

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 VIMO: A GENERATIVE VISUAL GUI WORLD MODEL FOR APP AGENTS

Anonymous authors

Paper under double-blind review

ABSTRACT

App agents, which autonomously operate mobile Apps through GUIs, have gained significant interest in real-world applications. Yet, they often struggle with long-horizon planning, failing to find the optimal actions for complex tasks with longer steps. To address this, world models are used to predict the next GUI observation based on user actions, enabling more effective agent planning. However, existing world models primarily focus on generating only textual descriptions, lacking essential visual details. To fill this gap, we propose **ViMo**, the first **Visual world Model** designed to generate future App observations as images. For the challenge of generating text in image patches, where even minor pixel errors can distort readability, we decompose GUI generation into graphic and text content generation. We propose a novel data representation, the **Symbolic Text Representation (STR)**, to overlay text content with symbolic placeholders while preserving graphics. With this design, ViMo employs a **STR Predictor** to predict future GUIs' graphics and a **GUI-text Predictor** for generating the corresponding text. Moreover, we deploy ViMo to enhance agent-focused tasks by predicting the outcome of actions. Experiments show that ViMo establishes visual world models as a compelling alternative to language-based approaches, producing visually plausible and functionally effective GUIs that empower App agents with more informed decisions.

1 INTRODUCTION

Recent advancements in Large Language Models (LLMs)¹ have unlocked new possibilities for deploying AI agents across diverse fields (Li et al., 2023; Gou et al., 2023; Rawles et al., 2024b). A notable application is the smartphone application (App) agents (Rawles et al., 2024a; Wang et al., 2024a), designed to directly interact with Graphical User Interfaces (GUIs) to perform tasks autonomously and efficiently in a mobile operating system.

However, existing agents struggle with making decisions for tasks requiring longer steps (Chae et al., 2024). To address this "long-horizon" limitation, an increasing number of studies have introduced world models, which predict how GUIs evolve in response to user actions (Gu et al., 2024). Yet, these models typically rely on language to describe future observations. These language-based descriptions often fail to capture the intricate visual details, such as the location and colour of GUI elements, necessary for a precise representation (Chae et al., 2024). A seemingly straightforward solution is to execute action candidates on App emulators. However, real-world execution is impractical for scalable planning since actions like payments or repeated purchases are difficult to backtrack. Similar concerns have motivated the broader world-model community to explore ML-based simulators (Li et al., 2025; Hafner et al., 2019; Hu et al., 2022). We tackle the problem in the GUI domain by designing a GUI world model capable of predicting hypothetical observations in the visual modality.

To build a visual GUI world model capable of generating plausible future GUI observations that are visually consistent with user actions, a straightforward approach involves generating each pixel of a GUI using image generations (Brooks et al., 2023; Rombach et al., 2022). Although these methods demonstrate promising results, such as the GUI graphic generation on the location, style, and colour of GUI elements (Wei et al., 2024), or scene-text generation in a style that aligns the visual

¹By LLMs, we refer to the concept of foundation models that accept various modalities (e.g., visual language models (VLMs), multimodal LLMs (MLLMs)) while producing textual sequences (W. contributors, 2024).



062
063 Figure 1: GUIs generated by image-based methods (UI-Diffuser (Wei et al., 2024), TextDiffuser-
064 2 (Chen et al., 2024c), and IP2P (Brooks et al., 2023) fine-tuned on GUI dataset, denoted as IP2P*).

065
066 context (Chen et al., 2024c; Zhang et al., 2024c), they still display distortions in the text rendering,
067 particularly for small-sized text where each pixel is critical for accurately identifying and representing
068 the text (see Fig. 1 for an illustration).

069 To address the challenges of accurately generating high-fidelity text content within a GUI, we propose
070 **ViMo**, the first visual GUI world model. ViMo decouples the generation of graphic and text content
071 into distinct processes, using a novel data representation named Symbolic Text Representation (**STR**).
072 In STR, each text content is replaced (overlaid) with a text symbol, a rectangle-shaped placeholder
073 with a defined border and fill colours, functioning as a special GUI element. Thus, we simplify
074 the task of text content generation to text symbol generation, which reframes the problem to the
075 localisation of the text within a GUI. Based on STR, ViMo employs a **STR Predictor** and a **GUI-text**
076 **Predictor** to generate the graphic and the text content respectively. Specifically, the STR predictor
077 is implemented as a diffusion model, taking the current STR, extracted from the given GUI, and
078 a user action as inputs to generate the STR of the next GUI. Meanwhile, the GUI-text predictor,
079 implemented based on an LLM, leverages the STR generated by the STR predictor to produce the
080 corresponding text for each text symbol. Finally, the predicted STR and the generated text are
081 combined to produce the next GUI.

082 We evaluated ViMo in three distinct scenarios to comprehensively demonstrate its effectiveness.
083 First, we assessed its world model capability, where the quality of the generated GUIs was measured
084 using visual similarity, instructional accuracy, and action readiness scores. Each score was examined
085 through both automatic metrics and user studies. These assessments provided a robust and holistic
086 evaluation of how visually precise and contextually plausible the generated GUIs were. Second,
087 we tested ViMo in an agent-focused task to evaluate its benefits for existing App agents and its
088 superiority over other language-based and image-based world models. In this setup, given a goal
089 and the current App observation, the agent selected optimal actions to achieve the goal (Wang et al.,
090 2024a). By accurately predicting the next GUI based on the current observation and an action,
091 ViMo enabled the agent to better anticipate action outcomes and make more informed decisions.
092 This experiment demonstrated the model’s effectiveness in enhancing decision-making for App
093 agents. Finally, we evaluated ViMo’s real-world applicability under online navigation and zero-
094 shot generalisation settings. These scenarios assessed the model’s ability to perform in real-time
095 interactions and to generalise to previously unseen Apps, further demonstrating its generalisation
096 capabilities and practical value in dynamic environments.

097 Our main contributions are summarised as follows:

- 098 • We propose ViMo, the first generative visual GUI world model that predicts App observations
099 in a visual modality, capable of more realistic and concrete visual GUI predictions compared to
100 contemporary language-based methods.
- 101 • To address the challenge of strict pixel-level accuracy required to avoid distorted or blurred text
102 generation in a GUI, we propose a Symbolic Text Representation (STR), overlaying text with uniform
103 text symbols (placeholders) to simplify text content generation to text location generation. Then
104 ViMo leverages an LLM to generate the corresponding text content for each text symbol.
- 105 • Extensive experiments demonstrated the effectiveness of ViMo in both world model evaluation
106 and agent-focused tasks. Specifically, ViMo achieved an average 29.14% and 182.74% relative
107 improvement over existing world models in terms of automatic metrics and user studies, respectively.
Moreover, ViMo boosted the step-wise action prediction accuracy of App agents, achieving a 14.07%

108 relative performance gain. In the online navigation setup, ViMo increased the task completion rate
 109 from 33.19% to 40.95%, yielding a substantial improvement of 7.76%.
 110

112 2 RELATED WORKS

114 2.1 APP AGENT

116 App agents, powered by LLMs, have advanced task automation on mobile Apps (Wen et al., 2024b;
 117 Chen et al., 2024a; Zhang et al., 2024a;b; Lee et al., 2023). These agents interact with GUIs by
 118 emulating human actions. Approaches in this domain are broadly divided into *language-based* and
 119 *multi-modality-based* methods. Language-based methods rely on textual description of the App
 120 observation and the user goal to generate appropriate actions (Wen et al., 2024a; Li et al., 2024),
 121 while multi-modality-based methods enhance this capability by incorporating GUIs for a more
 122 comprehensive understanding of the interface (Christianos et al., 2024; Wang et al., 2024b). However,
 123 these approaches struggle with long-horizon tasks that require multiple interdependent actions and a
 124 deep understanding of dynamic environments (Chae et al., 2024). For this challenge, a straightforward
 125 solution is to use real-world emulators to simulate GUI changes from user actions, enabling App
 126 agents to navigate complex scenarios and improve decision-making accuracy. However, emulators
 127 face significant drawbacks, including the safety risks from real-world interactions, such as repeatedly
 128 sending messages or making purchases. To overcome these, world models have gained attention as a
 129 more efficient alternative, not only in agents (Chae et al., 2024; Gu et al., 2024), but also in broader
 130 domains such as robotics (Li et al., 2025; Zhou et al., 2024) and self-driving (Hu et al., 2022).
 131

132 2.2 WORLD MODEL

134 By observing the real world, world models can predict how the environment evolves in response to
 135 an action (LeCun, 2022; Ding et al., 2024). For instance, GameNGen (Valevski et al., 2024) predicts
 136 how a game system will respond to user actions. Notably, the ability to anticipate potential outcomes
 137 of actions has proven to be highly beneficial in informing decision-making processes (Pascanu et al.,
 138 2017; Yang et al., 2024; Schriftwieser et al., 2020; Hafner et al., 2019). Inspired by their success,
 139 world models have emerged to predict the next observation on websites. These models (Chae et al.,
 140 2024; Gu et al., 2024; Liu et al., 2023) typically take a website observation and an action as inputs to
 141 generate a textual description of the next observation. While websites provide multiple sources of
 142 information, including the actual site and their CSS or HTML source files, mobiles present a more
 143 limited context, as only the GUIs are typically accessible. Moreover, text-only descriptions **such**
 144 **as Burns et al. (2024)** often lack the precise visual details required for accurately predicting future
 145 observations, highlighting the need for a visual world model capable of generating high-fidelity future
 146 GUI images.

147 2.3 GUI GENERATION

149 With the rapid advancements in image generation techniques (Rombach et al., 2022; Kumari et al.,
 150 2023; Cao & Gong, 2024), previous methods have explored generating GUI directly in pixel space.
 151 For instance, layout generation methods generate the location of GUI elements (Lu et al., 2023; Zheng
 152 et al., 2023; Sobolevsky et al., 2023; Zhao et al., 2019), scene-text generation methods generate
 153 text that aligns with the visual context (Chen et al., 2024c; Zhang et al., 2024c; Chen et al., 2024b;
 154 Zeng et al., 2024), UI-diffuser (Wei et al., 2024) fine-tune a stable diffusion model to generate
 155 mobile GUIs conditioned on text prompts. For the next GUI generation conditioned on current GUI
 156 observation and a user action, it seems straightforward to resort to an image-and-text-conditioned
 157 approach (Brooks et al., 2023). However, we find that pixel-based image generation struggles with
 158 rendering text accurately, as even minor pixel prediction errors can lead to distortions, particularly for
 159 small-sized text (see Fig. 1 for examples).

160 In this work, we advance beyond existing approaches that generate GUI entirely (graphic and text)
 161 at the pixel level. Instead, we render graphics as image pixels and generate text as language tokens,
 enabling a more accurate method for GUI generation.




Figure 2: Framework of our ViMo. We first detect text content (actual words) in the current GUI (x_k) and overlay it with text symbols (rectangle-shaped placeholders with a black border and white fill), to create STR_{x_k} . Then STR_{x_k} and the user action (a) are input to the STR predictor to generate the STR of the next GUI ($\text{STR}_{x_{k+1}^a}$). Next, text symbols within $\text{STR}_{x_{k+1}^a}$ are located and assigned unique ID token. Then the LLM predicts the text content corresponding to each token. Finally, the next GUI image is constructed by overlaying the predicted text into the STR.

3 METHOD

In this section, we first define our setup in Subsection 3.1. Then, we introduce ViMo in Subsection 3.2. Finally, we demonstrate how ViMo can be applied to enhance existing App agents in real-world scenarios (Subsection 3.3). [All the prompts in this section are listed in Appendix H.1](#).

3.1 PROBLEM SETUP

In general, a GUI world model processes a given GUI Image x_k at step k , and user action a , to predict the effect of a on x_k and simulate the next GUI. Formally, this can be expressed as:

$$x_{k+1}^a = f(x_k, a), \quad (1)$$

where $f(\cdot)$ represents the world model, and x_{k+1}^a denotes the predicted next ($k+1$) GUI image after applying a to x_k . In the following, we explain in detail of our world model.

3.2 ViMo: GENERATIVE VISUAL GUI WORLD MODEL

To tackle the limitation of existing methods (Wei et al., 2024; Chen et al., 2024c; Brooks et al., 2023) in generating visually plausible text for a GUI, as shown in Fig. 1, we propose ViMo, a novel generative visual GUI world model that decouples the graphic and text content generation. As shown in Fig. 2, we first detect and remove all the text in the GUI by overlaying it with a text symbol to create the Symbolic Text Representation (STR). Then a STR predictor is leveraged for determining the STR representation of the next GUI with a pixel-based diffusion process. Finally, a GUI-text predictor is proposed to generate the text content for each symbol using an LLM, followed by a handcrafted design to overplay the text into the STR image to create the next GUI. Their details are specified in the following.

3.2.1 STR: SYMBOLIC TEXT REPRESENTATION

To develop a GUI prediction model that eliminates the need to generate specific text content, we propose the Symbolic Text Representation (STR), where all the text content (actual words) within the GUI image is symbolised (overlaid) with uniform text symbols (placeholders). To be specific, we create an STR representation from a given GUI image with three steps: 1) using an OCR model (Shi et al., 2016; Qiao et al., 2020) to detect text within the GUI; 2) masking the detected text by overlaying it with a box filled with white and bordered in black; 3) we leverage an LLM to filter out static text displayed on static GUI elements and preserve it in the image, as it does not involve any semantic evolution or dynamic changes and remains unchanged as part of specific elements such as a keyboard

216 Algorithm 1 Enhancing App Agent with Generative Visual GUI World Model

217 **Input:** Current GUI Observation x_k , A goal g , A visual world model ViMo, A selection model
 218 $S(\cdot)$.
 219 **Output:** Action to be applied on x_k to achieve g .
 220 Generate action options \mathcal{A} with n actions $\{a^i\}$ (Eq. (4)).
 221 **for** $i = 1$ **to** n **do**
 222 Leverage ViMo to synthesise the next GUI observation conditioned on a^i and x_k , denoted as
 223 $x_{k+1}^{a^i}$ (Eq. (5)).
 224 **end for**
 225 Use $S(\cdot)$ to identify the optimal action with predicted observation (Eq. (6)).
 226
 227
 228

229 or a clock face. Additionally, we empirically find that predicting this static text with complex spatial
 230 patterns poses significant challenges for the LLM.

231 Through the above process, GUI images are transformed into the STR representation, where the text
 232 content is abstracted into a text symbol, relaxing the task of generating semantic text content into
 233 predicting text symbols that indicate the location and size.

235
 236 3.2.2 STR PREDICTOR
 237

238 Building on the powerful generative capability of diffusion-based models (Rombach et al., 2022),
 239 we introduce a STR predictor specifically trained to understand a given STR and a user action,
 240 enabling it to generate the corresponding next STR effectively. In particular, we fine-tune a pre-
 241 trained stable diffusion model (Rombach et al., 2022) to predict the next STR, conditioned on the
 242 STR of the current GUI and the user action. Given a STR representation (STR_{x_k}) extracted from
 243 GUI (x_k), the process starts with the encoding of STR_{x_k} into a latent representation (Kingma &
 244 Welling, 2013): $z = \mathcal{E}(\text{STR}_{x_k})$. Gaussian noise is then added to this representation to create z_t at
 245 timestep t . A denoising autoencoder is subsequently trained to predict the Gaussian noise in the latent
 246 representation, aiming to reverse the noise addition. The objective is defined as:
 247

$$L = \mathbb{E}_{\mathcal{E}(\text{STR}_x), \epsilon \sim \mathcal{N}(0, I), t} \left[\|\epsilon - \epsilon_\theta(z_t, \mathcal{E}(\text{STR}_{x_k}), t, a)\|_2^2 \right], \quad (2)$$

248 where ϵ_θ is a U-Net (Ronneberger et al., 2015) architecture conditioned on a timestep t , a text prompt
 249 a (action), the visual input z_t and the image condition STR_{x_k} . To support the condition on images,
 250 we follow IP2P (Brooks et al., 2023) to add additional input channels to the first convolutional layer,
 251 concatenating the image condition $\mathcal{E}(\text{STR}_{x_k})$ with the noised latent z_t . After training, our STR
 252 predictor is capable of synthesising the next STR ($\text{STR}_{x_{k+1}^a}$) for STR_{x_k} with action instruction a .
 253
 254

255 3.2.3 GUI-TEXT PREDICTOR
 256

257 Given a STR representation generated by our STR predictor, we design a GUI-text predictor to
 258 generate plausible text for the text symbols in the STR based on its graphics. Specifically, we first
 259 locate the text symbols in the STR by colour matching and boundary detection. This outputs the
 260 location of text symbols, along with their unique ID tokens assigned via enumeration, denoted as \mathcal{T} .
 261 Then we leverage the image processing and task understanding ability of LLM to predict the text
 262 content based on its context in STR. This process can be formulated as:
 263

$$\text{text}_{x_{k+1}^a} = \text{LLM}(\text{STR}_{x_{k+1}^a}, x_k, a, \mathcal{T}), \quad (3)$$

264 where $\text{STR}_{x_{k+1}^a}$ denotes the STR representation of x_{k+1}^a . $\text{text}_{x_{k+1}^a}$ contains the predicted text content
 265 for each text symbol, associated with its ID token. This design ensures flexible and accurate
 266 text generation tailored to the predicted GUI STRs as the context. Finally, we overlay each text
 267 content ($\text{text}_{x_{k+1}^a}$) to $\text{STR}_{x_{k+1}^a}$ to reconstruct the predicted GUI image (x_{k+1}^a). To be specific,
 268 text symbols are replaced with the corresponding text based on coordinates, with dynamic styling
 269 determined by the symbol's size and surrounding colours. [More details are provided in Appendix A.1.](#)

270 Table 1: GUI quality evaluation. s_{gc} indicates the GUI consistency, s_{ia} instructional accuracy, s_{ar} action
 271 readiness score and s_h the harmonic average between the 3 metrics. Δs_h is the relative
 272 performance gains of our ViMo over other methods. IP2P* denotes finetuning of IP2P on our dataset.
 273

Method	Automatic Metric					User Study				
	s_{gc}	s_{ia}	s_{ar}	s_h	Δs_h	s_{gc}	s_{ia}	s_{ar}	s_h	Δs_h
HTML-vision	0.70	85.77	62.79	0.72	5.39%	0.31	11.32	9.01	0.23	282.61%
IP2P*	0.74	63.57	70.15	0.69	10.20%	0.82	58.92	52.81	0.63	39.68%
UI-diffuser	0.60	39.61	38.75	0.44	71.82%	0.36	14.32	8.56	0.27	225.93%
ViMo (Ours)	0.74	75.39	78.68	0.76	-	0.89	91.12	84.71	0.88	-

3.3 VI MO ENHANCED APP AGENT

Motivated by that App agents usually face limitations in long-horizon planning to make optimal decisions on action selection (Chae et al., 2024), we leverage the proposed ViMo to enhance the decision-making of App agents.

To be specific, we break down the process into three steps: action option generation, action outcome synthesis, and action selection. In the first step, the App agent generates n action options, as follows:

$$\mathcal{A} = \text{Agent}(x_k, g), \quad (4)$$

where $\mathcal{A} = \{a^1, a^2, \dots, a^n\}$ denotes the action option set, x_k is the current GUI image at step k , and g is the given user goal. With these action options, our world model ViMo is leveraged to synthesise the outcome (next GUI) of these actions as follows:

$$x_{k+1}^{a^i} = \text{ViMo}(x_k, a^i), \quad (5)$$

where $x_{k+1}^{a^i}$ denotes the synthesised next GUI of applying action a^i on x_k . Finally, each action a^i and its corresponding predicted outcome $x_{k+1}^{a^i}$ are fed into an LLM-based selection model, which identifies the optimal action based on the generated GUIs. This process can be formulated as:

$$a_{se} = S\left(\{(a^i, x_{k+1}^{a^i})\}_{i=1}^n\right), \quad (6)$$

where a_{se} denotes the selected action, and $S(\cdot)$ is the selection model. This procedure is outlined in Algorithm 1. By predicting the next GUI, ViMo provides the agent with the potential outcome of an action, enabling it to make more informed decisions. We build the selection model to identify the best action in two steps. First, we query an LLM to evaluate all the action candidates, providing a judgment—either *valid* or *invalid*—and a confidence score for each action. Second, we query the LLM again to select the best action from the two highest-scoring actions. This process is motivated by our observation that, in over 70% of tasks, the difference between the top two scores is equal to or less than 0.1, indicating that both are likely optimal. By explicitly prompting the LLM to compare the top candidates, we go beyond coarse scoring and enable more detailed decision-making.

4 EXPERIMENTS

In this section, we begin by summarising our proposed STR dataset discussed in Subsection 3.2.1. Next, we tested the core capability of our ViMo, focusing on its GUI generation ability. Building on this, we demonstrated how the powerful GUI generation capability of ViMo can enhance the decision-making of App agents. Then, we studied our effectiveness in real-world App navigation tasks. Finally, we carried out the ablation study to validate the effectiveness of our model design. Specific setups and experiment details are elaborated in subsequent sections. Unless explicitly stated otherwise, GPT-4o (Hurst et al., 2024) was employed as the default LLM in the following sections.

4.1 DATASET SUMMARISATION

Our STR dataset was constructed using data from two widely recognised and large-scale sources: Android Control (Li et al., 2024) and Android in the Wild (AITW) (Rawles et al., 2024b). From these sources, we respectively sampled 12 and 7 Apps, selecting those with rich data samples while filtering out noise. Android Control provides two types of user actions: 1) *action commands*:




Figure 3: GUI generation comparison in graphic generation (Top) and text generation (Bottom).

predefined actions (e.g., click, scroll) accompanied by specific parameters such as coordinates (x, y); 2) *action instructions*: actions described in natural languages, such as "click the plus icon". We used action instructions as conditions for our world model as this approach was more concrete and better utilised the pre-trained model in understanding natural language. For AITW, action commands were converted into action instructions using GPT-4o (Hurst et al., 2024). In total, we collected 19 Apps with 3,550 episodes, 23,620 images and 18,450 actions. To ensure both time-efficient and cost-efficient experiments, we followed prior App agent (Rawles et al., 2024b) on partial split evaluation. Specifically, we randomly sampled 57 episodes across 19 distinct Apps. [Details of dataset collection, split summarisation, and full-split experiments are provided in Appendices A.3 and B.](#)

4.2 WORLD MODEL ABILITY

We evaluated the GUI generation capability of ViMo by GUI quality evaluation. We included IP2P (Brooks et al., 2023) and UI-diffuser (Wei et al., 2024), both originally designed for image editing and GUI generation. We fine-tuned IP2P on our dataset to generate everything of the GUI, including the text content and the graphic, denoted as IP2P*. We also leveraged an LLM to predict App observations in an HTML format, which were rendered into images, denoted as HTML-vision.

We leveraged 3 evaluation metrics: The *GUI consistency score* (s_{gc}) assessed the visual similarity between the ground truth and the generated next GUI; *Instructional accuracy score* (s_{ia}) determined whether the generated GUI adheres to the user action; *Action readiness score* (s_{ar}) evaluated whether the generated GUI retains valid elements essential for subsequent actions required to achieve the user goal. Both automatic evaluation and user studies were conducted. For the automatic evaluation, we used DINO (Caron et al., 2021) as the visual encoder to compute s_{gc} , and an LLM to evaluate s_{ia} and s_{ar} . For the user study, we invited 70 voluntary participants to complete questionnaires based on 80 GUI samples, generated by all 4 compared methods. For each sample, participants were asked three questions, one each for evaluating s_{gc} , s_{ia} , and s_{ar} . [Details of the prompts for the LLM-based evaluations and the full instructions for the user study are described in Appendix A.4.](#)

As shown in Table 1, ViMo achieved the highest score on the harmonic average of the three automatic metrics, surpassing other methods with an average relative performance improvement of 29.14%. The results of the user study were consistent with the automatic evaluations, where ViMo demonstrated the best performance. Notably, HTML-vision and UI-diffuser performed significantly worse in human evaluation compared to their scores in the LLM-based assessment. This discrepancy likely

378
379
380
Table 2: Performance comparison of app agents on step accuracy. Apps are categorised into "Leisure",
"Work", and "System".

Agent Type	App Agent	Leisure	Work	System	Overall
Language-Based	ER (Li et al., 2024)	31.76	46.15	34.13	34.50
	AutoDroid (Wen et al., 2024a)	35.81	46.15	31.75	35.46
	T3A (Rawles et al., 2024a)	41.22	51.28	42.86	43.13
	T3A + ViMo (Ours)	50.00	58.97	45.24	49.20
Multi-Modality-Based	APP-Agent (Zhang et al., 2023)	43.24	51.28	39.68	42.81
	Mobile-Agent-v2 (Wang et al., 2024a)	43.92	53.85	39.68	43.45
	M3A (Rawles et al., 2024a)	46.62	51.28	43.65	46.01
	M3A + ViMo (Ours)	53.38	53.85	45.24	50.16

390
391
Table 3: Compare World Models.392
393
Table 4: Zero-shot Evaluation.

Modality	World Model	Step Acc.
w/o WM	w/o WM	46.01
Language	Change-text	47.28
	HTML-text	46.65
Vision	HTML-vision	48.89
	UI-diffuser	47.60
	IP2P*	48.56
		50.16

App Agent	LLM	Step Acc.
SeeAct	GPT-4-Turbo	33.9
M3A	GPT-4-Turbo	42.1
ER	Gemini 1.5 Pro	24.4
T3A	Gemini-2.0-Flash	41.4
T3A+ViMo	Gemini-2.0-Flash	46.8
M3A	Gemini-2.0-Flash	44.2
M3A+ViMo	Gemini-2.0-Flash	47.6

402
403
arose because human evaluators perceived the outputs of these methods as visually unrealistic or
functionally incoherent, leading to lower subjective scores in s_{gc} , s_{ia} and s_h .

404
405
406
407
408
409
Qualitative comparisons are presented in Fig. 3, under two scenarios: GUI graphic changes (Top)
and text generation (Bottom, cropped for space efficiency). Experiments revealed that while the
HTML-vision method exhibited greater flexibility in responding to user actions (as shown in the
bottom examples), it failed to produce concrete details necessary for future actions (top). Conversely,
IP2P* generated plausible GUI graphics but lacked flexibility in text content generation (also reflected
by s_{gc} and s_{ia} in Table 1). This trade-off highlighted the superior balance of ViMo.

4.3 WORLD MODEL ENHANCED APP AGENT

411
412
413
This section demonstrates that: 1) ViMo enhanced the performance of App agents in decision-making;
2) ViMo outperformed other world models in enabling App agents to make more accurate decisions.

414
415
416
417
418
419
420
421
422
423
424
425
Comparison with App Agents. In this experiment, we collected 6 LLM-based App agents, which
included three language-based methods: ER, AutoDroid, and T3A, as well as three multi-modality-
based methods: APP-Agent, Mobile-Agent-v2, and M3A. We applied our ViMo into M3A and
T3A following the process in Subsection 3.3. Moreover, we followed the previous works (Rawles
et al., 2024b; Li et al., 2024) to use the step accuracy (the number of correct actions divided by
the number of overall actions) to quantify the model performance. To provide more detailed results, we
categorised the Apps into three groups: "Leisure", "Work" and "System". Table 2 demonstrates that
ViMo was beneficial to the App agent, achieving a relative performance gain of 9.01% for M3A
and 14.07% for T3A. These findings highlighted the effectiveness of our proposed world model in
providing App agents with enhanced decision-making capability. [Additional information about the
categorisation and experiments with more App agents are provided in the Appendix A.4.](#)

426
427
428
429
430
431
Comparison with World Models. We evaluated the ability of ViMo to enhance App agent decision-
making by comparing it against existing world models. In addition to vision-based world models
discussed in Subsection 4.2, we also incorporated two language-based world models (Gu et al.,
2024; Chae et al., 2024), utilising Change-text to generate textual descriptions capturing differences
between consecutive observations and HTML-text to predict App observations in an HTML format.
Then, we applied each world model to M3A App agents. Table 3, together with Appendix Table 9,
illustrates that vision-based methods consistently outperform language-based world models, thereby

432
433 Table 5: Online Evaluation on Android
434 World (Rawles et al. (2024a)).

435 App Agent	436 LLM	437 Task Acc.
436 SeeAct	437 GPT-4-Turbo	438 15.50
437 M3A	438 GPT-4-Turbo	439 25.40
438 M3A	439 Gemini-1.5-Pro	440 22.80
439 T3A	440 GPT-4-Turbo	441 30.60
440 T3A	441 Gemini-1.5-Pro	442 19.40
441 T3A	442 Gemini-2.0-Flash	443 33.19
442 T3A + ViMo	443 Gemini-2.0-Flash	444 40.95

444 reinforcing our motivation for developing visual GUI world models. Moreover, our approach achieves
445 superior performance over existing world models, underscoring its effectiveness and advantage.

447 4.4 REAL-WORLD APPLICATIONS

449 **Practical Deployment.** ViMo was designed to be lightweight and easily deployable. The minimum
450 requirement for deployment is a GPU with 16 GB of memory. Moreover, ViMo was implemented
451 as a plug-and-play API that required only a single function call, making integration straightforward.
452 Inference time on V100 GPU is 8 seconds on a STR image generation and 30 seconds on GUI-text
453 prediction. [We collected and compared the inference time with existing methods in the Appendix C.](#)

455 **Generalisation to New Apps.** Generalisation is a crucial capability for real-world applications. To
456 assess the generalisation performance of our method on new Apps that were unseen during training,
457 we conducted a zero-shot evaluation using data from the Android Control dataset (Li et al., 2024),
458 explicitly excluding Apps encountered during training. As shown in Table 4, ViMo substantially
459 outperformed the baseline and achieved 47.6%, underscoring its robustness and adaptability to novel
460 App environments. [Additional visualisations of unseen scenarios are provided in the Appendix F.](#)

462 **Online Evaluation.** To further demonstrate the effectiveness of ViMo in realistic App navigation
463 scenarios, we conducted an online evaluation using the AndroidWorld dataset (Rawles et al., 2024a),
464 which comprises 116 distinct navigation tasks. Performance was measured using the task success
465 rate (Task Acc.). As illustrated in Table 5, ViMo achieved a notable improvement of 7.76% over the
466 baseline method, highlighting its effectiveness and reliability in real-world settings.

467 4.5 ABLATION STUDY

470 In this section, we ablated on three key components
471 of ViMo: 1) preserving static text within the image
472 to simplify the text generation task; 2) using action
473 instructions instead of action commands as the
474 conditioning input for ViMo; and 3) varying the number
475 of iterations, where each iteration corresponds to one
476 roll-out step into the future during GUI prediction.

477 Firstly, for the challenge of predicting static text from
478 specific GUI elements, such as keyboard, number
479 pad or clock face, which typically did not involve text
480 changes and exhibited complex spatial patterns, we retained static text within the image (Subsection
481 3.2.1). This approach eliminated the need for the LLM to generate such static text while generating
482 in pixels instead. Secondly, we proposed conditioning STR prediction on action instruction rather
483 than action commands (Subsection 4.1). Ablation results are presented in Table 6, where "Static
484 Text" indicates whether static text was retained in the images, and "Action Instr." denotes whether
485 natural language instructions ("√") or abstract action commands ("-") were used as conditioning
486 input to ViMo. The first row indicates the baseline where ViMo was not applied. The table shows that
487 both components contributed significantly to performance improvements across the two App agents,

Table 6: Ablations on preserving static text
and using action instructions, [measured by step accuracy](#).

Static Text	Action Instr.	App Agent	
		T3A	M3A
N/A	N/A	43.13	46.01
√	–	42.81	45.05
–	√	47.28	48.88
√	√	49.20	50.16

Table 7: Ablation on the number of iterations.

Method	Iterations	Step Acc. (%)
T3A	N/A	39.94
	1	46.06
T3A+ViMo	2	46.65
	3	45.05

486 highlighting their critical roles in enabling ViMo to generate high-quality GUIs. [Visual comparison](#)
 487 [examples are provided in the Appendix F for further illustration.](#)

488 Our ViMo predicted future GUI observations, which could be recursively fed back as input to simulate
 489 further into the future. In this ablation study, we varied the iteration number to evaluate how extended
 490 roll-outs impact prediction accuracy. We took Gemini-2.0-Flash (Hassabis & Kavukcuoglu, 2024) as
 491 the LLM in this study. As shown in Table7, performing two iterations yielded the highest accuracy.
 492 However, this also led to increased computational cost. Therefore, we selected one step as a practical
 493 trade-off between performance and efficiency. We also observed a slight decline in performance
 494 at iteration 3 relative to iterations 1 and 2, indicating that extending the prediction horizon did not
 495 necessarily improve agent behaviour. This was likely due to that longer horizons introduced not only
 496 additional foresight but also a greater accumulation of prediction errors, whose detrimental effect
 497 could outweigh the potential benefits. [Further analysis of error accumulation, user examples of ViMo](#)
 498 [with App agent, and comparisons with various world models are provided in Appendices B and F.](#)

500 5 CONCLUSION

501 In this work, we introduced ViMo, a novel generative visual GUI world model designed to predict
 502 App observations in a visual modality, providing a more realistic and concrete approach compared
 503 to contemporary language-based models. To address the unique challenges of GUI generation,
 504 ViMo was equipped with the STR representation to simplify text content generation to text location
 505 prediction by overlaying text content with placeholders and delegating content generation to LLM.
 506 This innovation ensured high visual fidelity and avoided artefacts like distorted or blurred text.
 507 Through extensive experiments, we demonstrated that ViMo generated both visually plausible and
 508 functionally effective GUIs. Notably, ViMo boosted step-wise action prediction accuracy by a relative
 509 performance gain of 14.07%, underscoring its potential to enhance decision-making of App agents.
 510 Furthermore, real-world experiments demonstrated the strong generalisation ability of ViMo to unseen
 511 Apps, along with its robust performance in online navigation tasks under real-time environment
 512 interaction. Together with its superiority over language-based world models, these results highlighted
 513 the value of ViMo in advancing GUI world modelling in visual modality.

514 REFERENCES

515
 516 Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image
 517 editing instructions. In *CVPR*, pp. 18392–18402, 2023.

518 Andrea Burns, Kate Saenko, and Bryan A Plummer. Tell me what’s next: Textual foresight for
 519 generic ui representations. *arXiv preprint arXiv:2406.07822*, 2024.

520 Yu Cao and Shaogang Gong. Few-shot image generation by conditional relaxing diffusion inversion.
 521 In *ECCV*, pp. 20–37. Springer, 2024.

522 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and
 523 Armand Joulin. Emerging properties in self-supervised vision transformers. In *CVPR*, pp. 9650–
 524 9660, 2021.

525 Hyungjoo Chae, Namyoung Kim, Kai Tzu-iunn Ong, Minju Gwak, Gwanwoo Song, Jihoon Kim,
 526 Sunghwan Kim, Dongha Lee, and Jinyoung Yeo. Web agents with world models: Learning and
 527 leveraging environment dynamics in web navigation. *arXiv preprint arXiv:2410.13232*, 2024.

528 Jingxuan Chen, Derek Yuen, Bin Xie, Yuhao Yang, Gongwei Chen, Zhihao Wu, Li Yixing, Xurui
 529 Zhou, Weiwen Liu, Shuai Wang, et al. Spa-bench: A comprehensive benchmark for smartphone
 530 agent evaluation. In *NeurIPS 2024 Workshop on Open-World Agents*, 2024a.

531 Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser:
 532 Diffusion models as text painters. *NeurIPS*, 36, 2024b.

533 Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser-2:
 534 Unleashing the power of language models for text rendering. In *ECCV*, pp. 386–402. Springer,
 535 2024c.

540 Filippos Christianos, Georgios Papoudakis, Thomas Coste, Jianye Hao, Jun Wang, and Kun Shao.
 541 Lightweight neural app control. *arXiv preprint arXiv:2410.17883*, 2024.
 542

543 Jingtao Ding, Yunke Zhang, Yu Shang, Yuheng Zhang, Zefang Zong, Jie Feng, Yuan Yuan, Hongyuan
 544 Su, Nian Li, Nicholas Sukiennik, et al. Understanding world or predicting future? a comprehensive
 545 survey of world models. *arXiv preprint arXiv:2411.14499*, 2024.

546 Zhibin Gou, Zhihong Shao, Yeyun Gong, Yujiu Yang, Minlie Huang, Nan Duan, Weizhu Chen, et al.
 547 Tora: A tool-integrated reasoning agent for mathematical problem solving. In *ICLR*, 2023.

548

549 Yu Gu, Boyuan Zheng, Boyu Gou, Kai Zhang, Cheng Chang, Sanjari Srivastava, Yanan Xie, Peng Qi,
 550 Huan Sun, and Yu Su. Is your llm secretly a world model of the internet? model-based planning
 551 for web agents. *arXiv preprint arXiv:2411.06559*, 2024.

552 Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning
 553 behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.

554 Demis Hassabis and Koray Kavukcuoglu. Introducing gemini 2.0: our new ai model for the agentic
 555 era, December 2024. URL [https://blog.google/technology/google-deepmind/](https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/)
 556 google-gemini-ai-update-december-2024/. Accessed: 2025-05-04.

557

558 Anthony Hu, Gianluca Corrado, Nicolas Griffiths, Zachary Murez, Corina Gurau, Hudson Yeo, Alex
 559 Kendall, Roberto Cipolla, and Jamie Shotton. Model-based imitation learning for urban driving.
 560 *NeurIPS*, 35:20703–20716, 2022.

561 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 562 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*
 563 *arXiv:2410.21276*, 2024.

564

565 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint*
 566 *arXiv:1312.6114*, 2013.

567

568 Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept
 569 customization of text-to-image diffusion. In *CVPR*, pp. 1931–1941, 2023.

570

571 Yann LeCun. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open*
 572 *Review*, 62(1):1–62, 2022.

573

574 Sunjae Lee, Junyoung Choi, Jungjae Lee, Munim Hasan Wasi, Hojun Choi, Steven Y Ko, Sangeun
 575 Oh, and Insik Shin. Explore, select, derive, and recall: Augmenting llm with human-like memory
 576 for mobile task automation. *arXiv preprint arXiv:2312.03003*, 2023.

577

578 Chenhao Li, Andreas Krause, and Marco Hutter. Robotic world model: A neural network simulator
 579 for robust policy optimization in robotics. *arXiv preprint arXiv:2501.10100*, 2025.

580

581 Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbulin, and Bernard Ghanem. Camel: Commu-
 582 nicative agents for "mind" exploration of large language model society. *NeurIPS*, 36:51991–52008,
 583 2023.

584

585 Wei Li, William Bishop, Alice Li, Chris Rawles, Folawiyo Campbell-Ajala, Divya Tyamagundlu,
 586 and Oriana Riva. On the effects of data scale on computer control agents. *arXiv preprint*
 587 *arXiv:2406.03679*, 2024.

588

589 Anthony Liu, Lajanugen Logeswaran, Sungryull Sohn, and Honglak Lee. A picture is worth a
 590 thousand words: Language models plan from pixels. In *EMNLP*, pp. 16450–16459, 2023.

591

592 Yuwen Lu, Ziang Tong, Qinyi Zhao, Chengzhi Zhang, and Toby Jia-Jun Li. Ui layout generation
 593 with llms guided by ui grammar. *arXiv preprint arXiv:2310.15455*, 2023.

594

595 Razvan Pascanu, Yujia Li, Oriol Vinyals, Nicolas Heess, Lars Buesing, Sébastien Racanière, David
 596 Reichert, Théophane Weber, Daan Wierstra, and Peter Battaglia. Learning model-based planning
 597 from scratch. *arXiv preprint arXiv:1707.06170*, 2017.

598

599 Zhi Qiao, Yu Zhou, Dongbao Yang, Yucan Zhou, and Weiping Wang. Seed: Semantics enhanced
 600 encoder-decoder framework for scene text recognition. In *CVPR*, pp. 13528–13537, 2020.

594 Christopher Rawles, Sarah Clinckemaillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth
 595 Fair, Alice Li, William Bishop, Wei Li, Folawiyo Campbell-Ajala, et al. Androidworld: A dynamic
 596 benchmarking environment for autonomous agents. *arXiv preprint arXiv:2405.14573*, 2024a.
 597

598 Christopher Rawles, Alice Li, Daniel Rodriguez, Oriana Riva, and Timothy Lillicrap. An-
 599 droidinthewild: A large-scale dataset for android device control. *NeurIPS*, 36, 2024b.

600 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 601 resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022.

602 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 603 image segmentation. In *MICCAI*, pp. 234–241. Springer, 2015.

604 Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon
 605 Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari,
 606 chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.

607 Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based
 608 sequence recognition and its application to scene text recognition. *TPAMI*, 39(11):2298–2304,
 609 2016.

610 Andrey Sobolevsky, Guillaume-Alexandre Bilodeau, Jinghui Cheng, and Jin LC Guo. Guiget: Gui
 611 layout generation with transformer. *arXiv preprint arXiv:2304.09012*, 2023.

612 Dani Vavlevski, Yaniv Leviathan, Moab Arar, and Shlomi Fruchter. Diffusion models are real-time
 613 game engines. *arXiv preprint arXiv:2408.14837*, 2024.

614 W. contributors. Large language model — Wikipedia, the free encyclopedia. [Online]. Available:
 615 [urlhttps://en.wikipedia.org/wiki/Large_language_model](https://en.wikipedia.org/wiki/Large_language_model), 2024. Accessed: 2024-11-25.

616 Junyang Wang, Haiyang Xu, Haitao Jia, Xi Zhang, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang,
 617 and Jitao Sang. Mobile-agent-v2: Mobile device operation assistant with effective navigation via
 618 multi-agent collaboration. *arXiv preprint arXiv:2406.01014*, 2024a.

619 Taiyi Wang, Zhihao Wu, Jianheng Liu, Jianye Hao, Jun Wang, and Kun Shao. Distrl: An asyn-
 620 chronous distributed reinforcement learning framework for on-device control agents. *arXiv preprint*
 621 *arXiv:2410.14803*, 2024b.

622 Jialiang Wei, Anne-Lise Courbis, Thomas Lambolais, Gérard Dray, and Walid Maalej. On ai-inspired
 623 ui-design. *arXiv preprint arXiv:2406.13631*, 2024.

624 Hao Wen, Yuanchun Li, Guohong Liu, Shanhui Zhao, Tao Yu, Toby Jia-Jun Li, Shiqi Jiang, Yunhao
 625 Liu, Yaqin Zhang, and Yunxin Liu. Autodroid: Llm-powered task automation in android. In
 626 *MobiCom*, pp. 543–557, 2024a.

627 Hao Wen, Shizuo Tian, Borislav Pavlov, Wenjie Du, Yixuan Li, Ge Chang, Shanhui Zhao, Jiacheng
 628 Liu, Yunxin Liu, Ya-Qin Zhang, et al. Autodroid-v2: Boosting slm-based gui agents via code
 629 generation. *arXiv preprint arXiv:2412.18116*, 2024b.

630 Chang Yang, Xinrun Wang, Junzhe Jiang, Qinggang Zhang, and Xiao Huang. Evaluating world
 631 models with llm for decision making. *arXiv preprint arXiv:2411.08794*, 2024.

632 Weichao Zeng, Yan Shu, Zhenhang Li, Dongbao Yang, and Yu Zhou. Textctrl: Diffusion-based scene
 633 text editing with prior guidance control. *arXiv preprint arXiv:2410.10133*, 2024.

634 Chaoyun Zhang, Shilin He, Jiaxu Qian, Bowen Li, Liqun Li, Si Qin, Yu Kang, Minghua Ma, Qingwei
 635 Lin, Saravan Rajmohan, et al. Large language model-brained gui agents: A survey. *arXiv preprint*
 636 *arXiv:2411.18279*, 2024a.

637 Chi Zhang, Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu.
 638 Appagent: Multimodal agents as smartphone users. *arXiv preprint arXiv:2312.13771*, 2023.

639 Jiwen Zhang, Jihao Wu, Yihua Teng, Minghui Liao, Nuo Xu, Xiao Xiao, Zhongyu Wei, and Duyu
 640 Tang. Android in the zoo: Chain-of-action-thought for gui agents. *arXiv preprint arXiv:2403.02713*,
 641 2024b.

648 Lingjun Zhang, Xinyuan Chen, Yaohui Wang, Yue Lu, and Yu Qiao. Brush your text: Synthesize any
649 scene text on images via diffusion model. In *AAAI*, volume 38, pp. 7215–7223, 2024c.
650

651 Bo Zhao, Lili Meng, Weidong Yin, and Leonid Sigal. Image generation from layout. In *CVPR*, pp.
652 8584–8593, 2019.

653 Guangcong Zheng, Xianpan Zhou, Xuewei Li, Zhongang Qi, Ying Shan, and Xi Li. Layoutdiffusion:
654 Controllable diffusion model for layout-to-image generation. In *CVPR*, pp. 22490–22499, 2023.
655

656 Siyuan Zhou, Yilun Du, JiaBen Chen, Yandong Li, Dit-Yan Yeung, and Chuang Gan. Robodreamer:
657 Learning compositional world models for robot imagination. *arXiv preprint arXiv:2404.12377*,
658 2024.

659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 In this Appendix, we first provide detailed explanations, including prompts related to our methods,
 703 descriptions of our STR dataset, and evaluation details. Then, we present additional experimental
 704 results. Finally, we present additional visualisations of our proposed ViMo for GUI generation.
 705

706 A EXPERIMENTAL DETAILS

709 A.1 GUI-TEXT PREDICTOR

710 This subsection elaborates on the design and functionality of the GUI-text predictor, summarising its
 711 key components and providing a detailed explanation of its underlying processes.
 712

713 Given a STR prediction, the GUI-text predictor starts by locating the text symbols. To be specific, we
 714 first detect black borders by identifying black pixels in the BGR colour space, generating a binary
 715 mask that indicates whether a pixel is black or not. A pixel is classified as black if its BGR values
 716 fall within the range $[0,0,0]$ to $[50,50,50]$. Next, we identify rectangular regions within this mask
 717 by computing the ratio of the actual contour area to its corresponding bounding rectangle area. If
 718 this ratio exceeds 0.8, the region is considered a valid rectangle, allowing us to extract rectangles
 719 with black borders. For these detected regions, we further analyse their internal colour distribution
 720 to determine whether they contain the desired white colour. Specifically, we define white pixels
 721 as those with BGR values within the range $[200,200,200]$ to $[255,255,255]$. If more than 50% of
 722 the pixels within a region fall within this range, the region is classified as a text symbol. Thus, the
 723 locations of text symbols are extracted, and we assign a unique identifier (ID) to each symbol through
 724 enumeration.

725 Building on this, we take as inputs the current GUI image x_k , an action a to be applied to this image,
 726 the predicted STR ($\text{STR}_{x_{k+1}^a}$), the location and unique ID token of the text symbols in the STR \mathcal{T} .
 727 Then we leverage an LLM to predict the text content for each text symbol. The process begins with
 728 preprocessing the STR by overlaying the ID token for each text symbol to the corresponding position
 729 in the STR image, resulting in a modified representation denoted as STR_{k+1}^{ID} . Next, we prompt an
 730 LLM to identify which text symbols will remain unchanged after the action a (see the prompt in
 731 Subsection H.1). These symbols are determined to not be affected by the action and have content
 732 identical to the previous GUI x_k . Based on the resulting ID list, we retrieve the corresponding pixels
 733 from the previous GUI x_k based on their location and update the STR representation. The updated
 734 image is still referred to as STR_{k+1}^{ID} for simplicity.

735 Subsequently, the LLM is prompted to determine the semantic role of each text symbol by analysing
 736 its context (see the prompt in Subsection H.2). This semantic information, combined with STR_{k+1}^{ID} ,
 737 is then used to predict the exact text content of each symbol (see the prompt in Subsection H.3).

738 Finally, to overlay a symbol with its actual text content, we perform the following steps: 1) For a
 739 given text symbol’s location and corresponding text, the average background colour is computed
 740 by the average colour of the area on the edge of text symbol’s coordinates; 2) The text colour is set
 741 to either white or black to ensure optimal contrast with the background colour, for better visibility;
 742 3) The font size is calculated as the maximum size that allows the text to fit entirely within the
 743 boundaries of the text symbol, ensuring optimal use of space and readability.

744 A.2 ACTION SELECTION

746 In practice, our selection model, described in Section 3.3, identifies the best action in two steps.
 747 First, we query an LLM to evaluate all the action options, providing a judgment—either *valid* or
 748 *invalid*—and a confidence score for each action (see the prompt in Subsection H.4). These judgments
 749 are transformed into scores: if an action is judged *valid*, its score equals the confidence; if judged
 750 *invalid*, its score is the confidence multiplied by -1 . This scoring reflects that higher confidence
 751 in a *valid* action yields a higher score, while higher confidence in an *invalid* action results in a
 752 lower (negative) score. Second, we query the LLM again to select the best action from the two
 753 highest-scoring actions (see the prompt in Subsection H.5). This step is motivated by our observation
 754 that, in over 70% of tasks, the difference between the top two scores is equal to or less than 0.1,
 755 indicating that both are likely optimal. By allowing the LLM to choose between them, we refine the
 selection beyond simply picking the action with the highest score.

756 Table 8: Summarisation of our STR dataset.
757

758 Split	759 App	760 Episode	761 Image	762 Instrucion
759 Train	760 19	761 2853	762 19010	763 14852
760 Val	761 19	762 349	763 2290	764 1774
761 Test	762 19	763 348	764 2320	765 1824
762 All	763 19	764 3550	765 23620	766 18450

763 Table 9: Decision optimisation comparisons on APP agent performance. Apps are categorised into
764 "Leisure", "Work", and "System".
765

766 App Agent	767 World Model Modality	768 World Model	769 Leisure	770 Work	771 System	772 Overall
773 T3A	w/o world model	w/o world model	41.22	51.28	42.86	43.13
	Langugae	Change-text	49.32	51.28	42.06	46.65
		HTML-text	47.30	48.72	43.65	46.01
		HTML-vision	50.68	53.85	43.65	48.24
	Vision	UI-diffuser	48.65	53.85	43.65	47.28
		IP2P	48.65	53.85	45.24	47.92
		ViMo (Ours)	50.00	58.97	45.24	49.20
774 APP-Agnet	w/o world model	w/o world model	43.24	51.28	39.68	42.81
	Langugae	Change-text	45.96	56.41	45.24	46.96
		HTML-text	44.59	56.41	45.24	46.33
		HTML-vision	47.97	56.41	46.03	48.24
	Vision	UI-diffuser	47.30	56.41	44.44	47.28
		IP2P	47.30	58.97	45.24	47.92
		ViMo (Ours)	50.68	58.97	43.65	48.89
775 Mobile-Agent-v2	w/o world model	w/o world model	43.92	53.85	39.68	43.45
	Langugae	Change-text	47.30	66.67	41.27	47.28
		HTML-text	47.30	66.67	38.89	46.33
		HTML-vision	50.00	66.67	41.27	48.56
	Vision	UI-diffuser	49.32	61.54	41.27	47.60
		IP2P	46.62	66.67	45.24	48.56
		ViMo (Ours)	50.00	66.67	44.44	49.84
783 M3A	w/o world model	w/o world model	46.62	51.28	43.65	46.01
	Langugae	Change-text	51.35	51.28	41.27	47.28
		HTML-text	50.68	51.28	40.48	46.65
		HTML-vision	52.03	48.72	45.24	48.89
	Vision	UI-diffuser	50.00	48.72	44.44	47.60
		IP2P	52.03	48.72	44.44	48.56
		ViMo (Ours)	53.38	53.85	45.24	50.16

788
789 A.3 DATA COLLECTION

790 To ensure the quality and diversity of data samples for each App, while minimising noise, we
791 collected App information from both Android Control (Li et al., 2024) and Android in the Wild
792 dataset (AITW) (Rawles et al., 2024b) datasets. To be specific, out of 15,274 episodes in the Android
793 Control, only 5,697 episodes include the "open_app" action. From these episodes, we extracted their
794 "app_name", identifying 758 unique applications. However, only 13 of these Apps had more than 50
795 samples. To enrich the dataset, we manually collected additional samples for these 13 Apps from
796 the rest of the dataset. For AITW, we extracted App names by using the package name listed under
797 the "current activity" field. After filtering out the noisy, 11 valid Apps remained. By combining the
798 overlapping applications from both datasets, we obtained a total of 19 unique Apps. We split our
799 dataset into "Train", "Validation" and "Test" splits, and we summarise our dataset under each split in
800 Table 8.
801

802 Furthermore, we converted action commands into action instructions for AITW with specific prompts
803 in Subsection H.6. We use Paddleocr (Shi et al., 2016) for STR generation.
804

805 A.4 EVALUATION
806

807 **World Model Ability.** For the results under automatic metrics presented in Table 1, we prompt LLM
808 for the instructional accuracy score s_{ia} and action readiness score s_{ar} , as shown in Subsection H.7
809 and Subsection H.8 respectively. A generation is considered successful if "success" appears under
810 "Status" for s_{ia} and "yes" under "ready for action" for s_{ar} . For the user study, we collected 80

810 Table 10: Trajectory synthesis evaluation. "T+L" denotes the accuracy of the whole trajectory with
 811 length L.

813 World Model	814 T+1	815 T+2	816 T+3	817 T+4
818 w/o world model	819 22.81	820 14.04	821 7.02	822 0
823 Change-text	824 52.63	825 26.32	826 10.53	827 5.26
828 HTML-text	829 38.60	830 14.04	831 12.28	832 7.02
834 HTML-vision	835 43.86	836 19.30	837 10.53	838 10.53
841 UI-diffuser	842 52.63	843 29.82	844 12.28	845 5.26
848 IP2P*	849 56.14	850 21.05	851 10.53	852 7.02
855 ViMo (Ours)	856 57.89	857 36.84	858 14.03	859 12.28

822 Table 11: Evaluation on randomness by running the experiment 3 times (r1-r3) on our sampled test
 823 split. "All" denotes the evaluation of the full test split. s_{gc} , s_{ia} and s_{ar} are the metrics same with
 824 Table 1. s_h denotes their harmonic score. STD denote the standard deviation from r1 to r3.

825 World Model	826 s_{gc}	827 s_{ia}	828 s_{ar}	829 s_h
830 r1	831 0.7421	832 75.08	833 78.29	834 0.7582
835 r2	836 0.7323	837 75.63	838 77.64	839 0.7546
840 r3	841 0.7423	842 75.39	843 78.68	844 0.7605
845 STD	846 0.0057	847 0.23	848 0.42	849 0.0025
850 ALL	851 0.7389	852 75.37	853 78.20	854 0.7578

833 generated samples—20 from each of the four world models. We then asked 70 participants to
 834 answer three questions on each sample designed to reflect the s_{ia} , s_{gc} and s_{ar} scores, as detailed in
 835 Subsection H.9. For the s_{gc} , participants are asked to rate on a scale from 1 to 5. These scores were
 836 then normalised to the [0,1] range in Table 1.

837 **World Model Enhanced App Agent.** In Table 2, we categorised APPs based on their primary
 838 functions into three groups: **Leisures**, **Work**, and **System**. The **Leisure** category includes APPs
 839 commonly used for relaxation and entertainment, such as *Decathlon*, *eBay*, *Flipkart*, *Amazon*, *Adidas*,
 840 *Kitchen Stories*, *Booking.com*, *YouTube*, and *Vimeo*. The **Work** category comprises APPs typically
 841 associated with professional or productivity-related activities, including *Gmail*, *Drive*, and *Chrome*.
 842 Lastly, the **System** category encompasses APPs pre-installed in the Android operating system, such
 843 as *com.android.contacts*, *com.google.android.dialer*, *com.google.android.googlequicksearchbox*,
 844 *com.android.settings*, *com.google.android.APPs.maps*, and *com.android.vending*.

846 **Ablation on Iteration Numbers.** ViMo predicts future GUI observations, which can be recursively
 847 fed back as input to simulate further into the future. Taking the generative GUI as the current GUI,
 848 an agent was prompted to generate the action instructions based on the user goal (see the prompt in
 849 Subsection H.10). Then the action instruction and the GUI were fed into ViMo to generate the next
 850 GUI. In this study, we defined the iteration number as the number of times ViMo was called. We
 851 only use the final output as the signals during the candidate action selection phase, guiding the final
 852 selection among potential actions.

853 B ADDITIONAL EXPERIMENTAL RESULTS

856 **Comparison with World Models.** Table 3 compares our ViMo with existing world models under
 857 M3A App agent. To further highlight our superiority, Table 9 presents additional results of ViMo
 858 applied to T3A, APP-Agent, and Mobile-Agent-V2. The experimental results indicate that vision-
 859 based methods consistently outperform their language-based counterparts, thereby substantiating
 860 our motivation for developing visual GUI world models. Moreover, our approach achieves superior
 861 performance over existing world models, underscoring its effectiveness and advantage.

863 **Generation Error Analysis.** As discussed in Subsec. 4.5, our method ViMo can iteratively generate
 864 future GUIs. However, as the number of iterations increases, the accumulated error also grows. In




Figure 4: Qualitative ablation studies. Left: Static text generation. (a) Generating static text via an LLM; (b) Preserving the original text in the image by rendering it as image pixels. Right: STR generation under two input formats—(a) action command and (b) action instruction.




Figure 5: GUI generation conditioned on a novel combination of current GUI observation and user action.

addition to the evidence presented in Table 7, we conduct further experiments to analyse this iteration error and compare our approach with existing world models.


To this end, we design a trajectory synthesis evaluation to assess how well the GUIs generated by ViMo align with those observations in real-world environments over longer iterations. In this setup, the generated GUI is leveraged as the input to an App agent to generate the subsequent action, with higher-quality trajectories indicating a GUI more aligned with the real-world environment. Specifically, the GUIs generated by ViMo serve as the observation input for the App agent, which generates actions aimed at achieving the user’s goal. These output actions are then evaluated to reflect whether the GUI representations offer concrete and reliable information for action prediction. This process is repeated for L steps, and we calculate the success rate of the entire L -step trajectory.

We employ an LLM as a judge to assess the alignment between the agent’s simulated actions and the ground truth actions within a given trajectory. Specifically, an agent was prompted to generate the action instructions based on the given GUI and the user goal (see the prompt in Subsection H.10) and the LLM evaluated whether the simulated action lead to the same outcome as the ground truth action (see the prompt in Subsection H.11), a "yes" of the "Status" is calculated as a match.

As shown in Table 10, we compared ViMo against both visual- and language-based world models and demonstrated that while performance decreases across all world models with more iterations, our model significantly outperformed the other methods by providing more accurate and reliable information. This was reflected in higher trajectory prediction accuracy, underscoring the ability of our model to generate GUIs that aligned with the real-world environment.

918
 919 Table 12: Comparison of inference runtime and step accuracy across different models. “Execution
 920 Time” denotes the per-module execution time for the vision and text components. The end-to-end
 921 time includes model execution, initialisation, loading, and communication overhead. “–” indicates
 922 that the corresponding module is not applicable.

923 Model	924 Execution Time (Vision / Text)	925 End-to-End Time	926 Step Accuracy (%)
927 Baseline (T3A)	928 – / –	929 4 min	930 43.13
931 Change-text	932 – / 5s	933 5 s	934 46.64
935 IP2P*	936 8s / –	937 1.5 min	938 47.92
939 ViMo (Ours)	940 8s / 30s	941 2 min	942 49.20



944 Figure 6: False examples where the text symbols are incorrectly represented, making them unrecog-
 945 nisable to indicate the location of text.

946
 947
 948 **Randomness Study and Evaluation on Full Test Split.** ViMo involves random factors, particularly
 949 from the use of LLMs. The LLM is prompted to generate plausible textual content, and in some
 950 cases, multiple reasonable options can be produced. For example, in Fig. 3, it shows "5:49" on the
 951 top left corner for "set timer for 20 minutes" command and shows "Timer" for "set timer for 12
 952 minutes", both are plausible and valid in the given context. However, the key functional element,
 953 the timer itself, is consistent with the user instructions in both cases. To evaluate their influence, we
 954 conducted the experiment three times, as summarised in Table 11 (r1-r3). The results demonstrate
 955 that the randomness does not significantly impact the performance or consistency of our method.
 956 Additionally, we focused on a randomly selected subset of examples for evaluation, with results
 957 from the full test set also included to illustrate that the observed differences are minor, as shown in
 958 Table 11 (compare ALL to r1-r3). We consider the subset results to provide an accurate and reliable
 959 approximation for our analysis.

960
 961 **Qualitative Ablation Analysis.** In addition to the quantitative ablation results presented in Table 6,
 962 we also provide qualitative comparisons. Fig. 4 (left) illustrates the challenges faced by the LLM in
 963 predicting static text under complex spatial layouts. Fig. 4 (right) displays the STR generation of
 964 the same user intent but with different action types. It demonstrated that models learned with action
 965 commands failed to predict STR that aligns with the user’s intent, whereas action instructions offered
 966 a more concrete description, enabling the model to better capture the intent.

967
 968 **Qualitative Generalisation Study.** We studied the generalisation of ViMo in Fig. 5 by providing
 969 user actions that were not typically encountered within the App’s standard context. For example,
 970 in the Clock App, a user action to "add a file" generated a Drive-style file selection window while
 971 retaining the Clock interface. Similarly, in the Kitchen Store App, ViMo can generate content
 972 corresponding to the action. These results emphasised ViMo’s generalisation ability facing novel
 973 combinations of App observations and user actions.




Figure 7: Example of how ViMo helps the App agent to select the correct action.

C PRACTICAL DEPLOYMENT

In this section, we report the computational efficiency of our method to demonstrate its practicality in real-world applications. The minimum hardware requirement is a GPU with 16 GB of memory. As shown in Table 12, with a V100 GPU, STR image generation (Vision) takes approximately 8 seconds, and GUI-text prediction (Text) takes around 30 seconds. These runtimes are practical and we have demonstrated an online inference setting in Table 5.

Although ViMo introduces additional latency, this is an intentional design choice that is closely aligned with our objective of improving the reliability of long-horizon decision making. For example, when the planning horizon is extended to 4 steps, our ViMo pipeline achieves 2.33 \times higher accuracy (12.28% vs. 5.26% in Table 10) compared to Change-text and 1.75 \times higher to IP2P*. We view the ability to construct a visual world model that can accurately predict the consequences of actions as a critical foundation for building robust GUI agents. While this design involves a higher computational cost, we believe this trade-off is justified by the substantial gains in reliability and planning accuracy, especially when the horizon is longer.




Figure 8: Visualisation of ViMo in generating GUIs given a single current GUI paired with different actions.




Figure 9: Visualisation of ViMo in generating the next GUI. For each example, the action is displayed at the top, with the current GUI shown on the left and the generated GUI on the right.

D ADDITIONAL DISCUSSION

D.1 USER PRIVACY PRESERVATION

We emphasise that our method neither collects sensitive user data nor has the potential to leak such information.

For GUI graphic generation, the diffusion model is trained on our STR representation (GUI graphic), in which all textual content is removed from the GUI. The model is trained solely to predict visual and layout information. As a result, no sensitive or personally identifiable information is included in the training data or generated by the model.




Figure 10: Visualisation of ViMo in generating the next GUI. For each example, the action is displayed at the top, with the current GUI shown on the left and the generated GUI on the right.

For text rendering, no model fine-tuning is performed. At test time only, an LLM (such as ChatGPT) is prompted to generate textual content based on contextual information such as the previous GUI state, the user action, and the user goal. Sensitive user information that is not displayed on the screen is never seen or accessed. Sensitive information that may appear on screen (e.g., contact lists or account details) is not collected, stored, or reused for training or any other purpose.

Therefore, no private user information is collected, stored, or leaked during either the training or inference stages.

D.2 FAIRNESS COMPARISONS

To evaluate whether the visual modality provides more precise information than the language modality for app agent tasks, we conduct controlled comparisons under a non-finetuned setting. As shown in Tables 3, 9, and 10, our two visual baselines (HTML-vision and UI-diffuser), which are not fine-tuned on our task, consistently outperform the language-based methods (Change-text and HTML-text), which are also not fine-tuned. This result reflects the inherent advantage of visual representations for modelling fine-grained GUI details.

1134 To evaluate the necessity of our ViMo design, we also finetune the vision-based baseline IP2P on our
 1135 task using the same training data as ViMo. This ensures a fair comparison and further verifies that the
 1136 observed improvements stem from our model architecture rather than advantages in the training data.
 1137

1138 E LIMITATION AND FUTURE WORKS

1140 Fig. 6 illustrates failure cases where text symbols are not represented as our rectangle-shaped
 1141 placeholders with a black border and white fill, making them unrecognisable as text symbols.
 1142 Improving the representation of text symbols remains a potential direction for future work.
 1143

1144 **Beyond visual quality, computational efficiency remains a practical limitation. Improving the**
 1145 **efficiency of robust GUI agents is a critical research direction for enabling real-world deployment.**

1146 F ADDITIONAL VISUALISATION

1149 Fig. 7 demonstrates how ViMo helps the M3A App agent make better action decisions. Fig. 8 show-
 1150 cases results generated from a single current GUI paired with different actions, further highlighting
 1151 the versatility of our approach. Diverse visualisations are presented in Fig. 9 and Fig. 10. These
 1152 examples illustrate how our method effectively generates the next GUI based on the given action
 1153 and current GUI observation, showcasing its ability to produce visually coherent and contextually
 1154 accurate GUI simulations.

1155 G STATEMENT ON LLM USAGE

1158 We disclose that large language model (LLM) tools were used solely for language refinement of the
 1159 manuscript, including improving grammar and polishing phrasing. LLMs were not used to generate
 1160 scientific content, research ideas, experiment designs, data, analyses, or code. All suggestions and
 1161 modifications from these tools were made under the direct supervision and final approval of the
 1162 authors, and all authors are fully aware of and consent to this usage.

1163 H PROMPTS

1166 H.1 PROMPT TO DECIDE THE TEXT SYMBOLS TO REMAIN UNCHANGED AFTER THE ACTION

```
1168 You are a professional UI/UX analyst and your goal is to compare the two UI screenshots
1169 and return their overlapping layout.
1170 ### Inputs:
1171 1. **Current Screenshot**: The current mobile UI as an image.
1172 2. **Next UI Layout Screenshot**:
1173 - An image of the next mobile UI layout with all text replaced by white boxes.
1174 - Each box has a unique red ID label.
1175 3. Use action: a user action described by language
1176 Next UI Layout Screenshot is a result of a user action on the current screenshot, but
1177 the text elements are masked.
1178 Please help me identify those layouts that are located in the same position, so I can
1179 predict their text directly from the current screenshot.
1180 Usually, the system bar information should be included. Exclude elements from the
1181 result if:
1182 The content (text) changes as a result of the user action, even if the element exists
1183 in both screenshots.
1184 Please be very very cautious about putting an ID on the list, which means you are very
1185 very confident with this task. If you are unsure about some elements, please ignore
1186 them and do not put them on the list.
1187 ### Output the list of existing elements : Return the results in the following JSON
1188 format: ['id1','id2',...]
1189 ### Notes:
1190 - Ensure the detected elements appear in both UI screenshots, which means their
1191 surrounding context is the same.
1192 - Ensure identify those elements that their text will change by the user action and
1193 exclude them from your response.
1194 - Ensure identify those elements that share a similar context layout, but their absolute
1195 are not the same, and them from your response.
1196 - Ensure only reply in pure JSON format, with no placeholders or comments.
```

1188
1189

H.2 PROMPT TO DETERMINE THE SEMANTIC ROLE OF EACH TEXT SYMBOL

1190

1191
1192

You are a professional UI/UX analyst assigned to structure and analyse the semantics of mobile UI screenshots.

1193

Your goal is to segment the UI and annotate box elements in a way that enhances understanding of their roles and relationships within the interface. Inputs:

1194

- Current Screenshot: A visual representation of the mobile UI.

1195

- Next UI layout screenshot: A visual representation of the next UI layout with all the text masked with a white box. Each box has an ID number on it in red colour.

1196

- User Action: An action put on the current UI will result in the next UI.

1197

- Box locations: a list of box locations to better help you to locate the boxes in the format of 'id': id, 'Location':[x1,y1, width, height]. ID indicates their ID number in the UI screenshot.

1198

- UI_size: the width/height of the input images. They are the same size. The image you received might be resized. Please scale it back for the locations.

1199

Task:

1200

Structure the boxes in the Next UI layout screenshot with semantics based on the visual input by following these steps:

1201

1. Divide the UI into Semantic Windows Group the UI into functional sections with a specific name (e.g., "Header Windows," "Time Selector Panel").

1202

2. Structure Text Elements in Each Semantic Window.

1203

- Assign box elements to windows based on logical, visual relationships or semantic roles.

1204

- For every element, structure output as :

1205

- **id: corresponding box retrieved from the box list and the Next UI layout screenshot.

1206

- **Role: A brief explanation of the role of this box. You should consider their [x1,y1] to indicate their location, [w,h] to indicate their size to decide the role. It is important to consider the context for the role prediction. The role should be in detail to distinguish it from other items in the same category.

1207

- Output Format: { "Window Name": "Category Name": ["id":id, "Role": "Role" ,

1208

- "id":id, "Role": "Role" , ...], "Category Name": ["id":id, "Role": "Role" , ...

1209

-]

1210

- ... }

1211

Key Guidelines:

1212

- Ensure to retrieve id from the given screenshot and box list.

1213

- Avoid duplicating or omitting IDs.

1214

- Every box element in the box location list must be included in the structured output.

1215

- Ensure there is no additional formatting, code blocks or placeholders in your response; return only a clean JSON without any comments.

1216

1217

1218

1219

H.3 PROMPT TO PREDICT THE EXACT TEXT CONTENT FOR EACH SYMBOL

1220

1221

1222

Task: Plan the content for the next UI screen based on the provided inputs and instructions.

1223

Inputs:

1224

Current Screenshot: A visual representation of the mobile UI.

1225

Next UI layout screenshot: A visual representation of the next UI layout with all light yellow boxes indicating a text place. Each box has an ID number on it.

1226

User Instruction: A specific action or command that transitions the current UI to the next UI state.

1227

Semantics for the masks in Next UI screenshot: A structured map.

1228

Goal:

1229

Predict the content (text) for each masked area in the next UI layout screenshot based on the following steps:

1230

Map Affected Elements to the Next UI.

1231

Align the affected elements with the yellow box coordinates on the next UI.

1232

Predict the text for each yellow box based on the user instruction and the context of the current UI.

1233

If you can not find any information about the text, predict a plausible text based on its context.

1234

Ensure to use the semantics to help you understand the layout and predict the text. If you think the semantics is wrong, please modify it in your

1235

Output:

1236

Return the predictions in JSON format with the structure: { "Window Name " : "Category Name " : ["id " : id, "text " : "text", "role " : "role" , "id " : "id", "text " : "text", "role " : "role " , ...] }

1237

Ensure to predict text based on the context.

1238

Do not include any special characters.

1239

Ensure there is no additional formatting, code blocks or placeholders in your response; return only a clean JSON without any comments.

1240

1241

1242
1243

H.4 PROMPT TO EVALUATE ACTIONS WITH A CONFIDENCE SCORE

1244

1245

You are an agent who can operate an Android phone on behalf of a user. When given a user request, you will try to complete it step by step. At each step, a list of descriptions for most UI elements on the current screen will be given to you (each element can be specified by an index), together with a history of what you have done in previous steps. Based on these pieces of information and the goal, you must choose to perform one of the actions in the following list (action description followed by the JSON format) by outputting the action in the correct JSON format: action options from the dataset

The overall user goal/request is: {goal}

Here is a history of what you have done so far:{history} This is the action you picked in the latest step: {action}, whose semantic description is: {sum}

Your goal is to judge **whether the action you picked in the latest step is on the right track to the successful execution of the overall user goal/request**.

You will be given the screenshots before and after you perform the action

- The first screenshot corresponds to the UI state before you performed the action.
- The second screenshot corresponds to the UI state after you performed the action.

Also here is the list of detailed information for some UI elements in the before screenshot: {before_elements}

Note that, the "after" screenshot is generated by the agent's world model. As such, it may not faithfully represent the real UI. For instance: Some UI elements in the simulated "after" screenshot may not exist in a real UI. Your evaluation should consider the reliability of the UI predictions. If the "after" screenshot contains unreasonable elements, this likely indicates a failure.

Now provide your judgment on the selected action in JSON format. Your response must include:

Reason: A detailed explanation of why the action is valid or invalid.

Judgment: Your judgment must be either "valid" or "invalid".

Confidence: A confidence score between 0.0 and 1.0, reflects how likely your judgment is correct.

You must follow this structure exactly:

{Reason: ..., Judgement: "valid" or "invalid", Confidence: ...}

Your Answer:

1267

1268

1269

1270

1271

1272

H.5 PROMPT TO SELECT THE OPTIMAL ACTIONS AMONG TWO HIGHEST-SCORING ACTIONS

1273

1274

1275

You are an agent who can operate an Android phone on behalf of a user. When given a user request, you will try to complete it step by step. At each step, a list of descriptions for most UI elements on the current screen will be given to you (each element can be specified by an index), together with a history of what you have done in previous steps. Based on these pieces of information and the goal, you must choose to perform one of the actions in the following list (action description followed by the JSON format) by outputting the action in the correct JSON format action options from the dataset

1281

The overall user goal/request is: {goal}

1282

Here is a history of what you have done so far:{history}

1283

Here is a list of descriptions for some UI elements on the current screen:{before_elements}

1284

Here are two candidate actions:

1285

Action 1: {action_0}, described semantically as {sum_0}. The rationale for this action is: {act_re_0}

1286

Action 2: {action_1}, described semantically as {sum_1}. The rationale for this action is: {act_re_1}

1287

Hints for making your decision: {GUIDANCE}

1288

- Both "more options" buttons and scrolling actions may reveal new content. Evaluate which is more suitable for the goal.

1289

- Consider the history of previous actions. If prior steps involved repeated "scroll down" actions, it is more likely that "scroll down" is the correct next step.

1290

- If the user goal involves viewing reviews or similar tasks and the current screen already displays such content, "scroll down" may reveal more information.

1291

Your task is to choose the best action from the two provided.

1292

Now, provide your judgment in JSON format with the following structure:

1293

Reason: A detailed explanation of your choice, considering the hints above.

1294

Choice: Action 1 or Action 2.

1295

Your output must exactly match this format:

{Reason: ..., Choice: Action 1 or Action 2}

1296
1297

H.6 PROMPT TO CONVERT ACTION COMMANDS INTO ACTION INSTRUCTIONS

1298

1299

You are a professional UI/UX analyst specializing in identifying the semantics of dual point actions between mobile UI screenshots.

Inputs:

Current Screenshot: A visual representation of the mobile UI.

Next Screenshot: A visual representation of the NEXT mobile UI.

Goal: A user intent on this Mobile interface.

touch_xy: the x,y coordinates for the touch point, as a percentage of the image dimensions.

lift_xy: the x,y coordinates for the lift point, as a percentage of the image dimensions.

Your task is to analyse these elements describe the precise user action in plain language and return your answer in plain string (e.g., "click the + icon", "scroll up").

If the two screenshots are identical, please return an empty string as "".

If the Next Screenshot does not seem to be one step away from the Current Screenshot, return an empty string as "". One step means only one interaction with the cell phone. Ensure there is no additional formatting, code blocks or placeholders in your response; return only a clean string without any comments

1311

1312

1313

1314

H.7 PROMPT FOR INSTRUCTIONAL ACCURACY SCORE (s_{ia})

1315

1316

1317

You are an expert in evaluating the performance of a mobile emulator. The mobile emulator is designed to navigate the UI change based on human instruction.

Inputs:

Current UI Screenshot: The present state of the cellphone's user interface.

Next UI Screenshot: The mobile emulator generated UI indicating the next state of the cellphone's user interface based on human instruction.

Human instruction: The action applied on the current UI screenshot.

Your goal is to determine whether the mobile emulator successfully predicts the next UI image with current information and layout based on the current UI and the user action.

IMPORTANT

Format your response into a JSON map as shown below:

```
{
  "Thoughts": <your thoughts and reasoning process>,
  "Status": "success" or "failure",
}
```

1327

1328

1329

1330

1331

H.8 PROMPT FOR ACTION READINESS ACCURACY SCORE (s_{ar})

1332

1333

You are an expert in evaluating the performance of a mobile emulator. The mobile emulator is designed to navigate the UI change based on human instruction.

Inputs:

UI Screenshot: The mobile emulator generated UI indicating the state of the cellphone's user interface.

User intent: The user goal to achieve.

Next action: the action will be applied to this UI.

Your goal is to determine whether the next action is validated on the UI Screenshot.

Please also indicate if it is still in the right App according to the goal.

IMPORTANT Format your response into a JSON map as shown below:

```
{
  "Thoughts": <your thoughts and reasoning process>,
  "In the right App": "yes" or "no",
  "ready for action": "yes" or "yes",
}
```

1343

1344

1345

1346

1347

H.9 INSTRUCTIONS FOR USER STUDY

1348

1349

The following prompt provides the instructions for the user study. An example screenshot is shown in Fig. 11.




Figure 11: Screenshot of user study example.

1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369 Question 1: You are presented with a mobile screen (Current GUI).
 1370 Your task is to evaluate whether the generated GUI correctly reflects the result of the
 1371 user action applied on the current GUI. Answer "Yes" or "No" to each sample.
 1372 Question 2:
 1373 You are presented with Ground Truth Next GUI, the correct screens that should appear
 1374 after the user performs the given action.
 1375 In this task, you will evaluate the visual similarity between the Ground Truth Next GUI
 1376 and AI-generated GUI, scoring from 1-5.
 1377 Question 3:
 1378 In this task, you will validate whether a specific user action is valid on the
 1379 generated GUI.
 1380 An action is considered valid if the required GUI element is visually presented on the
 1381 screen. Answer "Yes" or "No" to each sample.

1384 H.10 PROMPT TO GENERATE THE ACTION INSTRUCTION BASED ON THE GIVEN GUI AND THE 1385 USER GOAL

1386
 1387
 1388 You are an autonomous intelligent agent tasked with navigating a cell phone to
 1389 accomplish specific tasks. You will be provided with the following information:
 1390 1. Initial UI screenshot: A visual representation of the initial state of the cell
 1391 phone's interface.
 1392 2. User Objective: This is the task you are trying to complete.
 1393 3. Previous Action: An action sequence performed on the initial UI.
 1394 4. Current UI states: A visual representation of the current state of the cell phone's
 1395 interface, generated by a simulated environment.
 1396 The initial image is the screenshot before actually performing all the previous
 1397 actions.
 1398 The current cell phone UI is generated by applying previous actions on the initial
 1399 screenshot.
 1400 Your Task: Please predict a single next step action to complete the given task based
 1401 on current vision states.
 1402 To be successful, it is very important to follow the following rules:
 1403 1. You should only issue one action that is valid based on the current UI states.
 1404 2. You should only issue one action at a time. Avoid issuing multiple actions like
 1405 "do A and do B".
 1406 3. Generate the action in plain text. For example, Scroll down to set the minute as
 1407 15.
 1408 4. Issue "Stop." if you think the action is already completed. Ensure you only return
 1409 the action, not other formats, comments or placeholders

1404
 1405 **H.11 PROMPT TO EVALUATES WHETHER THE SIMULATED ACTION LEADS TO THE SAME**
 1406 **OUTCOME AS THE GROUND TRUTH ACTION**

1407 You are an expert in evaluating the performance of a cell phone navigation agent. The
 1408 agent is designed to help a human user navigate a cellphone to complete a task.
 1409 Inputs:
 1410 Current UI Screenshot: The present state of the cellphone's user interface.
 1411 User Intent: The goal the human user aims to achieve.
 1412 Action History: The sequence of actions taken so far for you to track the progress.
 1413 Agent Simulated Action: The action suggested by the agent to achieve the user's
 1414 intent.
 1415 Ground Truth Action: The correct action is needed to achieve the user's intent.
 1416 Your goal is to determine whether the agent's simulated action leads to the same
 1417 outcome as the ground truth action.
 1418 Additionally, if the simulated action does not exactly match the ground truth action
 1419 but is still progressing toward the correct outcome to achieve user intent, indicating
 1420 that the action is "on the right track."
 1421 *IMPORTANT*
 1422 Format your response into a JSON map as shown below:
 1423 {
 1424 "Thoughts": <your thoughts and reasoning process>,
 1425 "Status": "success" or "failure",
 1426 "On the right track to success": "yes" or "no"
 1427 }

1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457