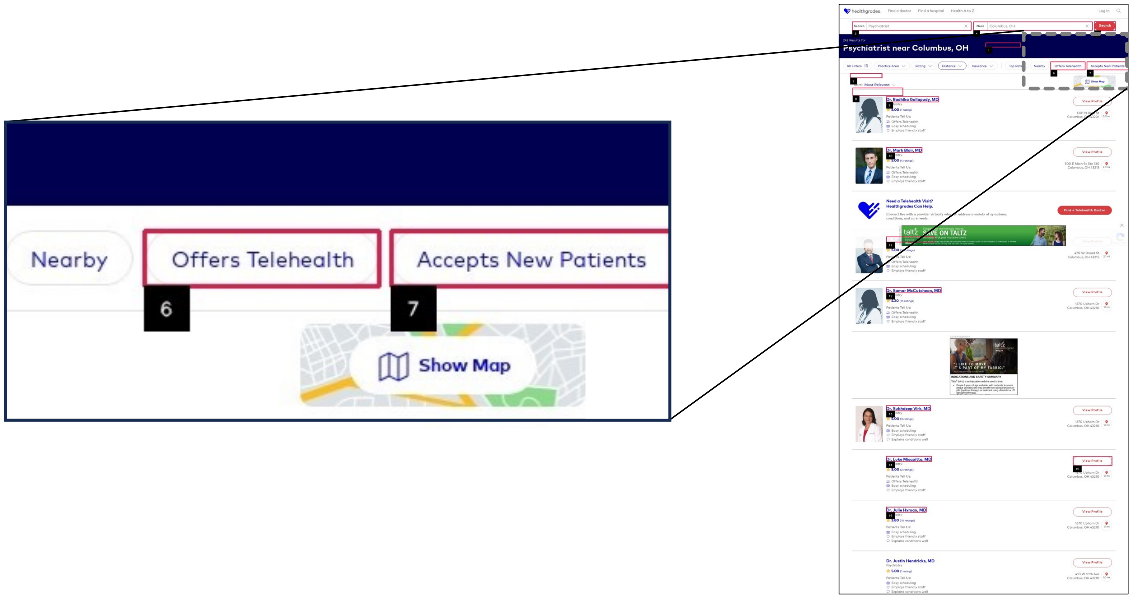

Action Generation: ...Next action:
- Click on the "Telehealth" filter located at the top of the page, in the row of filters, to apply this criterion to the search results.

Action Grounding:
ELEMENT: **7**

ACTION: CLICK

VALUE: None

Figure 13: In this case, while the model predicts the appropriate element, it incorrectly associates the element with the nearby label '7' instead of the correct label '6'.

Error in Action Generation

Task: Find a list of Tours that contain visits to the Louvre rated 5 stars.
Previous Actions:
[span] Top activities-> CLICK

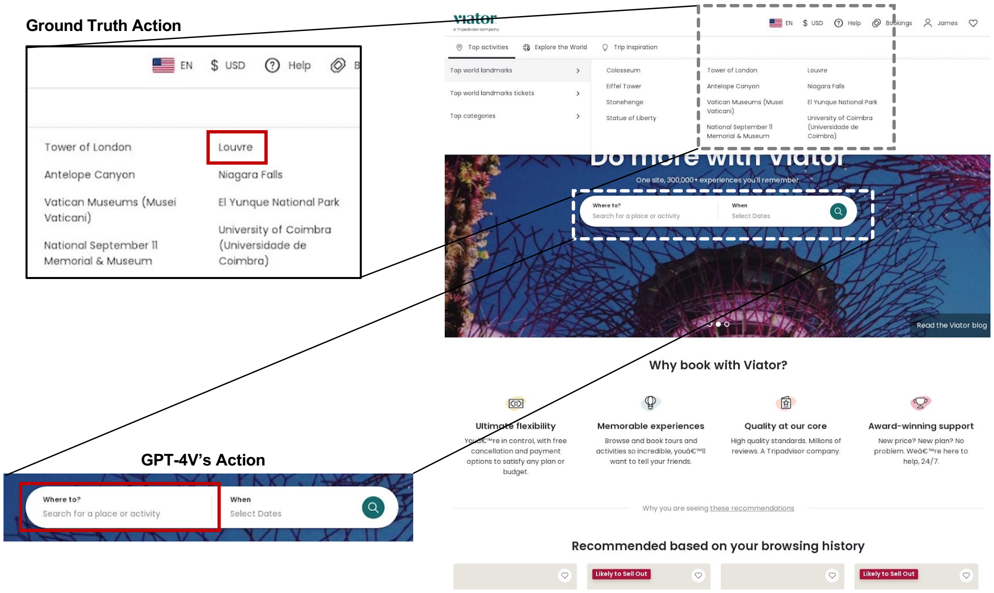

Ground Truth Action: Click the "Louvre" link at the top of the page.
GPT-4V's Prediction: Based on the user's task to find a list of tours that contain visits to the Louvre rated 5 stars, **the first next target element would be the search bar at the top of the page.** The user will likely need to click on the search bar, which is located in the upper portion of the webpage in the center, just below the main navigation menu. It has a placeholder text that reads "Search for a place or activity". The appropriate operation would be to click on this search bar and type "Louvre 5-star tours" to begin the search. This action is consistent with the user's goal of finding specific tours related to the Louvre with high ratings.

Figure 14: In this example, based on the previous actions, the model could take the action of clicking "Louvre" button. However, GPT-4V tends to call the search box for websites.

Strong Capability of Planning

Task: View the speakers that are bluetooth and wireless and filter the results to only show models that are on sale and cost less than $50.
Previous Actions:
None

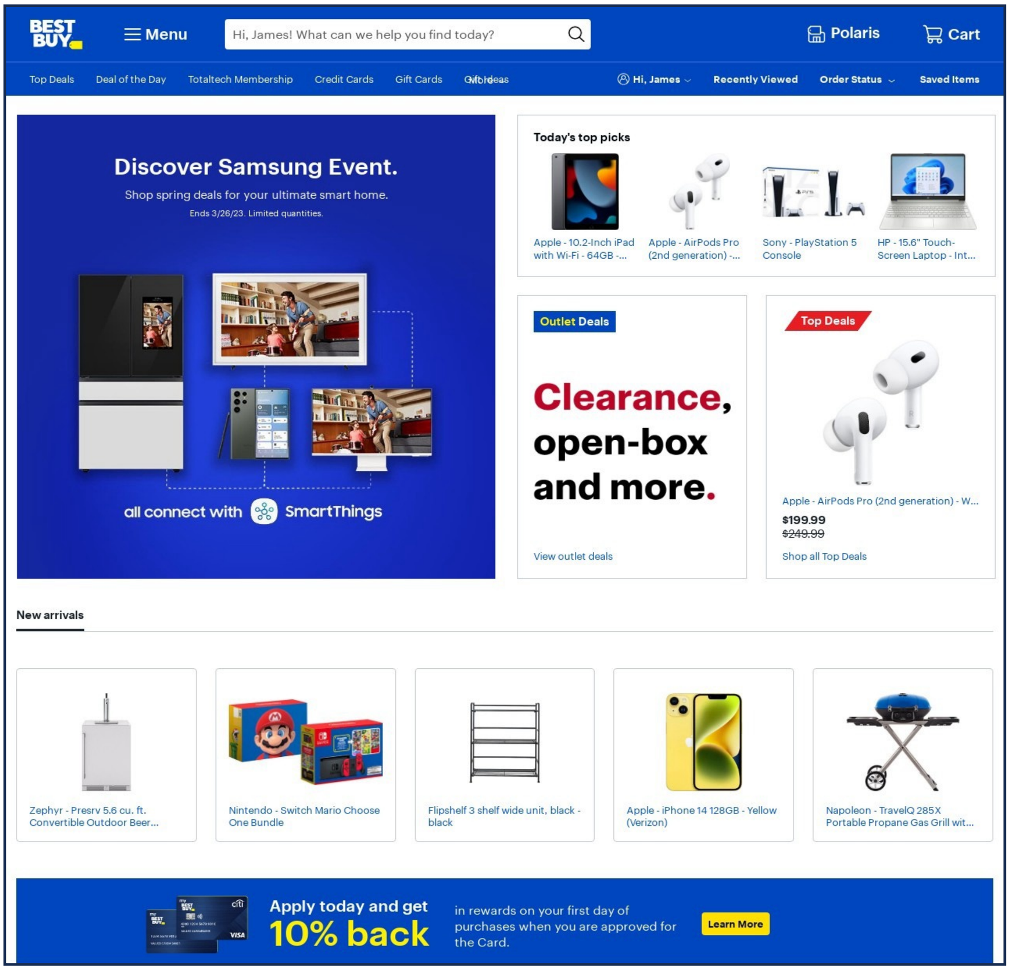
BEST BUY
Menu
Hi, James! What can we help you find today?
Polaris
Cart
Top Deals
Deal of the Day
Totaltech Membership
Credit Cards
Gift Cards
Gift Ideas
Hi, James
Recently Viewed
Order Status
Saved Items
Discover Samsung Event.
Shop spring deals for your ultimate smart home.
Ends 3/26/23. Limited quantities.
all connect with SmartThings
Today's top picks
Apple - 10.2-Inch iPad with Wi-Fi - 64GB -...
Apple - AirPods Pro (2nd generation) -...
Sony - PlayStation 5 Console
HP - 15.6" Touch-Screen Laptop - Int...
Outlet Deals
Clearance, open-box and more.
View outlet deals
Top Deals
Apple - AirPods Pro (2nd generation) - W...
$199.99
~~$249.99~~
Shop all Top Deals
New arrivals
Zephyr - Presrv 5.6 cu. ft. Convertible Outdoor Beer...
Nintendo - Switch Mario Choose One Bundle
Flipshelf 3 shelf wide unit, black - black
Apple - iPhone 14 128GB - Yellow (Verizon)
Napoleon - TravelQ 285X Portable Propane Gas Grill wit...
Apply today and get 10% back
in rewards on your first day of purchases when you are approved for the Card.
Learn More

GPT-4V: … To proceed with the task, the user should: **Locate the menu bar to find the category for audio equipment or speakers. This is typically found in a dropdown under a section such as "Products", "Categories", or directly as "Audio". Once the correct category is found, the user should click on it to navigate to the section of the website that lists speakers. After reaching the speakers section, the user will need to use the website's filtering options to select 'Bluetooth' and 'Wireless' features, and then set the price filter to show only the models that are on sale and cost less than $50.** …

Figure 15: In this example, the model generates a comprehensive plan for the task, including subsequent actions on following pages that are not currently visible.

Strong Capability of Understanding

Task: Rent a truck with the lowest rent with two dozen furniture pads for 100 miles at zip 08817 on April 12 at 2:30 pm rented truck is to be returned to the exact location and date, and the pickup and drop off will be at the nearest location.
Previous Actions:
…
[textbox] US City,State or Zip Code -> TYPE: 08817
[textbox] mm/dd/yyyy -> CLICK
[link] 12 -> CLICK
[select] Select Pick-up Time -> SELECT: 02:30 PM
[input] -> CLICK

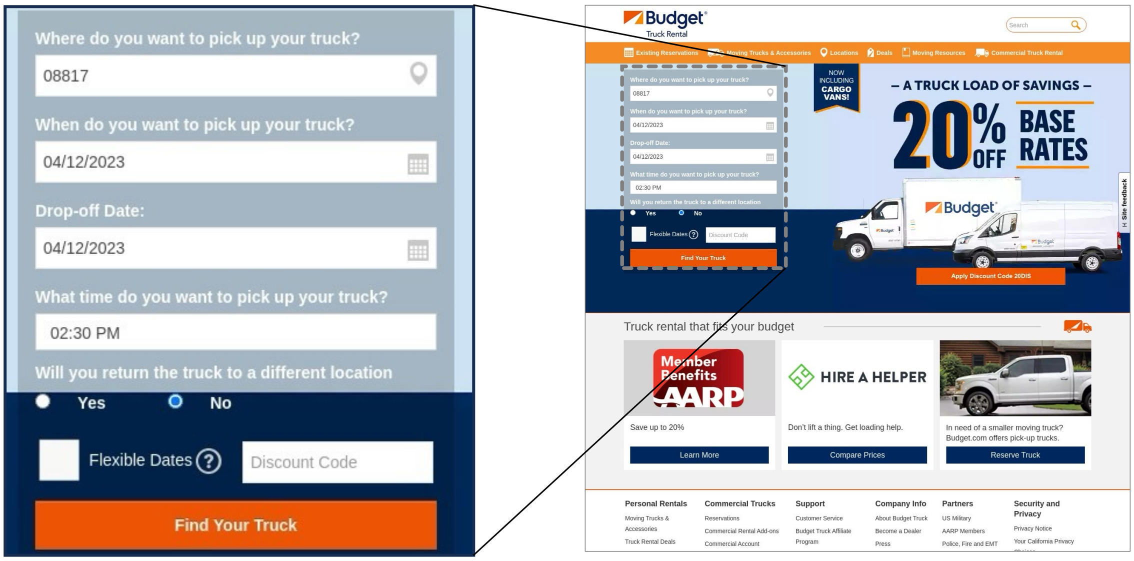

Where do you want to pick up your truck? 08817
When do you want to pick up your truck? 04/12/2023
Drop-off Date: 04/12/2023
What time do you want to pick up your truck? 02:30 PM
Will you return the truck to a different location ○ Yes ● No
☐ Flexible Dates ? Discount Code
Find Your Truck

GPT-4V: … The current webpage appears to be the truck rental reservation page of the Budget Truck Rental website. This page is designed for customers to input their rental details such as pickup location, dates, and times.
(Screenshot Details and Previous Action Analysis)
The screenshot shows that the user has already entered the pickup location (zip code 08817), selected the pickup date (April 12, 2023), and chosen the pickup time (2:30 PM). The drop-off date is also set to the same day, which aligns with the task requirement of returning the truck to the same location on the same date. The option "Will you return the truck to a different location" is set to "No," which is correct for the task at hand. The last action taken was a click, but it's not clear on what element the click was performed. However, based on the task, it's likely that the user is ready to proceed to the next step of finding a truck.
(Intended Action Based on Current Page and Previous Actions)
The next logical step for the user to take would be to click on the "Find Your Truck" button. This action will likely lead to a page where the user can select the type of truck they wish to rent and add any additional accessories, such as the two dozen furniture pads mentioned in the task. Since the task specifies renting the truck with the lowest rent, the user will need to compare the available options and select the most cost-effective one.
The "Find Your Truck" button is located in the central portion of the webpage, just below the input fields for the rental details. The corresponding operation would be to move the cursor over the "Find Your Truck" button and perform a click action.…

Figure 16: In this case, two critical pieces of information are inadequately captured by the textual history. Firstly, the website automatically set the drop-off date to the same day. Secondly, secondly, the 'No' button was selected (However it was missed in previous actions due to the button's lack of text). Nevertheless, the model discerns these details through a meticulous analysis of the screenshot, enabling it to make the correct decision for the next step.

Identical Elements

Task: Schedule a repair service for my gaming console closest to the zip code 10001 on March 24, anytime after 6 pm.
Previous Actions:
...
[[link] Schedule a Service -> CLICK
[p] Video Games & Entertainment -> CLICK
[heading] Schedule a repair or tune-up -> CLICK
[button] Update Location -> CLICK
[textbox] Enter Zip Code or City and State -> TYPE: 10001
[button] Find Stores -> CLICK
[button] Change Date -> CLICK
[button] March 24, 2023, is available -> CLICK
[combobox] Select a Time Slot -> SELECT: 6:20 p.m. – 6:40 p.m. (EDT)

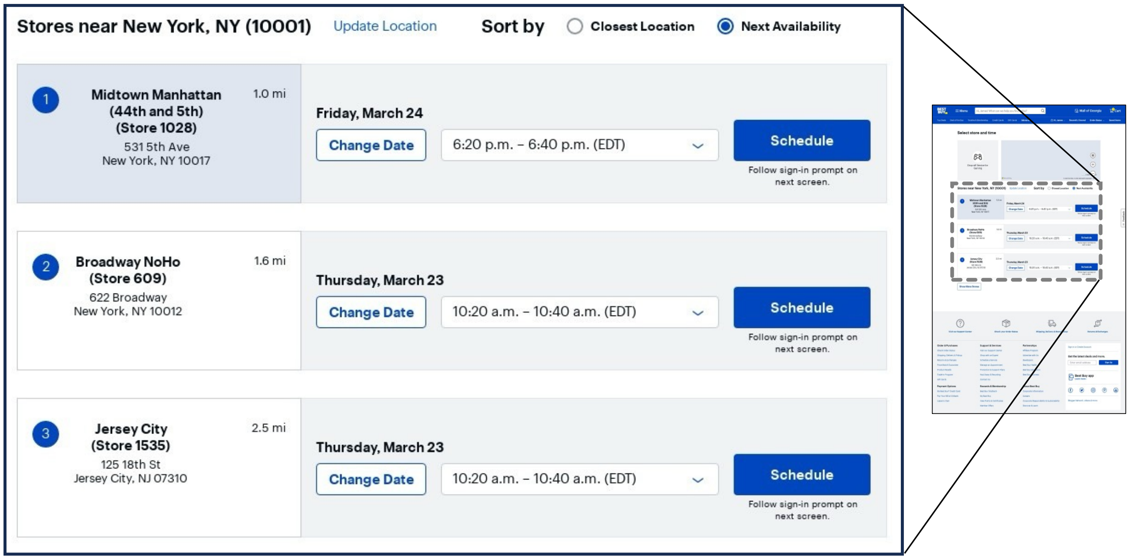
Stores near New York, NY (10001) Update Location Sort by ○ Closest Location ● Next Availability

1 Midtown Manhattan (44th and 5th) (Store 1028) 531 5th Ave New York, NY 10017	1.0 mi	Friday, March 24 Change Date 6:20 p.m. – 6:40 p.m. (EDT)	Schedule Follow sign-in prompt on next screen.
2 Broadway NoHo (Store 609) 622 Broadway New York, NY 10012	1.6 mi	Thursday, March 23 Change Date 10:20 a.m. – 10:40 a.m. (EDT)	Schedule Follow sign-in prompt on next screen.
3 Jersey City (Store 1535) 125 18th St Jersey City, NJ 07310	2.5 mi	Thursday, March 23 Change Date 10:20 a.m. – 10:40 a.m. (EDT)	Schedule Follow sign-in prompt on next screen.

GPT-4V: … The user is currently in the process of selecting a store location and time for the service.(Details and Previous Action Analysis)
The user initiated this process by navigating to the support & services section of the Best Buy menu, then selecting the option to schedule a service. They narrowed down their service type by selecting the "Video Games & Entertainment" category and chose to schedule a repair or tune-up. The user has also updated their location using the zip code 10001 to find stores near that area. They then selected the date of March 24, 2023, and chose the available time slot of 6:20 p.m. – 6:40 p.m. (EDT) for their service. This time slot fits the requirement of scheduling the service anytime after 6 pm on March 24th.From the screenshot, three store options near the zip code 10001 are listed. The closest one is the Midtown Manhattan location, which is 1.0 mile away from the provided zip code. This store already has the selected date and time of Friday, March 24, from 6:20 p.m. to 6:40 p.m. The last action from the previous action history indicates that the user has selected the time slot of 6:20 p.m. – 6:40 p.m. (EDT). This suggests that the user is now ready to finalize and confirm this scheduled service.(Intended Action Based on Current Page and Previous Actions)
Considering the details and the objective to schedule a repair service, the next logical step for the user would be to confirm and schedule the service at the selected store, date, and time. On the current page, the element that seems to be relevant to this action is the "Schedule" button associated with the selected store, date, and time. **Therefore, the next action step is:**
[button] Schedule (associated with the Midtown Manhattan location for the selected date and time) -> CLICK.

Figure 17: In this example, there are three identical 'Schedule' buttons, making it impossible for SEEACT$_{\text{Choice}}$ to distinguish among them. We empirically find that the model tends to choose the first one among the choices.

Knowledge Requirement – Example 1

Task: Find a driver training school in Dublin
Previous Actions:
…
[link] All -> CLICK
[link] Attending a Driver Training School -> CLICK
[link] Driver Training Schools -> CLICK
[link] Find a Driver Training School Near You -> CLICK

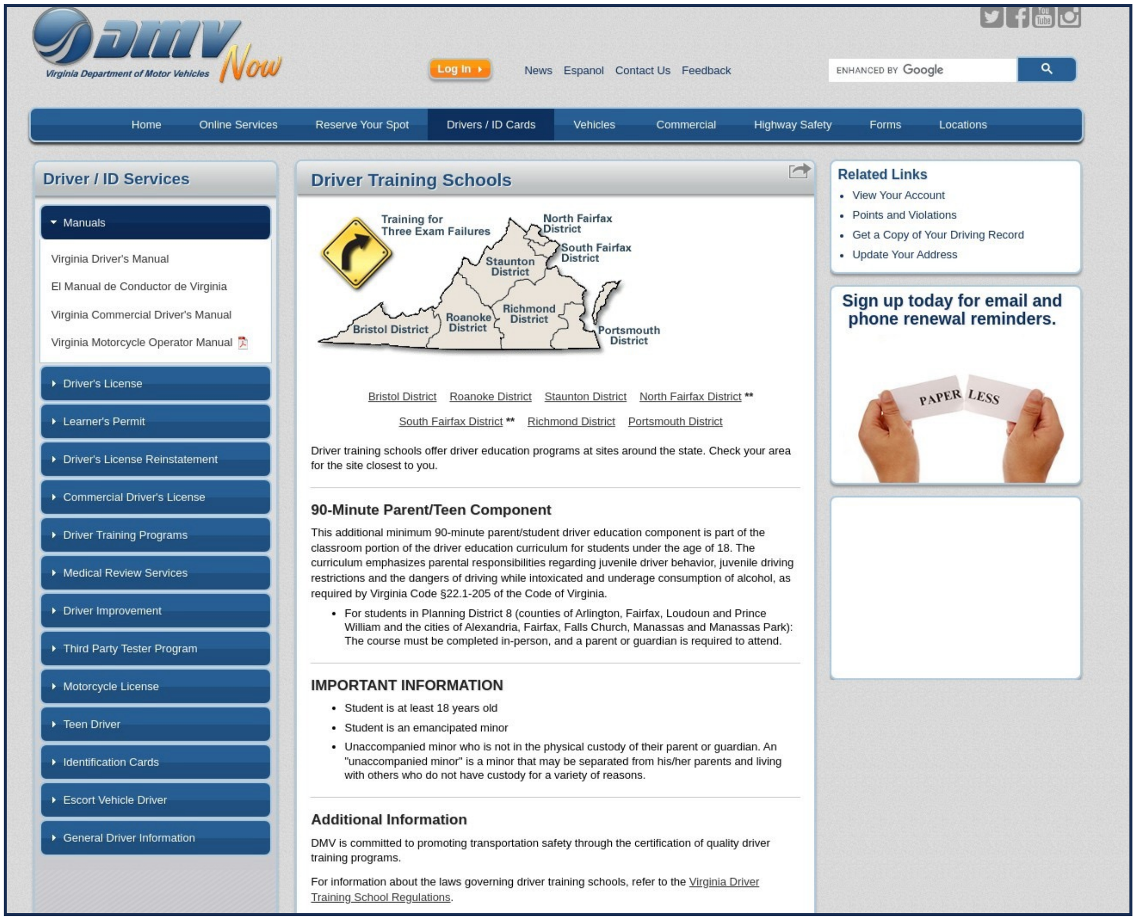

Driver Training Schools

Bristol District Roanoke District Staunton District North Fairfax District ** South Fairfax District ** Richmond District Portsmouth District

Driver training schools offer driver education programs at sites around the state. Check your area for the site closest to you.

90-Minute Parent/Teen Component

This additional minimum 90-minute parent/student driver education component is part of the classroom portion of the driver education curriculum for students under the age of 18. The curriculum emphasizes parental responsibilities regarding juvenile driver behavior, juvenile driving restrictions and the dangers of driving while intoxicated and underage consumption of alcohol, as required by Virginia Code §22.1-205 of the Code of Virginia.

- For students in Planning District 8 (counties of Arlington, Fairfax, Loudoun and Prince William and the cities of Alexandria, Fairfax, Falls Church, Manassas and Manassas Park): The course must be completed in-person, and a parent or guardian is required to attend.

IMPORTANT INFORMATION

- Student is at least 18 years old
- Student is an emancipated minor
- Unaccompanied minor who is not in the physical custody of their parent or guardian. An "unaccompanied minor" is a minor that may be separated from his/her parents and living with others who do not have custody for a variety of reasons.

Additional Information

DMV is committed to promoting transportation safety through the certification of quality driver training programs.

For information about the laws governing driver training schools, refer to the Virginia Driver Training School Regulations.

GPT-4V: … The current page has a list of districts, which implies that the user may need to select a district closest to Dublin, Virginia, to continue their search. **However, since there is no district explicitly named "Dublin," the user must know which district Dublin falls under or look for an option that allows for a more specific search.**
(Intended Action Based on Current Page and Previous Actions)
The first next target element that the user would likely need to operate with is the district name that includes or is closest to Dublin, Virginia. However, the screenshot does not provide enough detail to determine which district Dublin is in….

Figure 18: In this example, the task necessitates knowledge about which district Dublin is located in.

Knowledge Requirement – Example 2

Task: Find flights going from Indira Gandhi to Los Cabos.
Previous Actions:
[textbox] Origin -> TYPE: DEL

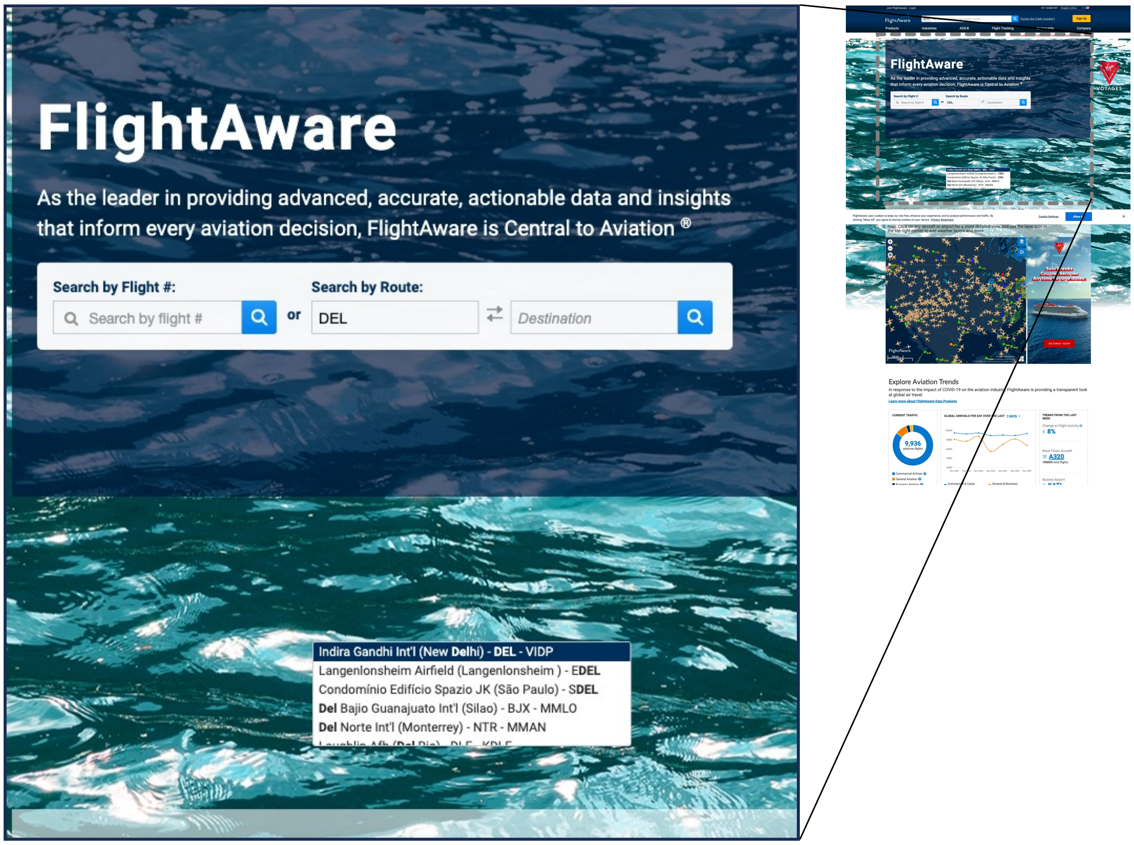

FlightAware
Search by Flight #:
Search by flight #
or
Search by Route:
DEL
Destination
Indira Gandhi Int'l (New Delhi) - DEL - VIDP
Langenlonsheim Airfield (Langenlonsheim) - EDEL
Condomínio Edifício Spazio JK (São Paulo) - SDEL
Del Bajio Guanajuato Int'l (Silao) - BJX - MMLO
Del Norte Int'l (Monterrey) - NTR - MMAN
Explore Aviation Trends

As the leader in providing advanced, accurate, actionable data and insights that inform every aviation decision, FlightAware is Central to Aviation ®

GPT-4V: ... Then, the user needs to identify the IATA code for Los Cabos International Airport, which is "SJD," and enter it into the destination field. Thus, the next target element is the destination field, which is located to the right of the origin field in the "Search by Route" section of the search interface.
Here is the action to be taken:
[textbox] Destination -> TYPE: SJD (after ensuring the origin field correctly contains "DEL").

Figure 19: In this example, the task requires knowledge of the IATA code for Los Cabos International Airport. GPT-4V accurately provides the correct code.

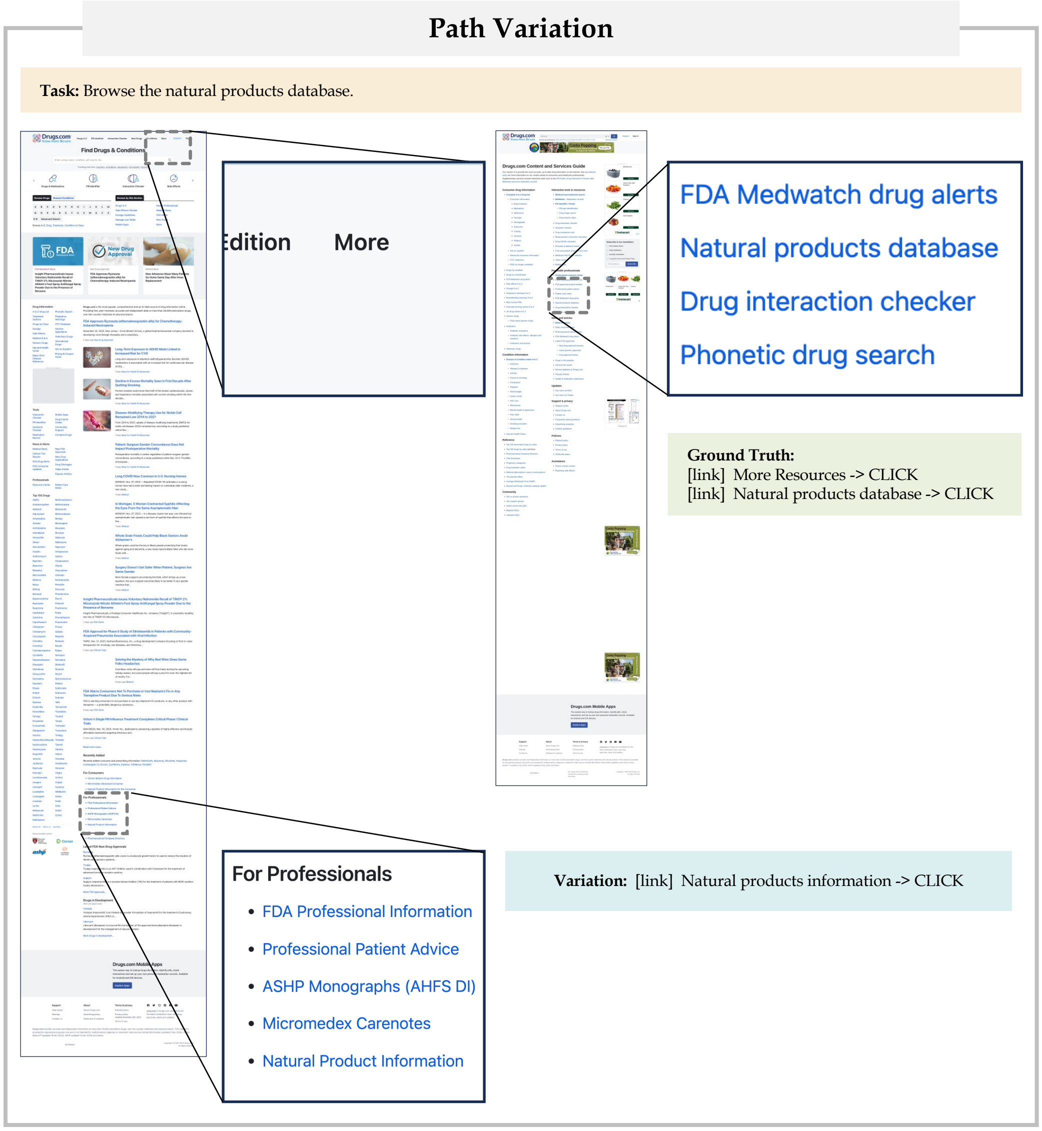
Edition
More
FDA Medwatch drug alerts
Natural products database
Drug interaction checker
Phonetic drug search
For Professionals
FDA Professional Information
Professional Patient Advice
ASHP Monographs (AHFS DI)
Micromedex Carenotes
Natural Product Information

Path Variation

Task: Browse the natural products database.

Ground Truth:
[link] More Resources -> CLICK
[link] Natural products database -> CLICK

Variation: [link] Natural products information -> CLICK

Figure 20: In this example, the ground truth in Mind2Web is to firstly click "More", then click "Natural products database" on the second page. In contrast, the model identifies a more direct approach, achieving the target page through 'Natural products information' on the first page.

Awareness of Error Correction

Task: Signup for virtual healthcare visit.
Previous Actions:
…
[span] James Smith -> CLICK
[button] Select patient and continue -> CLICK
[textbox] Date of birth -> CLICK
[textbox] Date of birth -> TYPE: 01/01/2000
[radio] Male -> CLICK
[radio] No -> CLICK
[textbox] Mobile phone number -> TYPE: 888888888
[textbox] Street address -> TYPE: 123 st

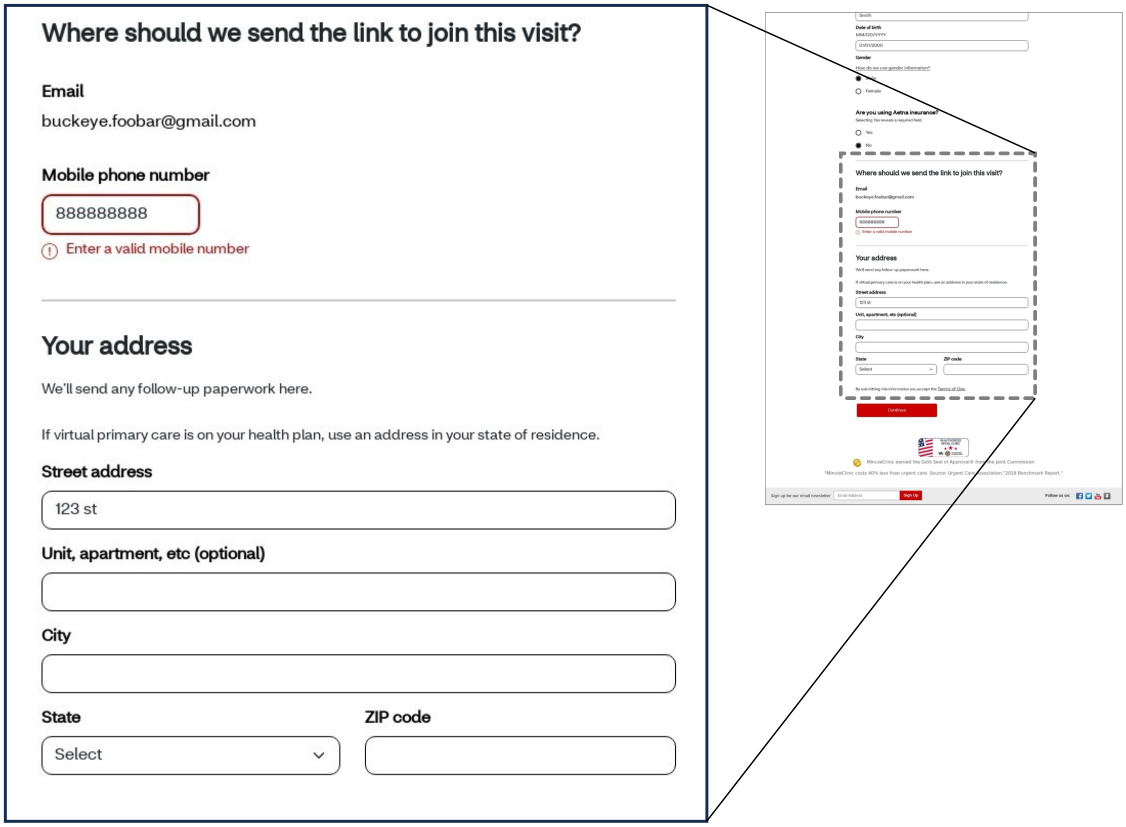

Where should we send the link to join this visit?

Email
buckeye.foobar@gmail.com

Mobile phone number
888888888
ⓘ Enter a valid mobile number

Your address

We'll send any follow-up paperwork here.

If virtual primary care is on your health plan, use an address in your state of residence.

Street address
123 st

Unit, apartment, etc (optional)

City

State
Select

ZIP code

GPT-4V: … Since the mobile phone number field is highlighted, it suggests there might be a problem with the entry—such as an incorrect format or not enough digits. **Therefore, the next action would be to correct the mobile phone number by clicking on the mobile phone number text box and entering a valid number.** Once that's corrected, the user should continue filling out the rest of the form with the city, state, and ZIP code before proceeding to click the "Continue" button to move forward with the signup process…

Figure 21: In this example, the webpage displays an error message indicating an invalid phone number, a consequence of prior actions. The model identifies this error and prioritizes its immediate rectification, foregoing the subsequent planned steps.