

---

# DigiRL: Training In-The-Wild Device-Control Agents with Autonomous Reinforcement Learning

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Pre-trained vision language models (VLMs), though powerful, typically lack training  
2 ing on decision-centric data, rendering them sub-optimal for decision-making tasks  
3 such as in-the-wild device control through Graphical User Interfaces (GUIs) when  
4 used off-the-shelf. While training with static demonstrations has shown some  
5 promise, we show that such methods fall short when controlling real GUIs due to  
6 their failure to deal with real world stochasticity and dynamism not captured in  
7 static observational data. This paper introduces a novel autonomous RL approach,  
8 called DigiRL, for training in-the-wild device control agents through fine-tuning a  
9 pre-trained VLM in two stages: offline and offline-to-online RL. We first build a  
10 scalable and parallelizable Android learning environment equipped with a VLM-  
11 based general-purpose evaluator and then identify the key design choices for simple  
12 and effective RL in this domain. We demonstrate the effectiveness of DigiRL using  
13 the Android-in-the-Wild (AitW) dataset, where our 1.5B VLM trained with RL  
14 achieves a 49.5% absolute improvement – from 17.7 to 67.2% success rate – over  
15 supervised fine-tuning with static human demonstration data. It is worth noting that  
16 such improvement is achieved without any additional supervision or demonstration  
17 data. These results significantly surpass not only the prior best agents, including  
18 AppAgent with GPT-4V (8.3% success rate) and the 17B CogAgent trained with  
19 AitW data (14.4%), but also our implementation of prior best autonomous RL  
20 approach based on filtered behavior cloning (57.8%), thereby establishing a new  
21 state-of-the-art for digital agents for in-the-wild device control.

## 22 1 Introduction

23 Advances in vision-language models (VLMs), especially in regards to their remarkable common-  
24 sense, reasoning, and generalization abilities imply that realizing a fully autonomous digital AI  
25 assistant, that can simplify human life by automating day-to-day activities on computer devices  
26 via natural language interfaces, is no longer a distant aspiration [16, 45, 55]. An effective device  
27 control AI assistant should be able to complete tasks in-the-wild through Graphical User Interfaces  
28 (GUIs) on digital devices: make travel plans; experiment with presentation designs; and operate a  
29 mobile device autonomously, all while running amidst stochasticity and distractors on the device, the  
30 Internet, and the tools it interacts with. However, enhanced reasoning or common-sense abilities do  
31 not directly transfer to intelligent assistant behavior: ultimately we want AI assistants to accomplish  
32 tasks, exhibit rational behavior, and recover from their mistakes as opposed to simply producing a  
33 plausible completion to a given observation based on the data seen during pre-training. This implies  
34 that a mechanism to channel abilities from pre-training into a deployable AI “agent” is lacking.

35 Even the strongest proprietary VLMs, such as GPT-4V [24] and Gemini 1.5 Pro [7], still struggle to  
36 produce the right actions when completing tasks on devices. While general-purpose vision-language  
37 abilities help these models still make meaningful abstract deductions about novel scenes when  
38 deployed, these deductions do not transfer to accurate reasoning for control [47, 45, 54, 44]. As a

39 result, most prior work for building device agents construct complex wrappers around proprietary  
 40 VLMs, combining them with prompting, search, or tool use [47, 44, 51, 50, 45]. While building  
 41 prompting or retrieval wrappers to improve decision-making performance of existing VLMs provides  
 42 a “stop-gap” solution in the short run, without updating the weights, the effectiveness of resulting  
 43 agents is inherently limited by the capabilities of the base model [49, 3]. For example, we found that  
 44 off-the-shelf VLMs make reasoning failures that derail the agent (e.g., Figure 2 and Figure 11), and  
 45 these are a direct consequence of the base model. A different solution is to fine-tune the model on  
 46 demonstrations via imitation learning. However, the dynamic nature of the web and device means  
 47 that models trained to mimic actions in stale data can result in sub-optimality as the eco-system  
 48 changes [26]. Additionally, agents trained in this way struggle to recover from out-of-distribution  
 49 states resulting from the agents’ own mistakes [8, 12].

50 If we can instead build an interactive  
 51 approach to *train* a VLM to directly  
 52 adapt and learn *from its own experi-*  
 53 *ence* on the device and the Internet,  
 54 that can be used to build a robust and  
 55 reliable device-control agent, without  
 56 needing wrappers on top of propri-  
 57 etary models. However, this learning-  
 58 based approach must satisfy some  
 59 desiderata. First, it must use online  
 60 interaction data since static demon-  
 61 stration data would not be represen-  
 62 tative of the task when the model is  
 63 deployed: for instance, even in the  
 64 setting of web navigation alone, dy-  
 65 namic nature of in-the-wild websites  
 66 means that the agent will frequently  
 67 encounter website versions that differ  
 68 significantly from the scenarios seen

69 during training and will need to behave reliably despite changes in visual appearance and distractions.  
 70 Second, learning on-the-fly means the approach must learn from multi-turn interaction data from  
 71 the model itself, a large of chunk of which would consist of failures. Proper mechanisms must be  
 72 designed to automatically pick out the correct actions while filtering the wrong ones.

73 To this end, **our main contribution** is a novel autonomous RL approach, DigiRL (i.e., RL for Digital  
 74 Agents), for training device control agents. The resulting agent attains state-of-the-art performance  
 75 on a number of Android device control tasks. To train this agent, our approach operates in two phases:  
 76 an initial offline RL phase to make maximal use of existing data, followed by an in-the-wild, offline-  
 77 to-online RL phase, that further fine-tunes the model obtained from offline RL on online rollout data.  
 78 Online RL training requires access to an environment that the agent can interact with and obtain  
 79 reliable reward signals, all in a reasonable amount of wall-clock time. To do so, we build a scalable  
 80 and parallelizable Android learning environment equipped with a robust VLM-based general-purpose  
 81 evaluator [26] (average error rate 2.8% against human judgement) that supports running up to 64  
 82 real Android emulators at the same time to make online RL real-time. Then, to effectively learn  
 83 autonomously, we develop an online RL approach that retains the simplicity of supervised learning,  
 84 but incorporates several key deep RL insights to enable fast fine-tuning. Concretely, our approach is  
 85 a variant of advantage-weighted regression (AWR) [28], equipped with: (i) an automatic curriculum  
 86 that uses a value function to order tasks so as to extract maximal learning signal, which is inspired by  
 87 prioritized replay methods [11, 32, 23], and (ii) a value-function trained via effective cross-entropy  
 88 loss [17, 5] to extract low-variance and less-biased gradient signal amidst stochasticity and diverse  
 89 tasks. This RL approach allows us to fine-tune VLMs to attain state-of-the-art after training on only  
 90 stale data, as well as sample-efficient learning with online data.

91 We evaluate our agent trained with DigiRL in carrying out diverse instructions from **Android in**  
 92 **the Wild dataset** [31] on real Android device emulators and find that our agent can achieve a  
 93 **49.5% improvement** over the existing state-of-the-art agents (from 17.7% to 67.2% success rate)  
 94 AutoUI [52] and CogAgent [9], and over 9% improvement over our implementation of the prior  
 95 best autonomous learning approach based on Filtered Behavior Cloning. The performance of our

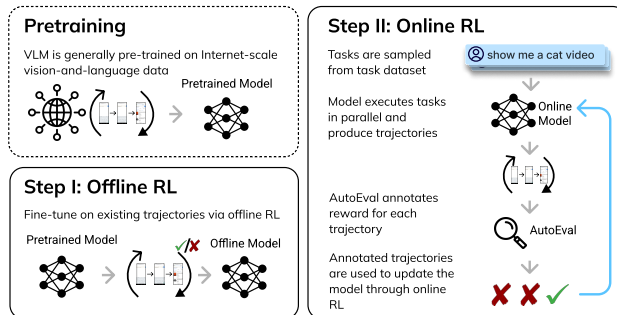


Figure 1: **DigiRL overview.** DigiRL is built upon a VLM that has been pre-trained on extensive web data to develop fundamental skills such as common knowledge, reasoning, and visual grounding. Initially, we employ offline RL to fine-tune the VLM using stale task-specific data, which helps in eliciting goal-oriented behaviors. Subsequently, our agent engages with real-world graphical user interfaces, continuously enhancing its performance through online RL and autonomous performance evaluations.

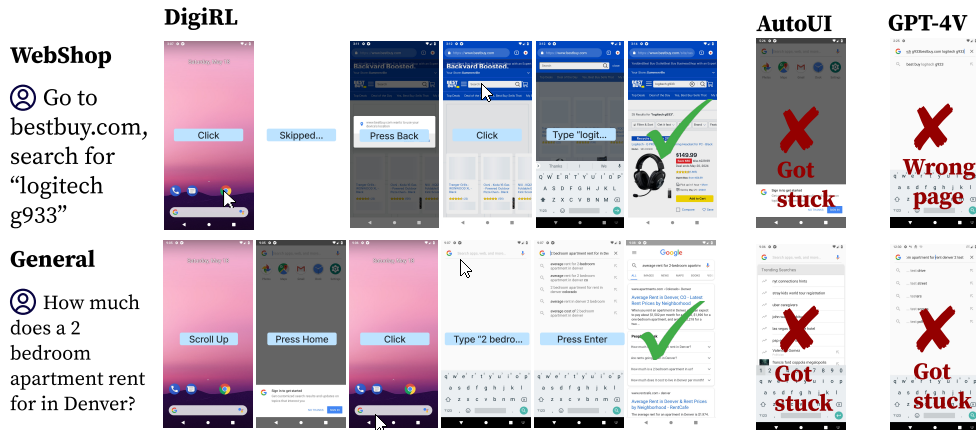


Figure 2: **Qualitative comparison between DigiRL and other approaches.** AutoUI trained from static human demonstrations can easily get stuck in out-of-distribution states while GPT-4V often get on a wrong goal (searched “logitech g933bestbuy.com logitech g933” in Google instead of bestbuy.com). In contrast, DigiRL can recover from such states and complete complex instruction as requested.

96 agent also significantly surpasses wrappers on top of state-of-the-art proprietary VLMs such as  
 97 GPT-4V [24] and Gemini 1.5 Pro [7] (17.7% success rate), despite using a significantly smaller model  
 98 (with 1.5B parameters). To our knowledge, *this is the first work to successfully build an autonomous*  
 99 *offline-to-online RL approach to enable state-of-the-art performance on device-control problems.*

## 100 2 Related works

101 **Multi-modal digital agents.** As opposed to language-only agents that largely interact with both  
 102 text or code inputs and outputs [33, 49, 3, 30, 46, 20, 13], training multi-modal agents capable  
 103 of controlling devices presents different challenges: first, device control is done directly at the  
 104 pixel-level and in a coordinate-based action space, instead of natural language [31, 44], and second,  
 105 the ecosystem of a device and the Internet tends to be quite stochastic and unpredictable, which is  
 106 absent with high-level planning in language only. To handle these challenges, prior work largely  
 107 builds on strong proprietary VLMs [24, 7], and designs complex rule-based wrappers [47, 50, 45, 51]  
 108 to enhance the visual grounding capabilities of VLMs in GUI interfaces and convert text output  
 109 into pixel interactions. However, without any form of fine-tuning, this limits the room for possible  
 110 performance improvement [44, 47, 49, 3], especially when pre-training corpora only present limited  
 111 action-labeled data. A separate line of work fine-tunes VLMs with demonstration data [19, 15, 9, 52]  
 112 via imitation learning, but myopically maximizing single-step action accuracy without accounting for  
 113 consequences of these actions in subsequent steps may lead to poor solutions amidst stochasticity [26],  
 114 as agents trained in such ways will struggle to recover from out-of-distribution states not included  
 115 in the demonstration data [8, 12]. The third category, and perhaps the closest to us, is works that  
 116 run filtered imitation learning on autonomously-collected data to directly maximize the episode  
 117 success rate [26, 18]. In contrast, **ours is the first work to run autonomous, offline-to-online**  
 118 **RL** for device control at scale, producing an agent that outperforms prior agents built via imitation.  
 119 Even when compared to prior work running on-policy RL simplistic in web navigation settings  
 120 (MiniWob++ [37, 10]), our approach is 1000x more sample efficient, at the full scale.

121 **Environments for device control agents.** Recent works have introduced simulated environments  
 122 for building device control agents [48, 55, 16, 53, 4, 44]. However, these environments are primarily  
 123 designed for evaluation, and present only a limited range of tasks within fully deterministic and  
 124 stationary settings, infeasible for acquiring a diverse repertoire of skills needed for device control.  
 125 Alternatively, other works use environments with a greater diversity of tasks [48, 37], but these  
 126 environments often oversimplify the task complexity, thus failing to transfer to in-the-wild settings.  
 127 Conversely, our training environment utilizes autonomous evaluation [26] with Gemini 1.5 Pro [7]  
 128 to support diverse, open-ended tasks on parallel *actual* Android devices, at full scale unlike prior  
 129 environments. This also contrasts other prior works that use single-threaded Android emulators [26,  
 130 39, 19] and thus inefficient for support online RL at scale.

131 **Reinforcement learning for LLM/VLMs.** The majority of prior research employing reinforcement  
 132 learning (RL) for foundation models concentrates on decision-making tasks that must be solved in  
 133 a single turn, such as preference optimization [25, 57, 2] or reasoning [27]. However, “myopically”  
 134 optimizing for single-turn interaction may result in sub-optimal strategies for multi-step problems [56,

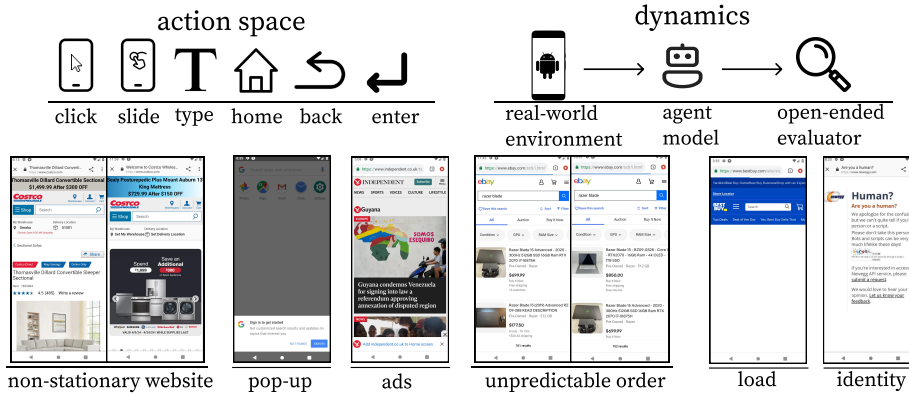


Figure 3: **Environment details.** *Top:* actions space and dynamics of the environment. *Bottom:* examples of the read-world non-stationarity and dynamism of the environment.

135 38, 42], especially admit a high degree of stochasticity. Therefore, we focus on building multi-  
 136 turn RL algorithms in this work. While prior work has developed value-based RL algorithms for  
 137 LLMs [42, 38, 1, 56], they typically require maintaining multiple models such as Q networks and  
 138 target value networks, and can be subjective to slow convergence and sensitivity to choices of hyper-  
 139 parameters. In contrast, we focus on identifying the key design choices for instantiating a simple yet  
 140 effective RL algorithm for practitioners to plug-and-play to substantially improve full-scale device  
 141 control, and our approach can serve as a base model for future research to build upon.

### 142 3 Problem setup and preliminaries

143 **Problem formulation.** We are interested in pixel-based interaction with virtual devices. We scope  
 144 our study in the control of Android devices: this is already significantly more challenging and more  
 145 general than previous learning-based environments that focus solely on web navigation [16, 55, 4],  
 146 where the web browser itself is merely one application within our broader environment, and link-based  
 147 device controls [47, 50] are inadequate for tasks like games that do not support link inputs.

148 Each episode begins with the emulator initialized to the home screen. Subsequently, a task is selected  
 149 from a predefined set of language instructions, some examples of which are shown in Appendix A.1.  
 150 An agent is then tasked with manipulating the emulator to fulfill this instruction. At each time step,  
 151 the agent receives a screenshot of the current screen as the observation. Following the action space  
 152 in prior literature [31], the available actions include tapping and sliding based on normalized  $(x, y)$   
 153 coordinates (ranging from 0 to 1 relative to the screen dimensions), typing text strings of variable  
 154 length, and pressing special buttons such as HOME, BACK, and ENTER, as illustrated in Figure 3.  
 155 Our train and test instructions comes from General and Web Shopping subsets in AitW [31]. These  
 156 tasks consist of information-gathering tasks like “What’s on the menu of In-n-Out?”, and shopping  
 157 tasks on the web like “Go to newegg.com, search for razer kraken, and select the first entry”.

158 **Challenges of stochasticity.** Real-world device control presents unique challenges of stochasticity  
 159 absent in simulated environments [55, 37] such as: (1) the dynamic nature of websites and applications,  
 160 which undergo frequent updates, causing the online observations to be different from stale offline data,  
 161 (2) various unpredictable distractors such as pop-up advertisements, login requests, and the stochastic  
 162 order of search results. (3) technical challenges and glitches such as incomplete webpage loading or  
 163 temporary access restrictions to certain sites. Examples of scenarios with such stochasticity from  
 164 our experiments are shown in Figure 3. We observe that these stochastic elements pose significant  
 165 challenges for pre-trained VLMs, including even those fine-tuned on device control data.

166 **Setup for reliable and scalable online RL.** As autonomous RL interleaves data collection and  
 167 training, to maximize learning amidst stochasticity, it is crucial to have a real-time data collection  
 168 pipeline to collect enough experience for gradient updates. While this is not possible in single-thread  
 169 Android emulator environments [26, 39] due to latency, we parallelize our Android emulator using  
 170 appropriate error handling as discussed in Appendix A.1. In addition, the environment must provide  
 171 a reward signal by judging whether the current observation indicates the agent has successfully  
 172 completed the task. To generalize our *evaluator* to support a wide range of tasks, we extend Pan  
 173 et al. [26]’s end-to-end autonomous evaluator that does not require accessing the internal states of the  
 174 emulator or human-written rules for each task. This contrasts previous works that manually write  
 175 execution functions to verify the functional completeness of each task [16, 48, 37, 44]. We adopt

176 Gemini 1.5 Pro [6, 7] as the backbone of the autonomous evaluator. We seed this model with few-shot  
 177 rollouts and the associated human-labeled success indicators to guide evaluation of novel queries.  
 178 This pipeline enables a single evaluator that can evaluate all AiTW tasks. The evaluator is highly  
 179 aligned with human annotations (average error rate 2.8%), validated in Figure 6.

## 180 4 DigiRL: autonomous RL for building a strong device control agent

181 We now present our autonomous RL framework for training device agents. We pose the device  
 182 control problem as a partially-observed Markov decision process (POMDP) and develop RL methods  
 183 for this POMDP. The core of our approach is based on a simple and scalable off-policy RL method,  
 184 advantage-weighted regression (AWR) [29], but we make crucial modifications to handle stochasticity  
 185 and highly-variable task difficulty, through the use of value functions trained with appropriate losses,  
 186 and an automatic curriculum, induced by an instruction-level value function to maximize learning.

187 **Device control and GUI navigation as a POMDP.** Device control is inherently a partially-observed  
 188 problem: there is often some hidden state information that is not observable within the current  
 189 screenshot (e.g., a background process running on the device, listings of other items on a webpage  
 190 that are important for decision-making but not visible together on one screen). These device control  
 191 agents should resolve their uncertainty pertaining to the task, and only then commit to an action. In  
 192 order to get this kind of behavior automatically from RL training, we conceptualize device control  
 193 guided by natural language instructions as a finite horizon Partially Observable Markov Decision  
 194 Process (POMDP) represented by  $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mu_0, \mathcal{R}, H\}$  and run policy gradient to solve  
 195 this POMDP. At the beginning, an initial state  $s_0$  and a natural language instruction  $c$  are sampled  
 196 from the initial state distribution  $\mu_0$ . A reward of 1 is given at the end if the agent successfully fulfills  
 197 the task per the evaluator, otherwise a reward of 0 is given. The sequence terminates either when the  
 198 agent accomplishes the task or when the maximum allowed number of interactions  $H$  is exceeded. In  
 199 principle, we should treat the state as the history of past observations to solve a POMDP correctly,  
 200 but in our experiments we find it enough to use the last two scenes as states.

201 **Definitions & notation.** To explain our approach in detail, we include several standard definitions  
 202 used in reinforcement learning (RL). The Q function for a policy  $\pi$  represents the expected long-  
 203 term return from taking a specific action at the current step and then following policy  $\pi$  thereafter:  
 204  $Q^\pi(s_h, a_h, c) = \mathbb{E}_\pi \left[ \sum_{t=h}^H r(s_t, a_t, c) \right]$ . The value function  $V^\pi(s_h, c)$  is calculated by averaging  
 205 the Q-value,  $Q^\pi(s_h, a_h, c)$ , over actions  $a_h$  drawn from the policy  $\pi$ . The advantage  $A^\pi(s_h, a_h, c)$   
 206 for a state-action pair is computed by subtracting the state’s value under the policy from its Q-value:  
 207  $A^\pi(s_h, a_h, c) = Q^\pi(s_h, a_h, c) - V^\pi(s_h, c)$ .

### 208 4.1 Backbone of our approach: off-policy RL via advantage-weighted regression

209 A starting point for our approach is the advantage-weighted regression (AWR) algorithm [29],  
 210 which says that we can improve the policy reliably by regressing the policy towards exponentiated  
 211 advantages induced by the reward function, as a proxy for optimizing the policy gradient while  
 212 staying close to the previous policy [14, 35, 34]:

$$\arg \max_{\pi} \mathbb{E}_{\nu} [\log \pi(a|s, c) \cdot \exp(A(s, a, c)/\beta)], \quad (4.1)$$

213 for some positive parameter  $\beta$  and the distribution of past experience  $\nu$ , and  $A(s, a, c)$  denotes the  
 214 advantage of a state-action pair  $(s, a)$  given a context  $c$ . To avoid tuning the hyperparameter  $\beta$ , we  
 215 consider an alternative that does “hard filtering” on the advantages instead of computing  $\exp(A)$ ,  
 216 similar to prior works [22, 43]. This leads to the following loss function for fine-tuning the model:

$$\mathcal{L}(\pi) = -\mathbb{E}_{\text{filter}(\nu)} [\log \pi(a|s, c)]. \quad (4.2)$$

217 Typically, these advantages are computed by running Monte-Carlo (MC) rollouts in the environment  
 218 to estimate the value of a given state-action pair, and subtracting from it an estimate of the value  
 219 of the state alone given by a learned value estimator. However, this approach is likely to produce  
 220 high-variance advantages given the stochasticity of the device eco-system that affects MC rollouts.

### 221 4.2 Obtaining reliable advantage estimates from doubly-robust estimators

222 To reliably identify *advantageous* actions given significant environment stochasticity, we construct a  
 223 per-step advantage estimator, inspired by doubly-robust estimators [40, 36]:

$$A^{\text{step}}(s_h, a_h, c) := \lambda^{H-h} r(s_H, a_H, c) + V^{\text{step}}(s_{h+1}, c) + r(s_h, a_h, c) - V^{\text{step}}(s_h, c), \quad (4.3)$$

224 where  $\lambda$  is a weighting hyper-parameter. This construction of the advantage estimator is a simplified  
 225 version of Generalized Advantage Estimation (GAE) [36], and balances an advantage estimator with  
 226 higher variance Monte-Carlo estimates  $\lambda^{H-h}r(s_H, a_H, c)$  (due to stochasticity) and an estimator  
 227 with higher bias  $V^{\text{step}}(s_{h+1}, c) + r(s_h, a_h, c) - V^{\text{step}}(s_h, c)$  (due to imperfect fitting of the value  
 228 function). We observed that combining both high-variance and high-bias estimators gave us a sweet-  
 229 spot in terms of performance. To implement the step-level hard filtering, we simply threshold this  
 230 doubly robust estimator as  $A^{\text{step}}(s_h, a_h, c) > 1/H$  to decide which actions progress towards the goal.

### 231 4.3 Automatic curriculum using an instruction-level value function

232 While the AWR update (Equation 4.1) coupled with a robust advantage estimator (Equation 4.3) is  
 233 likely sufficient on standard RL tasks, we did not find it to be effective enough for device control  
 234 in preliminary experiments. Often this was the case because the task set presents tasks with high-  
 235 variable difficulties that collecting more data on tasks that the agent was already proficient at affected  
 236 sample efficiency negatively. In contrast, maximal learning signal can be derived by experiencing the  
 237 most informative tasks for the agent during training. To this end, we design an instruction-level value  
 238 function  $V^{\text{instruct}}(c)$  to evaluate if a given rollout can provide an effective learning signal:

$$A^{\text{instruct}}(s_h, a_h, c) := \sum_{t=h}^H r(s_t, a_t, c) - V^{\text{instruct}}(c) = r(s_H, a_H, c) - V^{\text{instruct}}(c), \quad (4.4)$$

239 where  $\sum_{t=h}^H r(s_t, a_t, c)$  is a Monte-Carlo estimator of  $Q(s_h, a_h, c)$ . The equality holds because the  
 240 POMDP formulation only provides rewards at the end of a rollout. Intuitively, if a rollout attains a  
 241 high value of  $A^{\text{instruct}}(s_h, a_h, c)$ , it means the value function  $V^{\text{instruct}}$  is small. Therefore, this rollout  
 242 represents a valuable experience of the agent accomplishing a difficult task, and thus should be  
 243 prioritized, akin to ideas pertaining to prioritized experience [32] or level replay [11]. When training  
 244 the actor with a buffer of historical off-policy data, we first perform a filtering step to identify the  
 245 top- $p$  datapoints with highest  $A^{\text{instruct}}(s_h, a_h, c)$ . Then, we use it for AWR (Equation 4.1) with the  
 246 doubly-robust advantage estimator (Equation 4.3).

247 **Implementation details.** Inspired by the findings in some recent works [5, 17] that modern deep  
 248 learning architectures like transformers [41] are better trained with cross-entropy losses instead of  
 249 mean-squared losses, we utilize a cross-entropy objective based on the Monte-Carlo estimate of the  
 250 trajectory reward for training both of our value functions:

$$\begin{aligned} \mathcal{L}(V^{\text{traj}}) &= -\mathbb{E}_\nu[r(s_H, a_H, c) \log V^{\text{traj}}(c) + (1 - r(s_H, a_H, c)) \log(1 - V^{\text{traj}}(c))] \\ \mathcal{L}(V^{\text{step}}) &= -\mathbb{E}_\nu[r(s_H, a_H, c) \log V^{\text{step}}(s_h, a_h, c) + (1 - r(s_H, a_H, c)) \log(1 - V^{\text{step}}(s_h, a_h, c))] \end{aligned}$$

## 251 5 Experimental evaluation

252 The goal of our experiments is to evaluate the performance of DigiRL on challenging Android device  
 253 control problems. Specifically, we are interested in understanding if DigiRL can produce agents that  
 254 can effectively learn from autonomous interaction, while still being able to utilize offline data for  
 255 learning. To this end, we perform a comparative analysis of DigiRL against several prior approaches,  
 256 including state-of-the-art agents in Section 5.1. We also perform several ablation experiments to  
 257 understand the necessity and sufficiency of various components of our approach in Section 5.2.

258 **Baselines and comparisons.** We compare DigiRL with: (a) state-of-the-art agents built around  
 259 proprietary VLMs, with the use of several prompting and retrieval-style techniques; (b) running  
 260 imitation learning on static human demonstrations with the same instruction distribution, and (c) a  
 261 filtered BC approach [26]. For proprietary VLMs, we evaluate **GPT-4V** [24] and **Gemini 1.5 Pro** [7]  
 262 both zero-shot and when augmented with carefully-designed prompts. For the zero-shot setting, we  
 263 use the prompt from Yang et al. [47] and augment the observation with Set-of-Marks [54]. Set-of-  
 264 Marks overlays a number for each interactable element over the screenshot, so that a VLM can directly  
 265 output the number of the element to interact with in plain text instead of attempting to calculate pixel  
 266 coordinates, which is typically significantly harder. We also compare with AppAgent [47], which first  
 267 prompts the VLM to explore the environment, and appends the experience collected to the test-time  
 268 prompt. We also compare with two state-of-the-art fine-tuning methods for Android device control:  
 269 **AutoUI** (specifically AutoUI-Base [52]) and **CogAgent** [9]. AutoUI-Base uses an LM with 200M  
 270 parameters, and a vision encoder with 1.1B parameters. CogAgent has 11B parameters for its vision  
 271 encoder and 7B for its LM. The supervised training corpus for both AutoUI-Base and CogAgent  
 272 contains AitW, including the instruction set and the emulator configuration we use.

			AitW General		AitW Web Shopping	
			Train	Test	Train	Test
<b>Prompting</b>	SET-OF-MARKS	GPT-4V	5.2	13.5	3.1	8.3
		Gemini 1.5 Pro	32.3	16.7	6.3	11.5
	APPAGENT	GPT-4V	13.5	17.7	12.5	8.3
		Gemini 1.5 Pro	14.6	16.7	5.2	8.3
<b>Learning</b>	SUPERVISED	CogAgent	7.8	7.8	8.6	14.4
	TRAINING	AutoUI	12.5	14.6	14.6	17.7
	OFFLINE	Filtered BC	51.7 ± 5.4	50.7 ± 1.8	44.7 ± 1.6	45.8 ± 0.9
		<b>Ours</b>	46.9 ± 5.6	62.8 ± 1.0	39.3 ± 6.0	45.8 ± 6.6
	OFF-TO-ON	Filtered BC	53.5 ± 0.8	61.5 ± 1.1	53.6 ± 4.7	57.8 ± 2.6
		<b>Ours</b>	<b>63.5 ± 0.0</b>	<b>71.9 ± 1.1</b>	<b>68.2 ± 6.8</b>	<b>67.2 ± 1.5</b>

Table 1: **Main comparisons of different agents across various settings.** Each offline experiment is repeated three times and the mean and standard deviation are reported. Each online experiment is repeated two times. Results are evaluated with our autonomous evaluator with the first 96 instructions in the train and test set. Correlation of our correlation and human judgements can be found in Figure 6.

273 **Base VLM and offline dataset.** Both **Filtered BC** and **DigiRL** use trained AutoUI-Base checkpoints  
274 with the image encoder frozen. The instruction and step-level value functions for DigiRL employ  
275 this same frozen image encoder. The visual features output from the encoder are concatenated with  
276 instruction features derived from RoBERTa [21]. A two-layer MLP is then used to predict the value  
277 function. In the offline phase, the offline dataset is collected by rolling out the initial AutoUI-Base  
278 supervised trained checkpoint as policy. For fair comparisons, we keep the number of offline data  
279 collected in the pure offline training roughly the same as the total number of data collected in the  
280 offline-to-online training. Due to the dynamic nature of the Internet-device eco-system, our offline  
281 data was stale by the time we were able to run our offline-to-online experiments, and this presented  
282 additional challenge in offline-to-online learning. In both General and Web Shopping subsets, offline  
283 experiments make use of around 1500 trajectories while offline-to-online experiments start with  
284 around 500 offline trajectories and update with another 1000 online trajectories. In the offline phase,  
285 DigiRL skips instruction-level filtering and instead trains the actor with all successful trajectories to  
286 make full use of the offline data. See a detailed breakdown of our dataset in Appendix A.1.

## 287 5.1 Main results

288 Our main results are summarized in Table 1 and Figure 4.  
289 we find that in both AitW General and AitW Web Shopping  
290 subsets, our agent trained via DigiRL significantly out-  
291 performs prior state-of-the-art methods based on prompt-  
292 ing and retrieval (AppAgent + GPT-4V/Gemini 1.5 Pro) or  
293 training on static demonstrations (CogAgent and AutoUI),  
294 by a large margin with more than **49.5% absolute improve-**  
295 **ment** (from 17.7% to 71.9% on  
296 the General subset and from 17.7% to 67.2% on the Web Shopping subset). Notably, this improve-  
297 ment from DigiRL is realized *fully autonomously without making use of human supervision* (e.g.  
298 manually labeled demonstrations or hand-written verifiers).  
299

300 **Are inference-time prompting and retrieval techniques or supervised training enough for**  
301 **device control?** Delving into Table 1, we observe that off-the-shelf proprietary VLMs, even when  
302 supplemented with the set-of-marks mechanism, do not attain satisfactory performance: both GPT-4V  
303 and Gemini 1.5 Pro achieve success rates under 20%. One possible cause could be the under-  
304 representation of Android device data in the pre-training data. Moreover, inference-time adaptation  
305 strategies such as AppAgent [47] show minimal improvement, with gains not exceeding 5% for either





Figure 4: **Offline-to-online training curves for Filtered BC and DigiRL.** Curves are smoothed with exponential weighted averaging to start from the performance of supervised trained policy. Two runs for each model are started on two different dates with at least two days apart. Observe that DigiRL is able to improve faster with a fewer number of samples. Since the data collection frequency is the bottleneck, these performance trends directly reflect performance trends against wall-clock time as well.

312 model, suggesting a limited scope for improvement without fine-tuning of some sort. As illustrated in  
 313 Figure 5, the primary failures of these VLMs stem from hallucinatory reasoning that lead the VLMs to  
 314 land on a relevant but wrong page. This suggests that while state-of-the-art VLMs excel at high-level  
 315 reasoning in code or math problems, their reliability of reasoning in less familiar domains, such as  
 316 device control, remains inadequate. For example, for the instruction “Go to newegg.com, search for  
 317 ‘alienware area 51’, and select the first entry”, a GPT-4V based agent erroneously searched “alien  
 318 area 51 ebay” in Google.com and decided that it had made progress towards the task (Figure 11).

319 Training on domain-specific human  
 320 demonstrations, however, does boost  
 321 performance, allowing the smaller,  
 322 specialized VLM, AutoUI, to match  
 323 or surpass the larger, generalist VLMs  
 324 like GPT-4V and Gemini 1.5 Pro.  
 325 Nonetheless, this supervised imitation  
 326 learning approach still fall short, with  
 327 success rates on both subsets remain-  
 328 ing below 20%. This shortcoming is  
 329 not addressed via enhancements in  
 330 model scale or architecture, as evi-  
 331 denced by CogAgent [9], with 17 bil-  
 332 lion parameters still achieving similar  
 333 performance to AutoUI [52], which  
 334 has only 1.5 billion parameters. As  
 335 depicted in Figure 5, a predominant failure  
 336 mode for these agents is an inability to  
 337 rectify their own errors. An example  
 338 trajectory that we observed is that for the  
 instruction “what’s on the menu of In-n-Out”, the agent accidentally activated the voice input button, and failed to quit that page until the step limit. In contrast, DigiRL is able to recover from the errors more efficiently( Appendix B.2).



330 Figure 5: **Failure modes for each approach** on both the AiTW  
 331 General and Web Shopping subsets. We found that the failure  
 332 mode RL training is most effective at reducing compared to model  
 333 supervised trained on human data is “Fail to recover from mistakes”.  
 334 A more fine-grained decomposition can be found in Appendix C.

339 **Comparison of different RL approaches.** In Table 1 and Figure 4, we present a comparative analysis  
 340 of various RL approaches. Notably, both offline and offline-to-online configurations demonstrate  
 341 that our RL approach, when augmented with a continuous stream of autonomous interaction data  
 342 and reward feedback, substantially improves performance. This improvement is evident from an  
 343 increase in the success rate from under 20% to over 40%, as the agent learns to adapt to stochastic  
 344 and non-stationary device interfaces. Moreover, although the total sample sizes for offline and offline-  
 345 to-online settings are equivalent, the top-performing offline-to-online algorithm markedly surpasses  
 346 its offline counterpart (75% versus 62.8% on the General subset). This highlights the critical role and  
 347 efficacy of autonomous environment interaction, and establishes the efficacy of DigiRL in learning  
 348 from such uncurated, sub-optimal data. Lastly, DigiRL consistently outperforms the state-of-the-art  
 349 alternative, Filtered BC, across both the General and Web Shopping subsets, improving from 61.5%  
 350 to 71.9% and 57.8% to 61.4%, respectively, highlighting DigiRL’s performance and efficiency.

## 351 5.2 Analysis and ablations

352 **Failure modes analysis.** We conduct an additional user study to annotate the failure modes for each  
 353 agent as shown in Figure 5, and a more fine-grained breakdown can be found in Appendix C. At a  
 354 high level, we classify the major failure modes of all agents into the following three categories: (1)  
 355 **Failure to recover from mistakes** refers to the scenario where the agent made a mistake that led it to  
 356 states from which it failed to quickly recover and resume the task, such as a wrong search page. (2)  
 357 **Getting stuck midway** refers to the failure mode where the agent gets distracted on the right track to  
 358 completing the instruction and as a result fails to accomplish the task. For example, failing to click on  
 359 the right link or failing to search after typing the key words. (3) **Arriving at wrong goal** refers to the  
 360 failure mode where the agent arrives at a wrong page and mistakenly thinks that it had completed the  
 361 task. For e.g, the agent finds a macbook on costco.com instead of finding a macbook on ebay.com.

362 While all the types of failure modes benefit from offline and offline-to-online RL training as shown  
 363 in Figure 5, the most consistent and significant reduction is probably for the failure mode of failing  
 364 to recover from mistakes. This is because while pre-trained models, generating plausible future  
 365 tokens, can get distracted by the dynamic nature of the environment and, as a result, encounter at  
 366 never-before-seen states. With no clue of how to escape such states, these methods are unable to  
 367 recover and fail to solve the task. In contrast, by training on autonomously-collected rollouts, our  
 368 agent DigiRL is able to learn from its own mistakes and reduces failures to recover over training.






Figure 6: Correlation between our autonomous evaluator and human judgements for all policy models on General and Web Shopping subsets. For repeated offline and online runs, we report the correlation results for the run with the highest autonomous evaluation success rate.

369 **Ablation study of each component in DigiRL.** We conduct an ablation study on different components  
 370 of DigiRL in Figure 7 (right). We find that all the components used by our approach are necessary: (1)  
 371 using cross-entropy for training the value functions boosts performance by around 12% (compare Ours  
 372 and Ours w/ Regression); (2) using step-level advantages improves efficiency by 12% (comparing  
 373 Ours and Ours w/o step-level advantage); (3) the use of automatic curriculum improves the speed  
 374 of learning by around 25% (comparing Ours w/o step-level advantage and Filtered BC); (4) Ours  
 375 outperforms vanilla AWR that does not employ a doubly-robust advantage estimator or curriculum.

376 Additionally, we also observe no degradation in per-  
 377 formance as a result of “hard-filtering”, as show by  
 378 nearly comparable performance of our approach and  
 379 the best run of exponential filtering obtained via an  
 380 extensive tuning of the temperature hyperparame-  
 381 ter  $\tau$  in naïve AWR (comparing Ours and Ours w/  
 382 vanilla AWR reweighting), despite simplicity of im-  
 383 plementation in the hard filtering approach. Putting  
 384 together, these choices result in a new state-of-the-  
 385 art RL approach for device control.

386 **Evaluation of our autonomous evaluator.** In Fig-  
 387 ure 6, we present the findings from a user study  
 388 aimed at assessing the accuracy of our autonomous  
 389 evaluator. Our results indicate that the success rates  
 390 reported by our automatic evaluator are remarkably  
 391 consistent with those assessed by human evaluators  
 392 across almost all models, with differences less than 3%. Furthermore, we observed that evaluations on  
 393 the Web Shopping subset are more precise compared to those on the General subset. This increased  
 394 accuracy likely stems from the fact that tasks in the General subset are formulated in free-form  
 395 language, which can introduce ambiguity, whereas the Web Shopping subset features a narrower  
 396 range of language expressions, reducing potential variability.

## 397 6 Discussion, limitations, and broader impact

398 In this paper, we propose a novel autonomous RL approach, DigiRL, for training in-the-wild, multi-  
 399 modal, device-control agents that establish a new state-of-the-art performance on a number of Android  
 400 control tasks from Android-in-the-Wild dataset [31]. To achieve this, we first build a scalable and  
 401 parallelizable Android environment with a robust VLM-based general-purpose evaluator that supports  
 402 fast online data collection. We then develop a system for offline RL pre-training, followed by  
 403 autonomous RL fine-tuning to learn via interaction, amidst the stochasticity of the real-world Internet  
 404 and device eco-system. Our agent achieves a 280% improvement over the previous state-of-the-art  
 405 agents (from 17.7% to 68.2% in terms of task success rate), including AppAgent based on GPT-4V  
 406 and Gemini 1.5 Pro, and supervised trained models such as AutoUI and CogAgent.

407 Due to computational limitations, despite the fact that the parallel emulator and autonomous evaluator  
 408 can be easily extended to complicated tasks, our agent is trained only with tasks from AitW instead  
 409 of a generalist on device. Our design of the DigiRL algorithm aims for maximal implementation  
 410 simplicity, so we hope that our approach to serve as a base algorithm for future research to build  
 411 on. While our focus is on algorithmic framework, device-control agents would significantly impact  
 412 economy, society, and privacy due to data security and shared autonomy risks. Addressing these  
 413 concerns is crucial when integrating these agents, but this weakness is not specific to our approach.

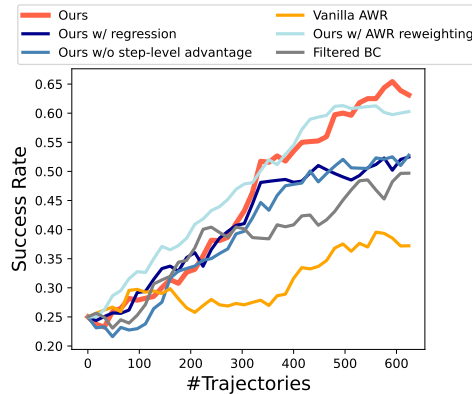


Figure 7: Ablation study results on the AitW Web Shopping subset.

## 414 References

- 415 [1] Marwa Abdulhai, Isadora White, Charlie Snell, Charles Sun, Joey Hong, Yuexiang Zhai, Kelvin  
416 Xu, and Sergey Levine. Lmrl gym: Benchmarks for multi-turn reinforcement learning with  
417 language models, 2023.
- 418 [2] Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier  
419 Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, Tony Wang, Samuel  
420 Marks, Charbel-Raphaël Segerie, Micah Carroll, Andi Peng, Phillip Christoffersen, Mehul  
421 Damani, Stewart Slocum, Usman Anwar, Anand Siththaranjan, Max Nadeau, Eric J. Michaud,  
422 Jacob Pfau, Dmitrii Krasheninnikov, Xin Chen, Lauro Langosco, Peter Hase, Erdem Bıyık,  
423 Anca Dragan, David Krueger, Dorsa Sadigh, and Dylan Hadfield-Menell. Open problems and  
424 fundamental limitations of reinforcement learning from human feedback, 2023.
- 425 [3] Baian Chen, Chang Shu, Ehsan Shareghi, Nigel Collier, Karthik Narasimhan, and Shunyu  
426 Yao. Fireact: Toward language agent fine-tuning. *ArXiv*, abs/2310.05915, 2023. URL <https://api.semanticscholar.org/CorpusID:263829338>.  
427
- 428 [4] Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam H. Laradji, Manuel Del Verme, Tom  
429 Marty, Léo Boisvert, Megh Thakkar, Quentin Cappart, David Vazquez, Nicolas Chapados, and  
430 Alexandre Lacoste. Workarena: How capable are web agents at solving common knowledge  
431 work tasks?, 2024.
- 432 [5] Jesse Farebrother, Jordi Orbay, Quan Vuong, Adrien Ali Taïga, Yevgen Chebotar, Ted Xiao,  
433 Alex Irpan, Sergey Levine, Pablo Samuel Castro, Aleksandra Faust, Aviral Kumar, and Rishabh  
434 Agarwal. Stop regressing: Training value functions via classification for scalable deep rl, 2024.
- 435 [6] 2023 Gemini Team. Gemini: A family of highly capable multimodal models, 2024.
- 436 [7] 2024 Gemini Team. Gemini 1.5: Unlocking multimodal understanding across millions of tokens  
437 of context, 2024.
- 438 [8] Dibya Ghosh, Jad Rahme, Aviral Kumar, Amy Zhang, Ryan P Adams, and Sergey Levine.  
439 Why Generalization in RL is Difficult: Epistemic POMDPs and Implicit Partial Observability.  
440 *NeurIPS*, 2021.
- 441 [9] Wenyi Hong, Weihang Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang,  
442 Zihan Wang, Yuxuan Zhang, Juanzi Li, Bin Xu, Yuxiao Dong, Ming Ding, and Jie Tang.  
443 Cogagent: A visual language model for gui agents, 2023.
- 444 [10] Peter C Humphreys, David Raposo, Toby Pohlen, Gregory Thornton, Rachita Chhaparia, Alistair  
445 Muldal, Josh Abramson, Petko Georgiev, Alex Goldin, Adam Santoro, and Timothy Lillicrap.  
446 A data-driven approach for learning to control computers, 2022.
- 447 [11] Minqi Jiang, Edward Grefenstette, and Tim Rocktäschel. Prioritized level replay. *CoRR*,  
448 abs/2010.03934, 2020. URL <https://arxiv.org/abs/2010.03934>.
- 449 [12] Yiding Jiang, J Zico Kolter, and Roberta Raileanu. On the importance of exploration for  
450 generalization in reinforcement learning. *Advances in Neural Information Processing Systems*,  
451 36, 2024.
- 452 [13] Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and  
453 Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues?, 2024.
- 454 [14] Sham M. Kakade and John Langford. Approximately optimal approximate reinforcement  
455 learning. In *International Conference on Machine Learning*, 2002. URL <https://api.semanticscholar.org/CorpusID:31442909>.  
456
- 457 [15] Raghav Kapoor, Yash Parag Butala, Melisa Russak, Jing Yu Koh, Kiran Kamble, Waseem  
458 Alshikh, and Ruslan Salakhutdinov. Omniact: A dataset and benchmark for enabling multimodal  
459 generalist autonomous agents for desktop and web, 2024.

- 460 [16] Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang,  
461 Graham Neubig, Shuyan Zhou, Ruslan Salakhutdinov, and Daniel Fried. Visualwebarena:  
462 Evaluating multimodal agents on realistic visual web tasks. *arXiv preprint arXiv:2401.13649*,  
463 2024.
- 464 [17] Aviral Kumar, Rishabh Agarwal, Xinyang Geng, George Tucker, and Sergey Levine. Offline  
465 q-learning on diverse multi-task data both scales and generalizes, 2023.
- 466 [18] Hanyu Lai, Xiao Liu, Iat Long Long, Shuntian Yao, Yuxuan Chen, Pengbo Shen, Hao Yu,  
467 Hanchen Zhang, Xiaohan Zhang, Yuxiao Dong, and Jie Tang. Autowebglm: Bootstrap and  
468 reinforce a large language model-based web navigating agent, 2024.
- 469 [19] Juyong Lee, Taywon Min, Minyong An, Changyeon Kim, and Kimin Lee. Benchmarking  
470 mobile device control agents across diverse configurations, 2024.
- 471 [20] Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding,  
472 Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui  
473 Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie  
474 Tang. Agentbench: Evaluating llms as agents, 2023.
- 475 [21] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy,  
476 Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT  
477 pretraining approach. *CoRR*, abs/1907.11692, 2019. URL [http://arxiv.org/abs/1907.](http://arxiv.org/abs/1907.11692)  
478 [11692](http://arxiv.org/abs/1907.11692).
- 479 [22] Ashvin Nair, Murtaza Dalal, Abhishek Gupta, and Sergey Levine. Accelerating online re-  
480 inforcement learning with offline datasets. *CoRR*, abs/2006.09359, 2020. URL [https:](https://arxiv.org/abs/2006.09359)  
481 [//arxiv.org/abs/2006.09359](https://arxiv.org/abs/2006.09359).
- 482 [23] OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew,  
483 Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider,  
484 Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba,  
485 and Lei Zhang. Solving rubik’s cube with a robot hand, 2019.
- 486 [24] 2023 OpenAI Team. Gpt-4 technical report, 2023.
- 487 [25] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin,  
488 Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton,  
489 Fraser Kelton, Luke E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis  
490 Christiano, Jan Leike, and Ryan J. Lowe. Training language models to follow instructions with  
491 human feedback. *ArXiv*, abs/2203.02155, 2022. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:246426909)  
492 [CorpusID:246426909](https://api.semanticscholar.org/CorpusID:246426909).
- 493 [26] Jiayi Pan, Yichi Zhang, Nicholas Tomlin, Yifei Zhou, Sergey Levine, and Alane Suhr. Au-  
494 tonomous evaluation and refinement of digital agents. *arXiv preprint arXiv:2404.06474*, 2024.
- 495 [27] Richard Yuanzhe Pang, Weizhe Yuan, Kyunghyun Cho, He He, Sainbayar Sukhbaatar, and  
496 Jason Weston. Iterative reasoning preference optimization, 2024.
- 497 [28] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression:  
498 Simple and scalable off-policy reinforcement learning. *CoRR*, abs/1910.00177, 2019. URL  
499 <http://arxiv.org/abs/1910.00177>.
- 500 [29] Xue Bin Peng, Aviral Kumar, Grace Zhang, and Sergey Levine. Advantage-weighted regression:  
501 Simple and scalable off-policy reinforcement learning, 2019.
- 502 [30] Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong,  
503 Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou,  
504 Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. Toolllm: Facilitating large language  
505 models to master 16000+ real-world apis, 2023.
- 506 [31] Christopher Rawles, Alice Li, Daniel Rodriguez, Oriana Riva, and Timothy Lillicrap. Android  
507 in the wild: A large-scale dataset for android device control. *arXiv preprint arXiv:2307.10088*,  
508 2023.

- 509 [32] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized experience replay,  
510 2016.
- 511 [33] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettle-  
512 moyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach  
513 themselves to use tools, 2023.
- 514 [34] John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. Trust  
515 region policy optimization. *CoRR*, abs/1502.05477, 2015. URL [http://arxiv.org/abs/  
516 1502.05477](http://arxiv.org/abs/1502.05477).
- 517 [35] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal  
518 policy optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL [http://arxiv.org/abs/  
519 1707.06347](http://arxiv.org/abs/1707.06347).
- 520 [36] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-  
521 dimensional continuous control using generalized advantage estimation, 2018.
- 522 [37] Tianlin Shi, Andrej Karpathy, Linxi Fan, Jonathan Hernandez, and Percy Liang. World of  
523 bits: An open-domain platform for web-based agents. In Doina Precup and Yee Whye Teh,  
524 editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of  
525 *Proceedings of Machine Learning Research*, pages 3135–3144. PMLR, 06–11 Aug 2017. URL  
526 <https://proceedings.mlr.press/v70/shi17a.html>.
- 527 [38] Charlie Snell, Ilya Kostrikov, Yi Su, Mengjiao Yang, and Sergey Levine. Offline rl for natural  
528 language generation with implicit language q learning, 2023.
- 529 [39] Daniel Toyama, Philippe Hamel, Anita Gergely, Gheorghe Comanici, Amelia Glaese, Zafarali  
530 Ahmed, Tyler Jackson, Shibl Mourad, and Doina Precup. Androidenv: A reinforcement learning  
531 platform for android. *arXiv preprint arXiv:2105.13231*, 2021.
- 532 [40] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double  
533 q-learning. *CoRR*, abs/1509.06461, 2015. URL <http://arxiv.org/abs/1509.06461>.
- 534 [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,  
535 Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- 536 [42] Siddharth Verma, Justin Fu, Mengjiao Yang, and Sergey Levine. Chai: A chatbot ai for  
537 task-oriented dialogue with offline reinforcement learning, 2022.
- 538 [43] Ziyu Wang, Alexander Novikov, Konrad Zolna, Jost Tobias Springenberg, Scott Reed, Bobak  
539 Shahriari, Noah Siegel, Josh Merel, Caglar Gulcehre, Nicolas Heess, and Nando de Freitas.  
540 Critic regularized regression, 2021.
- 541 [44] Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing  
542 Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. Osvorld: Benchmarking multimodal  
543 agents for open-ended tasks in real computer environments. *arXiv preprint arXiv:2404.07972*,  
544 2024.
- 545 [45] An Yan, Zhengyuan Yang, Wanrong Zhu, Kevin Lin, Linjie Li, Jianfeng Wang, Jianwei Yang,  
546 Yiwu Zhong, Julian McAuley, Jianfeng Gao, Zicheng Liu, and Lijuan Wang. Gpt-4v in  
547 wonderland: Large multimodal models for zero-shot smartphone gui navigation, 2023.
- 548 [46] John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. Intercode: Standardizing  
549 and benchmarking interactive coding with execution feedback, 2023.
- 550 [47] Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu.  
551 Appagent: Multimodal agents as smartphone users. *arXiv preprint arXiv:2312.13771*, 2023.
- 552 [48] Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable  
553 real-world web interaction with grounded language agents, 2023.
- 554 [49] Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang.  
555 Agenttuning: Enabling generalized agent abilities for llms, 2023.

- 556 [50] Chaoyun Zhang, Liqun Li, Shilin He, Xu Zhang, Bo Qiao, Si Qin, Minghua Ma, Yu Kang,  
557 Qingwei Lin, Saravan Rajmohan, et al. Ufo: A ui-focused agent for windows os interaction.  
558 *arXiv preprint arXiv:2402.07939*, 2024.
- 559 [51] Jiwen Zhang, Jihao Wu, Yihua Teng, Minghui Liao, Nuo Xu, Xiao Xiao, Zhongyu Wei, and  
560 Duyu Tang. Android in the zoo: Chain-of-action-thought for gui agents, 2024.
- 561 [52] Zhuosheng Zhang and Aston Zhang. You only look at screens: Multimodal chain-of-action  
562 agents, 2023.
- 563 [53] Ziniu Zhang, Shulin Tian, Liangyu Chen, and Ziwei Liu. Mmina: Benchmarking multihop  
564 multimodal internet agents. *arXiv preprint arXiv:2404.09992*, 2024.
- 565 [54] Boyuan Zheng, Boyu Gou, Jihyung Kil, Huan Sun, and Yu Su. Gpt-4v(ision) is a generalist  
566 web agent, if grounded, 2024.
- 567 [55] Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng,  
568 Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. Webarena: A realistic web  
569 environment for building autonomous agents. *ArXiv*, abs/2307.13854, 2023. URL <https://api.semanticscholar.org/CorpusID:260164780>.  
570
- 571 [56] Yifei Zhou, Andrea Zanette, Jiayi Pan, Sergey Levine, and Aviral Kumar. Archer: Training  
572 language model agents via hierarchical multi-turn rl. *arXiv preprint arXiv:2402.19446*, 2024.
- 573 [57] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei,  
574 Paul F. Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences.  
575 *CoRR*, abs/1909.08593, 2019. URL <http://arxiv.org/abs/1909.08593>.

# 576 Appendices

## 577 A Environment details

### 578 A.1 Post-processing of AiTW

579 The Android in the Wild (AiTW) task set is a large-scale dataset for android device control, containing  
580 five subsets: GoogleApps, Install, Web Shopping, General, and Single, where we select the General  
581 and Web Shopping subsets. Single subset is not considered here because all tasks in Single can be  
582 completed within one step and thus this subset fails to examine the multi-step challenges that we are  
583 interested in this paper. Install and GoogleApps are not considered due to security reasons as those  
584 tasks require an active Google account and parallel emulations can flag security concerns.

585 **General.** The General set focuses on searching for information and basic application usage. For  
586 example, it contains searching for latest news in Chile, search for flights from NYC to Sydney,  
587 opening Gmail, etc. We use all 545 tasks in the training set for training and the first 96 tasks in the  
588 test set for testing due to computational and budget constraints. The maximum allowed number of  
589 steps for this subset is 10. Offline data is collected by rolling our the initial AutoUI policy with tasks  
590 from the training set. The offline data used for the offline-to-online setting contains 608 trajectories  
591 while the offline data used for the offline setting contains 1552 trajectories. Some task examples are  
592 shown in Table 3.

---

Task Example
How do I get to the nearest Verizon Store?
How much does a 2 bedroom apartment rent for in Denver?
Search for flights from Barcelona to Boston
What's a good restaurant in New York?
What's on the menu at Burger King?

---

Table 2: Examples of task descriptions in the AiTW General task set.

593 **Web Shopping.** The Web Shopping subset comprises search instructions on various shopping  
594 websites, like searching for razer blader on ebay. As some websites (e.g. Amazon) and operations  
595 (e.g. adding items to cart) frequently require captcha verifications, we post-process the Web Shopping  
596 subset to exclude such operations and websites and also make the task easy to evaluate for our  
597 autonomous evaluator. The resulting task set involves navigating through five websites (costco.com,  
598 bestbuy.com, target.com, walmart.com, newegg.com) and three basic operations (go to website,  
599 search in the website, and select items from the searched results). Our post-processed training set  
600 contains 438 tasks and our testing set contains 96 tasks. Example tasks after post-processing can  
601 be found in Table 3. The maximum allowed number of steps for this subset is 20. Offline data is  
602 collected by rolling our the initial AutoUI policy with tasks from the training set. The offline data  
603 used for the offline-to-online setting contains 528 trajectories while the offline data used for the  
604 offline setting contains 1296 trajectories.

---

Difficulty	Task Example
1	Go to costco.com Go to walmart.com
2	Go to costco.com, search for "bose soundsport free" Go to walmart.com, search for "logitech g910"
3	Go to costco.com, search for "bose soundsport free" and select the first entry Go to walmart.com, search for "logitech g910" and select the first entry

---

Table 3: Examples of task descriptions in the AiTW Webshopping task set.

## 605 B Qualitative examples

### 606 B.1 Random sample of trajectories for different agents

607 In Figures 8 and 9, we provide trajectories of DigiRL, AutoUI, and GPT-4V randomly sampled  
608 from our test set to offer a qualitative understanding of the agents’ performance. As shown in these  
609 examples, DigiRL can efficiently carry out in-the-wild device control tasks and less likely to get stuck  
610 or get to a wrong page compared to AutoUI and GPT-4V.

### 611 B.2 Error Recovery

612 We observe that DigiRL is able to recover from its own mistakes. As shown in Figure 10, we find  
613 that DigiRL explores ways to get back to the original screen in order to perform a search. As a  
614 comparison, AutoUI fails to reset to the original screen and gets stuck at the diverged screen. Under  
615 the hood, we find DigiRL trying to maximize the state value, which usually induces it to reset to the  
616 original screen (that has a large value to success).

### 617 B.3 Reasoning failure of GPT-4V

618 The performance of GPT-4V failed on AiTW tasks predominantly due to not being able to carry out  
619 control actions as it plans on a high level, and then not being able to recover from these mistakes.  
620 Moreover, one of the main reasons why it is not able to recover from a mistake is that it might  
621 hallucinate and make itself believe that it is a wrong app or website. Indeed, GPT-4V constructs  
622 a plan of further actions when provided a task from either Web Shopping or General dataset of  
623 AiTW. Then, when it makes a misclick and fails to successfully proceed in an intermediate step,  
624 it might think that it actually solved that intermediate step and is in the correct app or website to  
625 execute further actions, causing the overall trajectory to fail. An example of this is provided in  
626 Figure 11. Here, we ask the model to search for an item in a webshopping website, in particular in  
627 “newegg.com”. However, the model fails to proceed to that website due to not being able to precisely  
628 locating the search button. Then, instead of trying to go to that website again, the model thinks it is  
629 already in that webshopping website, and mistakes the search bar of Google with the search bar of  
630 “newegg.com”. Hence, the rest of the trajectory also fails. Another slightly different phenomenon is  
631 illustrated in Figure 12. Here, the model is able to proceed to the correct website and search for an  
632 item, but this time it fails to tap on the search button on the website and clicks to an advertisement  
633 instead. Consequently, the model fools itself to think it successfully searched the item, and scrolls  
634 the page hoping to find that item, but it cannot do so because in reality it views the results of the  
635 advertisement. The primary reason of these failures is the challenge of grounding the control actions  
636 in GUI interfaces to realize the intermediary goals laid out by GPT-4V model’s thoughts. As an  
637 example, we provide an illustration of trying to set up an alarm task in Figure 13. Here, in the last  
638 frame, it fails to execute the precise movements in the necessary amount of rounds to correctly set up  
639 the alarm to the desired time, and in the last frame we see that the action taken does not align with  
640 the thought process of the model.

## 641 C Fine-grained failure modes

642 In Figure 14, we present a more fine-grained breakdown for all six failure modes provided in the user  
643 study. Those failure modes include:

- 644 • *Failure to recover from mistakes* refers to the scenario where the agent made a mistake that  
645 led it to states from which it failed to quickly recover and resume the task, such as a wrong  
646 google search page.
- 647 • *Failure to click on the right link or failure to click* refers to the failure mode where the agent  
648 either fails to locate the element that it tries to click on and keeps clicking on the nearby  
649 region, or fails to start typing in the string when it is supposed to do so.
- 650 • *Failure to take reasonable attempts at all* refers to the failure mode where there is no clear  
651 reason that the agent fails to complete the task and does not seem to be on the right track  
652 throughout the trajectory.

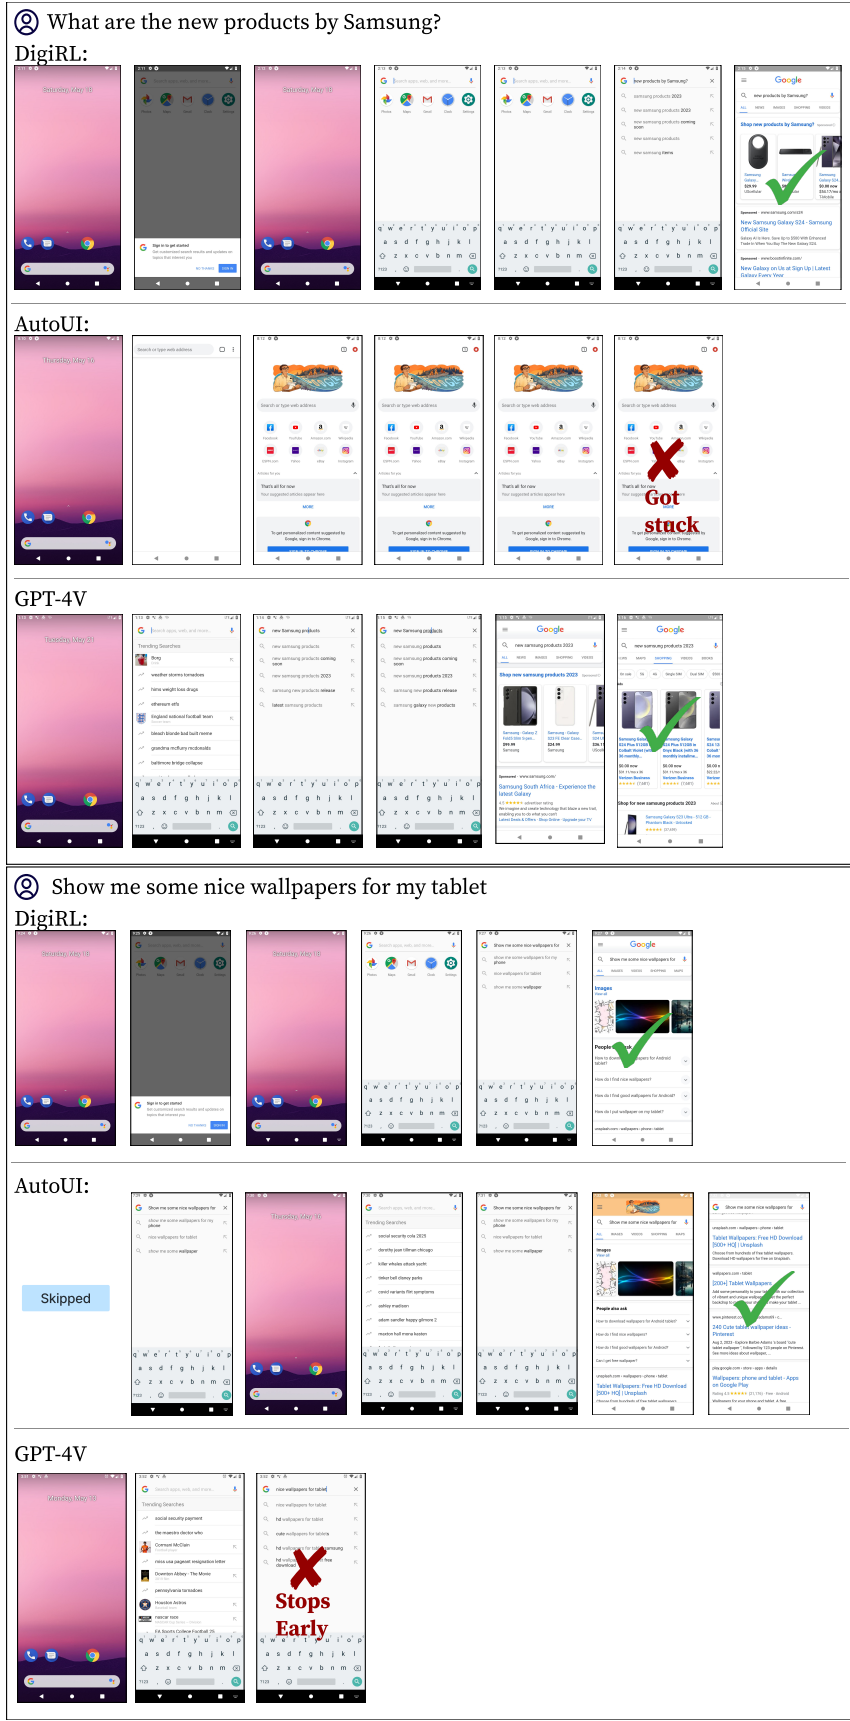



Figure 8: Agents’ trajectory on two randomly sampled tasks on the General split of AitW.



Go to costco.com, search for 'macbook pro', and select the first entry


DigiRL:



AutoUI:




GPT-4V:




Go to newegg.com, search for 'duracell triple a'

DigiRL:



AutoUI:



GPT-4V:




Figure 9: Agents' trajectory on two randomly sampled tasks on the WebShop split of AitW.

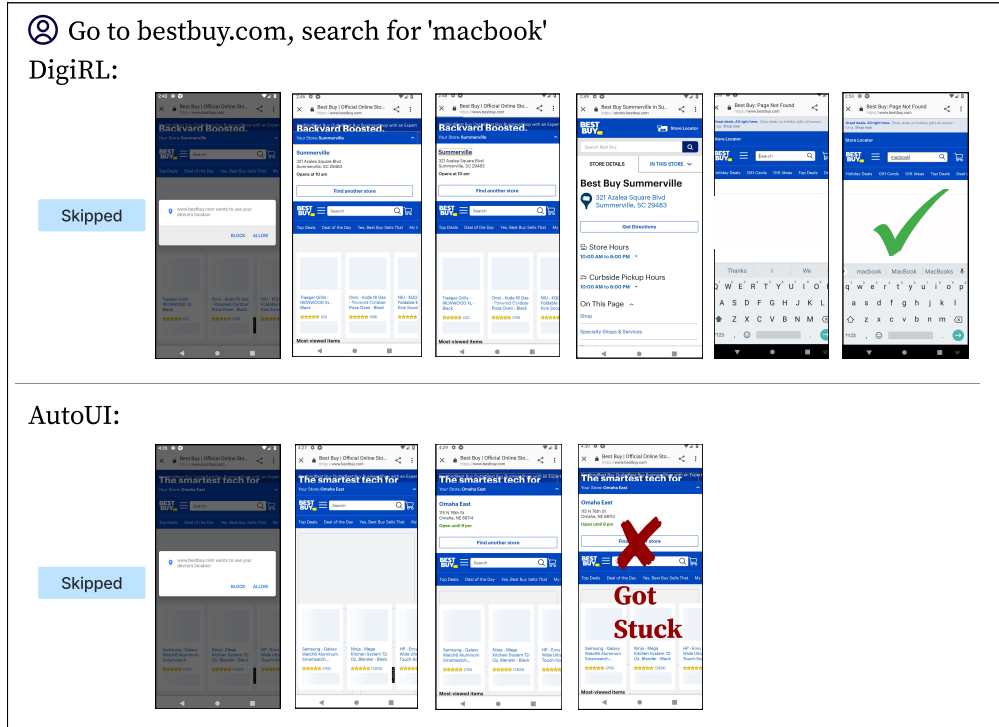


Figure 10: Error recovery cases. In bestbuy.com, we systematically find DigiURL able to recover from its own mistakes, while AutoUI fails to do so.

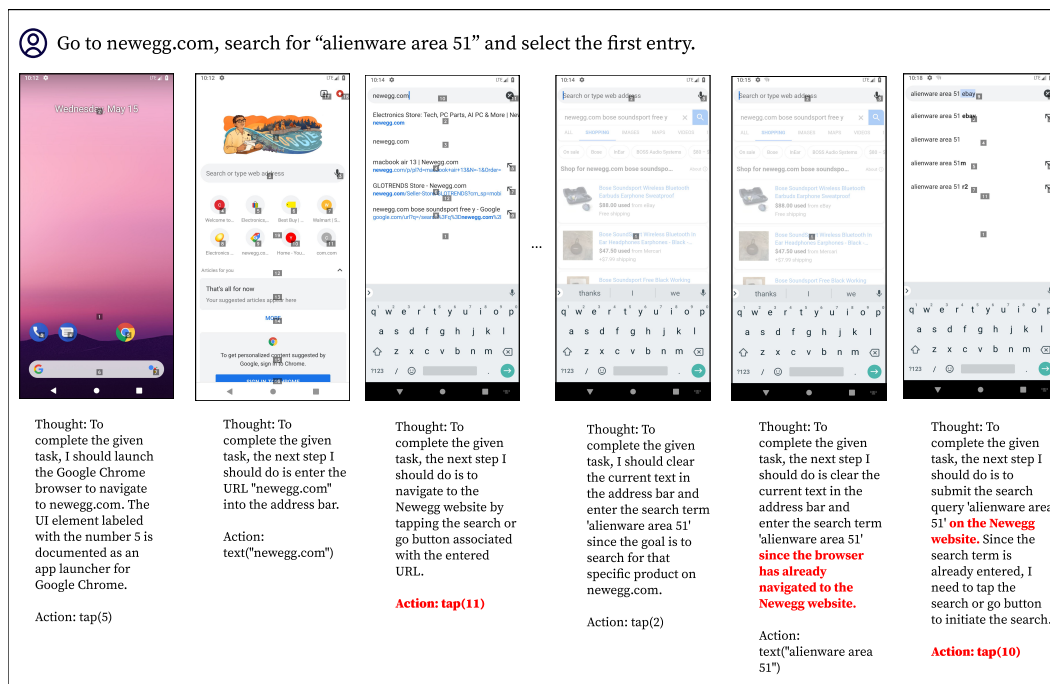


Figure 11: Failure of GPT-4V, with its thoughts and link-based actions given. A typical cause of failure is that it cannot tap on the correct “search” button after entering a query and mistakenly tapped onto the “x” symbol in the search bar as the “search” button. Here the goal is: Go to newegg.com, search for “alienware area 51” and select the first entry. As seen in red emboldened actions, it fails to press search button and deletes the query instead. Also, as seen in red highlighted parts in thoughts, it thinks it is in “newegg.com” website even though it is not.

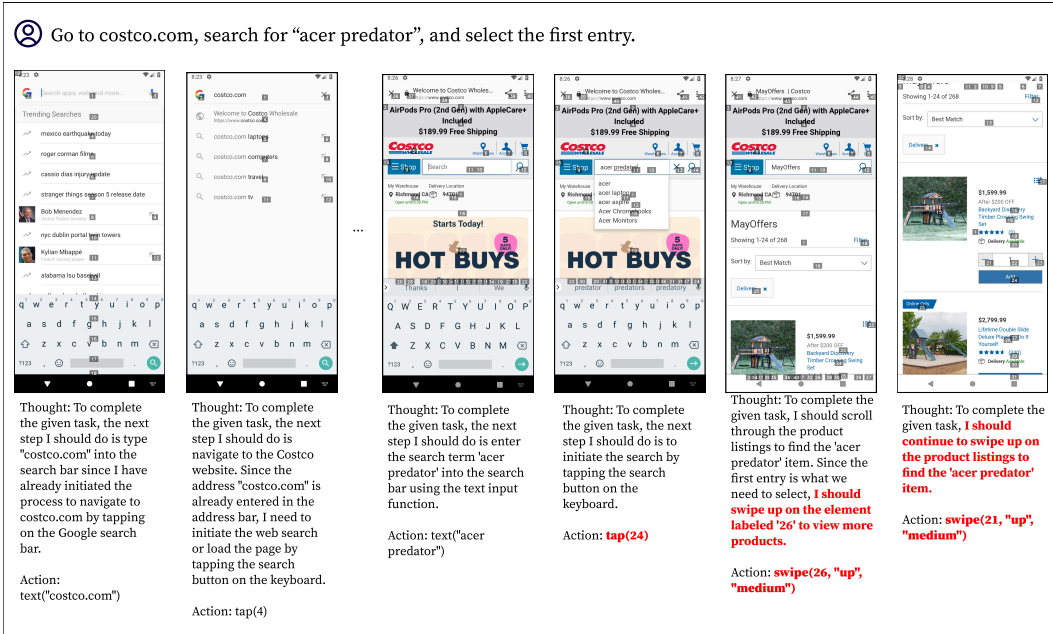


Figure 12: Failure of GPT-4V, with its thoughts and link-based actions given. This time the reason for failure is misclick on the wrong button. The task is “Go to costco.com, search for “acer predator”, and select the first entry”. Notice that up until the fourth frame in this Figure, the trajectory goes correct. But then it clicks on the generic advertisements on the Costco.com website, and it cannot recover back. It continues to scroll the page and takes wrong actions thereafter.

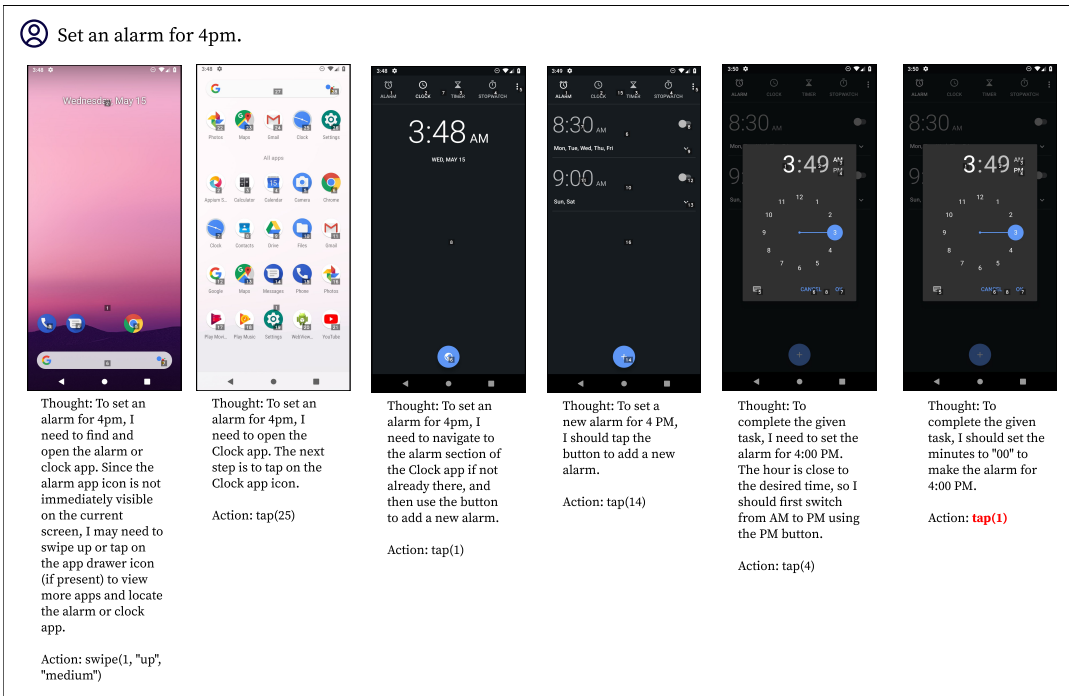


Figure 13: Failure of GPT-4V, with an example task on the AiTW general test set. The task is “Set an alarm for 4pm”. Here, GPT-4V is able to successfully navigate to the clock app, and the alarm settings of that app. However, it cannot take the correct precise actions to set the alarm quickly enough, and it fails due to maximum rounds reached. In the last round, notice that the action of tap(1) contradict with its own thought process of setting minutes to “00”.




Figure 14: Failure modes decomposition for each policy model for both General and Web Shopping subsets.

- 653 • *Quit or press HOME early* refers to the failure mode where the agent decided to finish the task or press HOME to start over before the task is actually finished.
- 654
- 655 • *Stops at wrong but relevant page* refers to the failure mode where the agent arrives at a wrong page and mistakenly thinks that it had completed the task. For example, the agent finds a macbook on costco.com while the instruction asked it to find a macbook on ebay.com.
- 656
- 657
- 658 • *Technical issues* refer to the failure mode that either the task is impossible (e.g. the tasks asks to open Amazon app but this app is not installed) or the agent is temporarily blocked from a certain website due to frequent visits.
- 659
- 660

661 The translation between fine-grained failure modes and coarse-grained failure modes is presented in Table 4.

Fine-Grained Failure	Coarse-Grained Failure
Fail to recover from mistakes	Fail to recover from mistakes
Fail to click on the right link or fail to type	Get stuck midway
Fail to take reasonable attempts at all	Get stuck midway
Quit or Press HOME early	Arrive at wrong goal
Stops at wrong but relevant page	Arrive at wrong goal
Technical Issues	None

Table 4: Examples of task descriptions in the AiTW Webshopping task set.

662

## 663 D Experiment machines

664 Our main experiments are conducted on VM instances from Google Cloud Platform. Each VM instance comes with 1x Tesla T4 GPU and 16x Intel(R) Xeon(R) CPU.

## 666 E Setup for parallel environment

667 Running multiple emulators in parallel can be challenging due to the inefficiency in thread synchronization and frequent fault propagation when one emulator runs into an unknown error. To address this challenge, we set up a server-client system where all emulator processes are running in independent server processes. Each emulator process communicates with the main training process through different UIAutomotor servers. The main training process sends high-level instructions to UIAutomotor servers (such as reset and step), while UIAutomotor servers parse high-level instructions into low-level UI commands (such as typing a character and tapping at a coordinate) and such UI commands are executed by the emulator processes. When an exception is thrown in the emulator, the UIAutomotor examines if it is recoverable (e.g. an UI command takes too long to execute in the emulator) and reset the emulator process if it is not. When an exception is thrown in the UIAutomotor server, the main training process stops and resets the UIAutomotor server to ensure data correctness.

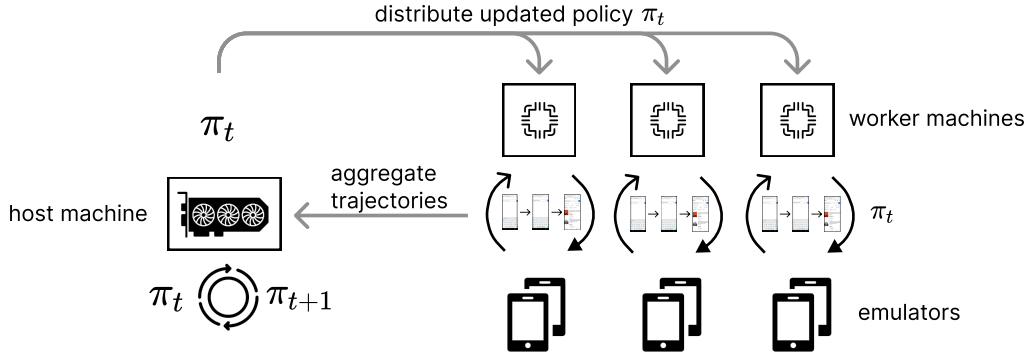


Figure 15: Multi-machine parallel emulator execution. The host machine is equipped with GPU accelerators and the worker machines are equipped only with CPUs. The policy update is executed on the worker machine and the trajectory collections are executed distributedly on the worker machines and aggregated by the host machine.

678 This design can easily be scaled up to a multi-  
 679 machine setting. As illustrated in Figure 15, one  
 680 host machine equipped with GPU accelerator has a  
 681 local copy of the current policy  $\pi_t$ , and distributes  
 682 the policy to all worker machines equipped with only  
 683 one GPU and multiple CPUs. Each worker machine  
 684 will then collect trajectories of different tasks using  
 685  $\pi_t$ . After all collection processes are synchronized,  
 686 the host machine gathers all the trajectories together  
 687 to update the policy to  $\pi_{t+1}$ . This process keeps  
 688 iterating until the policy converges.

689 The performance boost with respect to the number  
 690 of worker machines is nearly linear, as demonstrated  
 691 in Figure 16, where we conduct experiments that  
 692 examine the scaling performance of our parallel em-  
 693 ulator. Our distributed emulator that runs emulations  
 694 across multiple servers can reliably collect data with  
 695 up to 64 parallel emulators on 128 CPUs with near-  
 696 linear speedup. In contrast, a naive baseline that runs all parallel emulations on the same server  
 697 achieves much inferior performance (0.74 compared to 1.74 trajs/min using 64 CPUs).

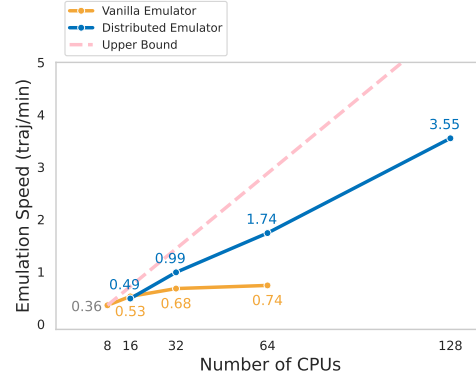


Figure 16: **Emulation speed w.r.t number of CPUs used.** The upper bound can only be achieved when there is no communication and error handling cost. Our design of distributed emulator can significantly improve the efficiency of emulation compared to the vanilla method of running all emulations over the same instance.

## 698 F Autonomous evaluator details

699 Our autonomous evaluator gives a reward to each observation we get. The observation is composed  
 700 of the current screenshot of device and the task. The evaluator gives a reward of 1 if the screenshot  
 701 shows a completion of the task, and will terminate the POMDP as a result result.

702 The optimized prompt is shown in Figure 17 and Figure 18 for General and Web Shopping subsets  
 703 respectively.

## Prompt

You're an expert in evaluating whether the Screenshot successfully completes the Task.

### ====Examples====

Screenshot: {train\_1.png}

Task: Open the settings.

Q: What should I expect to see on the screenshot if I've opened the settings?

A: I should expect to see I'm in the settings app. The screenshot shows the home screen of a mobile device, with various app icons displayed, including the settings app icon, but the settings app is not opened.

Status: failure

Screenshot: {train\_2.png}

Task: Find hotels in washington dc

Q: What should I expect to see on the screenshot if I've searched for hotels in Washington, DC?

A: I should expect to see I'm in a search results page for hotels in Washington, DC. The screenshot shows a Google search page with the search field populated with the query "hotels in washington dc" and a list of suggested searches related to hotels in Washington, DC, but it does not show any search results for hotels in Washington, DC.

Status: failure

Screenshot: {train\_3.png}

Task: What's a good restaurant in Portland?

Q: What should I expect to see on the screenshot if I've searched for a good restaurant in Portland?

A: I should expect to see I'm in a search results page for a good restaurant in Portland. The screenshot shows a Google search page with a search input field for "good restaurant in portland" and a map results preview showing business locations near Portland, like "Li Pigeon", "Portland City Grill", and "Higgins",

Status: success

... (more cases)

### ====Your Turn====

Screenshot: {test.png}

Task: {task\_this\_traj}

Respond in this format:

Q: What should I expect to see on the screenshot if I've <repeat the task>?

A: I should expect to see <first expectation, then what's in the given screenshot.>

Status: success or failure (don't return anything else)

Start with "Q:".

## Response

Q: What should I expect to see on the screenshot if I've searched for the price of a 12' ladder at Home Depot?

A: I should expect to see the price of a 12' ladder at Home Depot; the screenshot shows a search result page for the price of a 12' ladder, with some product advertisements showing prices from Home Depot.

Status: success

## Image Sources

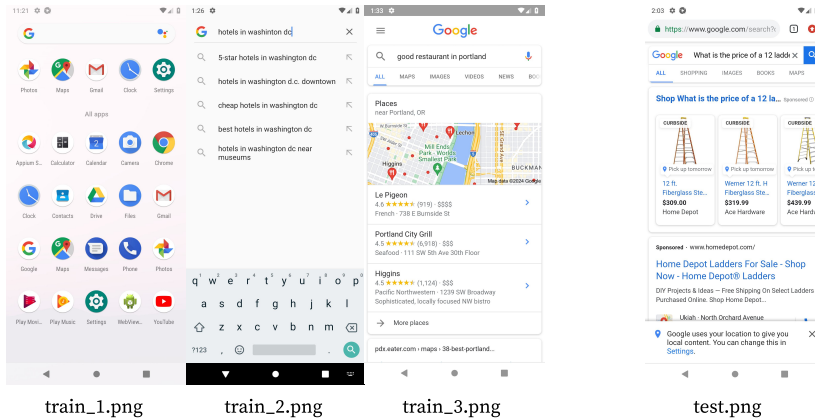


Figure 17: Prompt for our autonomous evaluator for tasks in AitW General subset.

## Prompt

You're an expert in evaluating whether the Screenshot successfully completes the Task.

### ====Examples====

Screenshot: {train\_1.png}

Task: Go to bestbuy.com

Q: What should I expect to see on the screenshot if I've gone to bestbuy.com?

A: I should expect to see I'm in the Best Buy website, which usually shows the best buy logo with some featured products and categories. The screenshot shows I'm searching for "bestbuy.com" in the Google search (with some search suggestions) instead of being in the Best Buy website.

Status: failure

Screenshot: {train\_2.png}

Task: Go to ebay.com, search for "corsair k70"

Q: What should I expect to see on the screenshot if I've gone to ebay.com AND searched for "corsair k70"?

A: I should expect to see I'm in the eBay website and search results for "corsair k70". The screenshot shows I'm in the eBay website with some search suggestions for "corsair k70", but it does not show search results of the product, which usually includes price and the product details.

Status: failure

Screenshot: {train\_3.png}

Task: Go to ebay.com, search for "lenovo thinkpad"

Q: What should I expect to see on the screenshot if I've gone to ebay.com AND searched for "lenovo thinkpad"?

A: I should expect to see I'm in the eBay website and search results for "lenovo thinkpad". The screenshot shows I'm in the eBay website and have several search results for "lenovo thinkpad".

Status: success

... (more cases)

### ====Your Turn====

Screenshot: {test.png}

Task: {task\_this\_traj}

Respond in this format:

Q: What should I expect to see on the screenshot if I've <repeat the task>?

A: I should expect to see <first expectation, then what's in the given screenshot.>

Status: success or failure (don't return anything else)

Start with "Q:".

## Response

Q: What should I expect to see on the screenshot if I've searched for the price of a 12' ladder at Home Depot?

A: I should expect to see the price of a 12' ladder at Home Depot; the screenshot shows a search result page for the price of a 12' ladder, with some product advertisements showing prices from Home Depot.

Status: success

## Image Sources

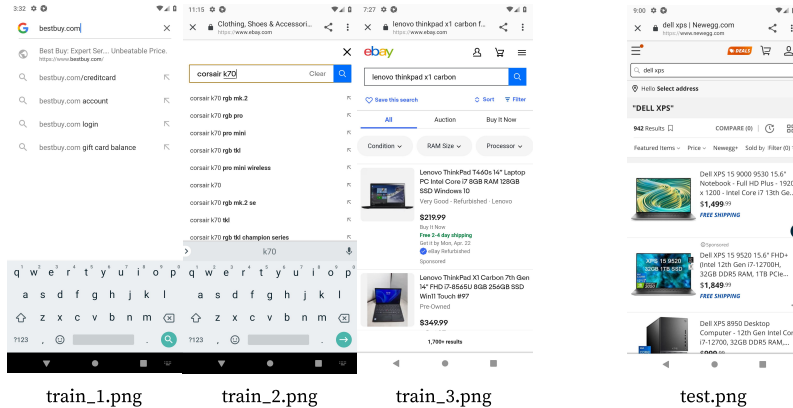


Figure 18: Prompt for our autonomous evaluator for tasks in AitW Web Shopping subset.

705 Figure 19 shows the prompt that we used for testing the Set-of-Marks performance for GPT-4V and  
706 Gemini 1.5 Pro. This prompt is directly taken from Yang et al. [47].

**Prompt**

"You are an agent that is trained to perform some basic tasks on a smartphone. You will be given a \nsmartphone screenshot. The interactive UI elements on the screenshot are labeled with numeric tags starting from 1. The \nnumeric tag of each interactive element is located in the center of the element.\n\nYou can call the following functions to control the smartphone:\n\n1. tap(element: int)\nThis function is used to tap an UI element shown on the smartphone screen.\n\n\"element\" is a numeric tag assigned to an UI element shown on the smartphone screen. \n\nA simple use case can be tap(5), which taps the UI element labeled with the number 5.\n\n2. text(text\_input: str)\nThis function is used to insert text input in an input field/box. text\_input is the string you want to insert and must \nbe wrapped with double quotation marks. A simple use case can be text(\"Hello, world!\"), which inserts the string \n\"Hello, world!\" into the input area on the smartphone screen. This function is usually callable when you see a keyboard \nshowing in the lower half of the screen.\n\n3. long\_press(element: int)\nThis function is used to long press an UI element shown on the smartphone screen.\n\n\"element\" is a numeric tag assigned to an UI element shown on the smartphone screen.\n\nA simple use case can be long\_press(5), which long presses the UI element labeled with the number 5.\n\n4. swipe(element: int, direction: str, dist: str)\nThis function is used to swipe an UI element shown on the smartphone screen, usually a scroll view or a slide bar.\n\n\"element\" is a numeric tag assigned to an UI element shown on the smartphone screen. \n\n\"direction\" is a string that \nrepresents one of the four directions: up, down, left, right. \n\n\"direction\" must be wrapped with double quotation \nmarks. \n\n\"dist\" determines the distance of the swipe and can be one of the three options: short, medium, long. You should \nchoose the appropriate distance option according to your need.\n\nA simple use case can be swipe(21, \"up\", \"medium\"), which swipes up the UI element labeled with the number 21 for a \nmedium distance.\n\n5. grid()\nYou should call this function when you find the element you want to interact with is not labeled with a numeric tag and \nother elements with numeric tags cannot help with the task. The function will bring up a grid overlay to divide the \nsmartphone screen into small areas and this will give you more freedom to choose any part of the screen to tap, long \npress, or swipe.

The task you need to complete is to **How much does a 2 bedroom apartment rent for in Denver?**.

Your past actions to proceed with this task are summarized as follows: **None**

Now, given the documentation and the following labeled screenshot, you need to think and call the function needed to proceed with the task. Your output should include three parts in the given format:  
 Observation: <Describe what you observe in the image>  
 Thought: <To complete the given task, what is the next step I should do>  
 Action: <The function call with the correct parameters to proceed with the task. When you are certain that the task is successfully done and the goal is reached as of the current observation, you should output FINISH. You cannot output anything else except a function call or FINISH \nin this field.>  
 Summary: <Summarize your past actions along with your latest action in one or two sentences. Do not include the numeric \ntag in your summary>\n\nYou can only take one action at a time, so please directly call the function."

Figure 19: Set-of-Marks prompting. The boldened inputs can be changed according to our goal. The task changes for every different task. The past actions change as we take actions (it is None now since this is the prompt for the first round).



## 707 H Other Experiments

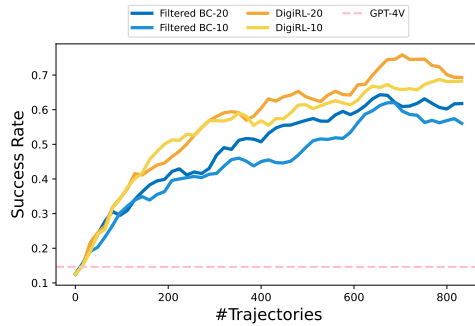


Figure 20: **Success rate with different horizon length** ( $H \in \{10, 20\}$ ) under different methods on the AiTW Google Search task set.

### 708 H.1 Horizon Limit

709 We investigate the horizon limit of filtered BC and DigiRL on the AiTW General subset. As most tasks  
710 can be effectively solved within 10 steps, we specify two horizon limits: a sufficient horizon  $H = 10$ ,  
711 and a redundant horizon  $H = 20$ . Results show that a redundant horizon introduces significantly  
712 faster learning speed for both filtered BC and DigiRL, presumably because longer horizon means  
713 more opportunity to try in a single trajectory. In both horizon settings, we observe the DigiRL offers  
714 a significant speedup of around 100 trajectories over Filtered BC.

## 715 I Hyperparameters

716 Hyperparameters for both Filtered BC and DigiRL are carefully tuned through binary search on the  
717 training set of General and Web Shopping subsets. The final choice of hyperparameters for both  
718 methods can be found in Table 5. As shown in the table, the only hyperparameters introduced by  
719 DigiRL are supervised training hyperparameters for the value function and instruction value function  
720 (including number of iterations and learning rate) and GAE  $\lambda$ .

Table 5: Hyperparameters for All Experiments

Method	Hyperparameter	Offline	Offline-to-Online
Filtered BC	actor lr	3e-3	3e-3
	batch size	128	128
	rollout trajectories	-	16
	replay buffer size	-	5000
	rollout temperature	-	1.0
	maximum gradient norm	0.01	0.01
	actor updates per iteration	20	20
	number of iterations for offline actor updates	10	10
DigiRL	actor lr	3e-3	3e-3
	value function lr	3e-3	3e-3
	instruction value function lr	3e-3	3e-3
	instruction value function lr	3e-3	3e-3
	batch size	128	128
	rollout trajectories	-	16
	replay buffer size	-	5000
	rollout temperature	-	1.0
	maximum gradient norm	0.01	0.01
	GAE $\lambda$	0.5	0.5
	actor updates per iteration	20	20
	value function updates per iteration	5	5
	instruction value function updates per iteration	-	5
	number of iterations for offline actor updates	10	10
	number of iterations for offline value function updates	20	20
number of iterations for offline instruction value function updates	-	20	

Table 6: Hyperparameters for DigiRL and Filtered BC on both General and Web Shopping subset of AitW..

## 721 NeurIPS Paper Checklist

### 722 1. Claims

723 Question: Do the main claims made in the abstract and introduction accurately reflect the  
724 paper’s contributions and scope?

725 Answer: [Yes]

726 Justification: The main claims in the abstract and introduction explicitly state the contribu-  
727 tions of the paper.

728 Guidelines:

- 729 • The answer NA means that the abstract and introduction do not include the claims  
730 made in the paper.
- 731 • The abstract and/or introduction should clearly state the claims made, including the  
732 contributions made in the paper and important assumptions and limitations. A No or  
733 NA answer to this question will not be perceived well by the reviewers.
- 734 • The claims made should match theoretical and experimental results, and reflect how  
735 much the results can be expected to generalize to other settings.
- 736 • It is fine to include aspirational goals as motivation as long as it is clear that these goals  
737 are not attained by the paper.

### 738 2. Limitations

739 Question: Does the paper discuss the limitations of the work performed by the authors?

740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791

Answer: [Yes]

Justification: Limitations are discussed in the last section of the paper.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This paper does not provide theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All loss functions and implementation details are provided in Section 4.

Guidelines:

- 792 • The answer NA means that the paper does not include experiments.
- 793 • If the paper includes experiments, a No answer to this question will not be perceived
- 794 well by the reviewers: Making the paper reproducible is important, regardless of
- 795 whether the code and data are provided or not.
- 796 • If the contribution is a dataset and/or model, the authors should describe the steps taken
- 797 to make their results reproducible or verifiable.
- 798 • Depending on the contribution, reproducibility can be accomplished in various ways.
- 799 For example, if the contribution is a novel architecture, describing the architecture fully
- 800 might suffice, or if the contribution is a specific model and empirical evaluation, it may
- 801 be necessary to either make it possible for others to replicate the model with the same
- 802 dataset, or provide access to the model. In general, releasing code and data is often
- 803 one good way to accomplish this, but reproducibility can also be provided via detailed
- 804 instructions for how to replicate the results, access to a hosted model (e.g., in the case
- 805 of a large language model), releasing of a model checkpoint, or other means that are
- 806 appropriate to the research performed.
- 807 • While NeurIPS does not require releasing code, the conference does require all submissions
- 808 to provide some reasonable avenue for reproducibility, which may depend on the
- 809 nature of the contribution. For example
  - 810 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
  - 811 to reproduce that algorithm.
  - 812 (b) If the contribution is primarily a new model architecture, the paper should describe
  - 813 the architecture clearly and fully.
  - 814 (c) If the contribution is a new model (e.g., a large language model), then there should
  - 815 either be a way to access this model for reproducing the results or a way to reproduce
  - 816 the model (e.g., with an open-source dataset or instructions for how to construct
  - 817 the dataset).
  - 818 (d) We recognize that reproducibility may be tricky in some cases, in which case
  - 819 authors are welcome to describe the particular way they provide for reproducibility.
  - 820 In the case of closed-source models, it may be that access to the model is limited in
  - 821 some way (e.g., to registered users), but it should be possible for other researchers
  - 822 to have some path to reproducing or verifying the results.

## 823 5. Open access to data and code

824 Question: Does the paper provide open access to the data and code, with sufficient instruc-  
 825 tions to faithfully reproduce the main experimental results, as described in supplemental  
 826 material?

827 Answer: [No]

828 Justification: We are still actively cleaning the code and make the environment more  
 829 accessible to a broader audience. Once we are done with that, we will open-source the code  
 830 along with the release of the paper.

831 Guidelines:

- 832 • The answer NA means that paper does not include experiments requiring code.
- 833 • Please see the NeurIPS code and data submission guidelines ([https://nips.cc/  
 834 public/guides/CodeSubmissionPolicy](https://nips.cc/public/guides/CodeSubmissionPolicy)) for more details.
- 835 • While we encourage the release of code and data, we understand that this might not be
- 836 possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not
- 837 including code, unless this is central to the contribution (e.g., for a new open-source
- 838 benchmark).
- 839 • The instructions should contain the exact command and environment needed to run to
- 840 reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- 841 • The authors should provide instructions on data access and preparation, including how
- 842 to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- 843 • The authors should provide scripts to reproduce all experimental results for the new
- 844 proposed method and baselines. If only a subset of experiments are reproducible, they
- 845 should state which ones are omitted from the script and why.
- 846

- 847
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- 848
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.
- 849
- 850

## 851 6. Experimental Setting/Details

852 Question: Does the paper specify all the training and test details (e.g., data splits, hyper-  
853 parameters, how they were chosen, type of optimizer, etc.) necessary to understand the  
854 results?

855 Answer: [\[Yes\]](#)

856 Justification: Dataset details are provided in Appendix [A.1](#) and hyperparameters are provided  
857 in Appendix [I](#).

858 Guidelines:

- 859 • The answer NA means that the paper does not include experiments.
- 860 • The experimental setting should be presented in the core of the paper to a level of detail  
861 that is necessary to appreciate the results and make sense of them.
- 862 • The full details can be provided either with the code, in appendix, or as supplemental  
863 material.

## 864 7. Experiment Statistical Significance

865 Question: Does the paper report error bars suitably and correctly defined or other appropriate  
866 information about the statistical significance of the experiments?

867 Answer: [\[Yes\]](#)

868 Justification: Repeated experiments are carried out with their means and standard deviations  
869 reported in Table [1](#).

870 Guidelines:

- 871 • The answer NA means that the paper does not include experiments.
- 872 • The authors should answer "Yes" if the results are accompanied by error bars, confi-  
873 dence intervals, or statistical significance tests, at least for the experiments that support  
874 the main claims of the paper.
- 875 • The factors of variability that the error bars are capturing should be clearly stated (for  
876 example, train/test split, initialization, random drawing of some parameter, or overall  
877 run with given experimental conditions).
- 878 • The method for calculating the error bars should be explained (closed form formula,  
879 call to a library function, bootstrap, etc.)
- 880 • The assumptions made should be given (e.g., Normally distributed errors).
- 881 • It should be clear whether the error bar is the standard deviation or the standard error  
882 of the mean.
- 883 • It is OK to report 1-sigma error bars, but one should state it. The authors should  
884 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis  
885 of Normality of errors is not verified.
- 886 • For asymmetric distributions, the authors should be careful not to show in tables or  
887 figures symmetric error bars that would yield results that are out of range (e.g. negative  
888 error rates).
- 889 • If error bars are reported in tables or plots, The authors should explain in the text how  
890 they were calculated and reference the corresponding figures or tables in the text.

## 891 8. Experiments Compute Resources

892 Question: For each experiment, does the paper provide sufficient information on the com-  
893 puter resources (type of compute workers, memory, time of execution) needed to reproduce  
894 the experiments?

895 Answer: [\[Yes\]](#)

896 Justification: This information is provided in Appendix [D](#).

897 Guidelines:

- 898
- The answer NA means that the paper does not include experiments.
  - 899
  - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
  - 900
  - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
  - 901
  - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).
  - 902
  - 903
  - 904
  - 905

## 906 9. Code Of Ethics

907 Question: Does the research conducted in the paper conform, in every respect, with the  
908 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

909 Answer: [Yes]

910 Justification: The research conducted in the paper conform, in every respect, with the  
911 NeurIPS code of Ethics.

912 Guidelines:

- 913 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 914 • If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- 915
- 916 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-  
917 eration due to laws or regulations in their jurisdiction).

## 918 10. Broader Impacts

919 Question: Does the paper discuss both potential positive societal impacts and negative  
920 societal impacts of the work performed?

921 Answer: [Yes]

922 Justification: The positive societal impacts are discussed in the Introduction while the  
923 negative societal impacts are discussed in Section 6.

924 Guidelines:

- 925 • The answer NA means that there is no societal impact of the work performed.
- 926 • If the authors answer NA or No, they should explain why their work has no societal  
927 impact or why the paper does not address societal impact.
- 928 • Examples of negative societal impacts include potential malicious or unintended uses  
929 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations  
930 (e.g., deployment of technologies that could make decisions that unfairly impact specific  
931 groups), privacy considerations, and security considerations.
- 932 • The conference expects that many papers will be foundational research and not tied  
933 to particular applications, let alone deployments. However, if there is a direct path to  
934 any negative applications, the authors should point it out. For example, it is legitimate  
935 to point out that an improvement in the quality of generative models could be used to  
936 generate deepfakes for disinformation. On the other hand, it is not needed to point out  
937 that a generic algorithm for optimizing neural networks could enable people to train  
938 models that generate Deepfakes faster.
- 939 • The authors should consider possible harms that could arise when the technology is  
940 being used as intended and functioning correctly, harms that could arise when the  
941 technology is being used as intended but gives incorrect results, and harms following  
942 from (intentional or unintentional) misuse of the technology.
- 943 • If there are negative societal impacts, the authors could also discuss possible mitigation  
944 strategies (e.g., gated release of models, providing defenses in addition to attacks,  
945 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from  
946 feedback over time, improving the efficiency and accessibility of ML).

## 947 11. Safeguards

948 Question: Does the paper describe safeguards that have been put in place for responsible  
949 release of data or models that have a high risk for misuse (e.g., pretrained language models,  
950 image generators, or scraped datasets)?

951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971  
972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002

Answer: [NA]

Justification: The capability of the model that we will be releasing is limited to simple tasks in Android in the Wild dataset, and therefore does not have a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have properly cited the assets that we are using.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

## 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This submission does not include new assets. New assets including open-sourced code, model checkpoints, and model trajectories will be released with documentation when we release the paper.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025  
1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037

#### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This research does not involve crowdsourcing or human subjects. Annotations of trajectories in Figure 5 and Figure 6 are carried out by authors alone.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

#### 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This research does not involve crowdsourcing or human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.