CAN VLMS PLAY ACTION ROLE-PLAYING GAMES? TAKE BLACK MYTH WUKONG AS A STUDY CASE

Anonymous authors

Paper under double-blind review

Abstract

Recently, large language model (LLM)-based agents have made significant advances across various fields. One of the most popular research areas involves applying these agents to video games. Traditionally, these methods have relied on game APIs to access in-game environmental and action data. However, this approach is limited by the availability of APIs and does not reflect how humans play games. With the advent of vision language models (VLMs), agents now have enhanced visual understanding capabilities, enabling them to interact with games using only visual inputs. Despite these advances, current approaches still face challenges in action-oriented tasks, particularly in action role-playing games (ARPGs), where reinforcement learning methods are prevalent but suffer from poor generalization and require extensive training. To address these limitations, we select an ARPG, "Black Myth: Wukong", as a research platform to explore the capability boundaries of existing VLMs in scenarios requiring visualonly input and complex action output. We define 13 tasks within the game, with 76.9% focusing on combat, and incorporate several state-of-the-art VLMs into this benchmark. Additionally, we will release a human operation dataset containing recorded gameplay videos and operation logs, including mouse and keyboard actions. Moreover, we propose a novel VARP (Vision Action Role-Playing) agent framework, consisting of an action planning system and a human-guided trajectory system. Our framework demonstrates the ability to perform basic tasks and succeed in 90% of easy and medium-level combat scenarios. This research aims to provide new insights and directions for applying multimodal agents in complex action game environments. The code and datasets will be made available at https://varp-agent.github.io/.

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

030

031

1 INTRODUCTION

037 In recent years, LLM-based agents have achieved significant breakthroughs across various fields 038 (Deng et al., 2023; Gur et al., 2024; He et al., 2024; Wang et al., 2024a; Zhang et al., 2023b), particularly with the integration of tools and memory modules (Zhou et al., 2024), as seen in AutoGPT and Reflection (Yang et al., 2023; Shinn et al., 2023). Among these, applying LLM-based agents in 040 video games has become one of the most popular areas of research.(Qian et al., 2024; Park et al., 041 2023; Li et al., 2024; Wang et al., 2023b; 2024b) These methods input information from video games 042 into LLMs, which then undergo complex reasoning and integration through agent frameworks, ulti-043 mately producing keyboard and mouse commands that can directly interact with the game to com-044 plete tasks. Previous works have mostly focused on accessing video game APIs to read in-game 045 environmental and action information. For instance, the framework proposed by Wang et al. (Wang 046 et al., 2023a) has been successfully applied in the game Minecraft. Agents can achieve autonomous 047 mining, building, and attacking enemies in the game. However, this approach does not align with 048 how humans play games, and most games do not offer open APIs, which limits the widespread application of this method. Recently, the emergence of vision language models (VLMs) like GPT-40 has further enhanced the visual understanding capabilities of these agents, showcasing broader po-051 tential in mobile apps and games. For example, the Cradle framework (Tan et al., 2024) has been implemented in Red Dead Redemption 2 (RDR2). It directly uses game screenshots from RDR2 as 052 input, rather than using an API to read game memory information. However, Cradle relies heavily on text-based guiding information in the game screenshots to create new skills. For tasks or games

Agents	Agents		Game	Game Type	Environment	
Reflexion (Shinn	et al., 2023)	LLM-based	ALFWorld	Text-based Adventure	API	
ReAct (Yao et al	., 2023)	LLM-based	ALFWorld	Text-based Adventure	API	
Voyager (Wang e	et al., 2023a)	LLM-based	Minecraft	Sandbox	API	
CreativeAgent (2	Zhang et al., 2023a)	VLM-based	Minecraft	Sandbox	API and Screen	
Cradle (Tan et al	., 2024)	VLM-based	RDR2	AAA Action Adventure	Screen	
DQN (analogand	ldigital, 2021)	RL-based	Sekiro	AAA Action Role-Playing	Screen	
Other Project (C	at, 2024; fange, 2024)	RL-based	BMW	AAA Action Role-Playing	Screen	
VARP (Ours)		VLM-based	BMW	AAA Action Role-Playing	Screen	

Table 1: Comparison of several existing agents. Among them, "API" refers to the model's use
 of video game APIs to access in-game environmental and action information, whereas "Screen"
 indicates that visual understanding is derived solely from game screenshots.

with weak textual guidance, such as some action role-playing games(ARPG), Cradle is unable to
leverage the effective performance of VLMs. For ARPGs, many researchers employ reinforcement
learning methods, where penalties and rewards are predefined for specific tasks. After extensive
training periods and numerous iterations, the trained agents can complete the given specific tasks.
However, RL-based agents can only accomplish tasks within the environment they were trained in
and find it challenging to transfer to other tasks. ARPGs contain a large number of specialized tasks,
which pose a significant challenge for RL-based agents with poor generalization capabilities. We
conducted a comparison of some representative methods in Tab. 1.

073 Thus, most of the existing research focuses on relatively simplified settings. This simplification 074 arises primarily from two significant challenges: 1) Immediate visual input. Since environmen-075 tal data is not always accessible through game APIs, learning from visual input becomes a more 076 straightforward strategy, especially in AAA games (characterized by A lot of time, A lot of re-077 sources, A lot of money), where understanding the immediate visual input is crucial. 2) Actionoriented tasks. Action games are immensely popular among players; however, in this domain, RL-based agents still dominate, which require extensive training time and have poor generalization 079 ability. For VLM-based agents, the game interfaces of ARPGs provide very few textual hints, and most of the actions need to be learned through experience and self-innovation. As a result, previous 081 agents have found it challenging to extract effective guidance information from visual inputs.

In this paper, we will select the "*Black Myth: Wukong*," abbreviated as BMW, an AAA ARPG,
as our research platform for extensive experimentation. We are dedicated to establishing a VLMbased agent framework to thoroughly investigate the capability boundaries of existing models (e.g.
GPT-40, Gemini) in scenarios requiring visual-only input and complex action output. Among them,
visual-only input refers to the model making decisions solely by understanding and analyzing the
game screenshot, while complex action output necessitates the model to perform intricate and continuous actions, such as precise operations in combat scenarios.

To achieve this goal, we define 12 tasks in the game "Black Myth: Wukong," with 75% of these tasks being combat-related. Several state-of-the-art VLM models, including GPT-40, will be incorpo-091 rated into this benckmark to comprehensively explore their performance boundaries. Subsequently, 092 to advance the development of VLM-based agents in AAA action games, we will open-source a human operation dataset, which includes records of mouse and keyboard commands as well as 094 gameplay recordings. Lastly, we innovatively propose a VARP(Vision Action Role-Playing) agent 095 framework, consisting of an action planning system and a human-guided trajectory system. Specif-096 ically, the action planning system is responsible for generating action combos that are suitable for 097 combat scenarios, while the human-guided trajectory system learns from human data via retrieval. 098 Through extensive evaluations, our proposed framework demonstrates the capability to accomplish basic tasks such as picking up items and opening treasure chests, while also succeeding in 90% of 099 esay and medium battles. We hope this research will provide new insights and directions for the 100 application of multi-modal agents in complex action game environments. The main contributions of 101 this paper are summarized as follows: 102

103 104

105

106

- **Benchmark**. We define 12 tasks based on the game "*Black Myth: Wukong*," with 75% of these tasks focused on combat. Several state-of-the-art VLM models, including GPT-40, will be incorporated into this benckmark to thoroughly explore their capability boundaries.
- **Dataset**. We release a dataset containing recorded gameplay videos along with relevant operation logs, including mouse movements, clicks, and keyboard actions, which includes 1000 records.

Framework. We propose a VARP agent framework, which comprises an action planning system and a human-guided trajectory system. Through these systems, the agent can execute complex action combos and learn from human operation.

2 RELATED WORK

112

113

129

130

131

132

133

134

135

136

154

155 156

157

114 115 2.1 LLM AND VLM-DRIVEN AGENTS

116 In recent years, various intelligent agents driven by large language models (LLM) and multimodal 117 language models (VLM) have gradually come to the forefront, demonstrating immense potential 118 in multitasking and autonomous learning. For LLM, Reflexion (Shinn et al., 2023) enhances the 119 decision-making ability of language agents through a framework of linguistic feedback and self-120 reflection, allowing agents to autonomously reflect in the face of feedback signals and maintain contextual memory during task execution and decision-making processes. ReAct (Yao et al., 2023), 121 on the other hand, emphasizes real-time information retrieval and strategy adjustment. By combin-122 ing reasoning with action and interacting with external knowledge sources (such as the Wikipedia 123 API), it adds dynamic information retrieval capabilities, providing greater interpretability and con-124 trollability. Voyager (Wang et al., 2023a) can explore and learn skills in an unsupervised manner 125 within the Minecraft environment, continuously exploring the world, acquiring diverse skills, and 126 making new discoveries without human intervention through a combination of automated courses, 127 skill libraries, and iterative prompting mechanisms. 128

For VLM, CreativeAgent (Zhang et al., 2023a) focuses on creative tasks, employing multimodal generation to achieve the construction of complex structures. Its combination of an imagination module and controller enables efficient planning and execution based on free-form language instructions and the generated task details. Cradle (Tan et al., 2024) takes video images displayed on a screen as input, extracting text and visual information to make decisions through a workflow of "reflecting on the past, summarizing the present, and planning for the future," outputting control signals for keyboard and mouse interaction, allowing AI agents to interact with software like humans without relying on any internal APIs.



Figure 1: **Pipeline of VARP.** We propose a novel framework named the VARP agent, which directly takes game screenshots as input and generates keyboard and mouse operations to play the ARPG.

3.1 OVERVIEW

We propose a novel framework named the VARP agent, which directly takes game screenshots as input. Through inference by a group of Vision-Language Models (VLMs), it ultimately generates actions in the form of Python code, which can directly operate the game character. Each action is a sequence that consists of various combinations of atomic commands. These atomic commands include light attack, dodge, heavy attack, restore health, and others. Meanwhile, the VARP agent

maintains three libraries: a situation library, an action library, and a human-guided library. These
libraries can be retrieved and updated to store intensive knowledge for self-learning and human
guidance. Overall, the VARP agent is divided into two systems: the action planning system and the
human-guided trajectory system, as shown in Fig. 1. In the action library, "def new_func_a()" represents
the new action generated by the action planning system. while "def new_func_h()" represents
the new action generated by the human-guided trajectory system. "def pre_func()" represents
the new action generated by the human-guided trajectory system. "def pre_func()" represents
the new actions. The following sections will elaborate on each system in detail.

169 170

171

3.2 ACTION PLANNING SYSTEM

The action planning system is primarily used for action reasoning and generation. This system 172 utilizes a situation library and an updatable action library as knowledge retrieval bases. Guided 173 by input game screenshots, the system employs a group of VLMs to select or generate actions 174 appropriate for the current situation. The generated situations and actions are stored or updated in 175 the two libraries. Additionally, We propose decomposable task-specific auxiliary modules to break 176 down large tasks into smaller subtasks, which are then distributed across multiple VLMs to reduce 177 the occurrence of model forgetting and hallucinations. We also introduce a self-optimizable action 178 generation module to encourage VLMs to generate new actions specific to some hard tasks, thereby 179 completing complex tasks more efficiently and with higher quality.

- 180
- 181 3.2.1 BASIC VLMS GROUP182

183 Inspired by Cradle (Tan et al., 2024), our main pipeline continues to adopt the five basic modules from Cradle, with some of these basic modules calling the VLM for reasoning, forming a basic 184 group of VLMs. During initialization, we manually predefined some actions and placed them into 185 the action library as the prior knowledge. Each action is a Python function with detailed textual annotations, and we computed the embeddings of these annotations for storage. Information Gath-187 ering is responsible for gathering information from sampled game screenshots, including textual and 188 visual information related to situations and actions. The textual information primarily includes text 189 guides, text labels, and notifications; the visual information mainly covers environmental positions, 190 character actions, and interface icons. The former is assisted by OCR tools for text recognition, while 191 the latter uses the object detection tools for visual localization. Self Reflection takes a few game 192 screenshots from the last video in the situation library as input to assess whether the last executed 193 action successfully produced the correct effect and whether the current task has been completed. If execution fails, the module needs to provide a reason for the failure to guide the next step in action 194 generation. Task Inference infers the current task to be executed based on the results of previous 195 modules, and generates the task description. Skill Curation calculates the similarity between the 196 task description's embedding and the embeddings of the textual annotations in the action library 197 to find some matching actions, which form the candidate action set. Decision Making utilizes the Chain of Thought (CoT) (Wei et al., 2023) approach to reason through and deeply analyze multiple 199 sequential questions (such as whether to enable combat mode, restore health, or select from avail-200 able spell skills, etc.). Finally, the module infers the most suitable action from the candidate action 201 set, executes the Python code, and operates the keyboard and mouse to control the player character 202 to complete the corresponding task. These five basic modules will record each intermediate product 203 into the situation library.

204 205

206

3.2.2 Self-Optimizable Action Generation Module

The basic VLMs group can only acquire actions from a predefined action library or from game 207 screenshots with clear textual prompts. For certain tasks in ARPGs that have weak textual guidance, 208 such as real-time combat, this method is unable to learn new actions. Therefore, we propose a 209 self-optimizable action generation module (SOAG) that allow the VARP agent to summarize the 210 enemy's actions during combat, thereby optimizing existing actions and generating new ones to 211 counter enemy attacks. The new actions are combinations of the two atomic commands: dodging 212 and light strikes. The optimization goal is to maximize the evasion of enemy attacks and the ability 213 to strike the enemy while minimizing the player character's health loss. 214

215 Specifically, in SOAG, we introduce a component responsible for action function generation. This component takes the information gathering and self reflection results, along with the current and

216 last game screenshots, as input. It analyzes the characteristics of the enemy under the current task, 217 such as name, appearance, weapon, etc. Most importantly, it needs to analyze the enemy's current 218 and previous actions. For example, for the hard enemy named Bullguard, its attack actions can be 219 roughly categorized as: "charging forward with the axe", and "chopping the axe downwards three 220 times consecutively", etc. Therefore, this component needs to inference new actions for dodging and counterattacking based on the current enemy actions. For instance, for "charging forward with the axe," the new action should be to dodge once and then attack continuously; for "chopping the axe 222 downwards three times consecutively," the new action should be to dodge three more times before 223 counterattacking. The generated new actions are permutations of the atomic operations "dodge" and 224 "light attack." The generated actions are stored in the action library with detailed textual annotations. 225

226 227

228

3.2.3 DECOMPOSABLE TASK-SPECIFIC AUXILIARY MODULES

In ARPGs, especially in BMW game, the VLM's inference involves a large number of tokens, including multiple images and long texts. The attention mechanism used by VLMs allocates attention to all tokens in long texts. As the input length increases, the attention distribution becomes increasingly diluted. In the basic VLMs group, due to the excessive number of input tokens for each module, the model may fail to effectively focus on key information, leading to errors such as forgetting and hallucination. This issue is particularly evident in the decision-making module, where the VLM frequently makes mistakes when answering multiple questions.

236 To address this problem, we decomposed the basic modules and added multiple parallel auxiliary 237 sub-modules for specific tasks, which are then integrated by the VLM. The structure is similar to an 238 MLP. Specifically, as shown in the workflow of action planning in Fig. 1, we decomposed the origi-239 nal decision-making module that handled multiple tasks into 5 sub-modules. 1) Enemy Sub-module 240 is used to analyze the enemy's status (such as its health, position, etc.) and action description, which 241 assists the agent in obtaining detailed information of the enemy. 2) Combat Sub-module determines 242 which combat method to use, including light attack or heavy attack, by observing the heavy-attack status in the bottom right corner of the game screen. 3) Health Sub-module is responsible for con-243 stantly monitoring the player's health bar. If the health is consumed excessively, it assists the agent 244 by prioritizing the action of recovering health. 4) Spell-skill Sub-module monitors the status of the 245 player's spell skills while simultaneously analyzing the situation in the combat state to determine 246 the appropriate time to use available spell skills. 5) The integration sub-module is responsible 247 for integrating the outputs of all sub-modules and reasoning to determine the best action from the 248 candidate action set for the current specific task. The decomposable task-specific auxiliary modules 249 decompose long tokens and focus on each individual question, significantly improving the accuracy 250 of the decision-making module.

251 252 253

254

3.3 HUMAN-GUIDED TRAJECTORY SYSTEM

255 Human actions are seen as valuable data, implicitly rich in knowledge of the physics and game 256 world, which can lead to advanced action combinations for very complex tasks, such as way-finding 257 tasks and high-difficulty combat tasks. To learn the human experience from this implicit data, we 258 first collected a human dataset and then used it to improve the performance of our VARP agent. The 259 collection process of human operation data and dataset analysis will be detailed in Sec. 3.1, which 260 consists of mouse keyboard logs, and recording game screenshots. In this section, we focus on how to use it to implement a human-guided trajectory system. In this section, we refer to our annotated 261 dataset as the human-guided library. It is a collection of pairs consisting of game screenshots and 262 human operations, with each pair having a unique timestamp. 263

For very hard tasks in the game, we first take a screenshot of the current game interface. Based on this game screenshot, we query the human-guided library for the screenshot with the highest similarity. We then input this screenshot along with the subsequent n-frame screenshots and their corresponding operations into the human-guided trajectory system. This system will utilize a VLM to analyze and summarize the input images and operations, ultimately outputting a new humanguided action, which is then stored in the action library for the action planning system to choose and execute.



Figure 2: Image examples of defined tasks.

4 EXPERIMENTS

4.1 DATASET COLLECTION

303 We collected a human operation dataset that includes mouse and keyboard logs, as well as recordings 304 of game screenshots. Specifically, we recruited 200 volunteers to play the BMW game and record 305 their operations, with approximately 70% of them experiencing this game for the first time. To 306 ensure the dataset's quality, we eliminated invalid data from volunteers who did not complete the 307 tasks. Over the course of two weeks, we ultimately gathered a total of 1,000 valid data entries. Specifically, over 90% of tasks 11 and 12 (i.e., Defeat Guangzhi and Defeat Wandering Wight) were 308 discarded, indicating that defeating these bosses in a single attempt poses a significant challenge for 309 players. Moreover, we observed that volunteers exhibited redundant actions during the annotation 310 process, such as excessive mouse clicks and scrolling. Therefore, some volunteers will be asked to 311 play the game again to identify the optimal actions, and this refined data will be labeled as "clean" in 312 our released dataset. Please refer to the supplementary material for more details about our dataset.

313 314 315

297

298 299

300 301

302

4.2 BENCHMARK AND TASK DEFINITION

316 To investigate the capabilities of existing VLMs in playing action games, we define 10 basic tasks 317 and 3 challenging tasks aligned with the game's narrative, with 76.9% of these tasks occurring in 318 combat scenes. As illustrated in Tab. 2 and Fig. 2, all tasks are concentrated in the first chapter 319 of the game, due to the limited understanding and reasoning abilities of VLMs. In terms of bench-320 marking, we allow the agent to test each task 5 times and calculate the success rate for each task. 321 For combat tasks, a task is deemed successful if the player's character defeats the enemy, while a task is considered a failure if the player's character is defeated and killed by the enemy. We have 322 manually assessed the difficulty of 13 tasks, categorizing them as easy, medium, hard, and very 323 hard. Due to the absence of maps and guidance, and the presence of numerous "invisible walls" in

Task ID	Task Name	Description	Diffuculty	
1	Guidance	Defeat Erlang, the Sacred Divinty	Easy	
2	Combat 1	Defeat WolfScout	Easy	
3	Gather	Gather	Easy	
4	Combat 2	Defeat WolfStalwart	Easy	
5	Combat 3	Defeat WolfSwornsword	Easy	
6	Open	Open	Easy	
7	Combat 4	Defeat WolfSoldier	Easy	
8	Combat 5	Defeat Croaky	Easy	
9	Combat 6	Defeat Crow Diviner	Middle	
10	Combat 7	Defeat Bullguard	Hard	
*11	Combat 8	Defeat Guangzhi	Very Hard	
*12	Combat 9	Defeat Wandering Wight	Very Hard	
*13	Move	Autonomous Navigation	Very Hard	

Table 2: Task definitions in *Black Myth: Wukong* (BMW), where "*" indicates the challenging task.

the BMW game, we classify task 13, autonomous navigation (i.e., moving from the spawn point to the Bullguard's location within five minutes), as a very hard task. This is a challenging task even for human novices. We utilize the success rates from this benchmark to evaluate the performance of the VARP agent and various VLMs.

4.3 IMPLEMENTATION DETAILS

346 All evaluations are performed on a machine equipped with an NVIDIA RTX 4090 GPU running 347 the Windows operating system. We use three of the most popular VLMs to drive our agent: GPT-4o-2024-05-13(OpenAI, 2024), Claude 3.5 Sonnet(Anthropic, 2024), and Gemini 1.5 pro(Google, 348 2024). We also utilize OpenAI's text-embedding-ada-002(Sonnet, 2022) model to generate embed-349 dings for each action. The size of game interface for the BMW game is set to 1920×1080 . During 350 the inference of VLMs, we pause the game using the photo mode. We employ Grounding DINO(Liu 351 et al., 2024) for object detection of people and objects in game screenshots to assist the VLMs in 352 better extracting useful information. 353

354 355

324

340

341

342

343 344

345

4.4 PERFORMANCE EVALUATION

To evaluate the performance of the VARP agent without human guidance, we conducted experiments on our proposed benchmark while disabling the human-guided trajectory system of the VARP agent. In this performance evaluation, we only tested the benchmark and compared the VARP agent with human novice players.

We calculated the success rates of the VARP agent driven by different VLMs and human novice 361 players when completing each task. As shown in Fig. 3, both the VARP agent and human novice 362 players achieved high success rates on tasks 1 to 8, reaching nearly 100% on most tasks. In task 363 9, the VARP agent's average success rate was 40%, which also confirms its "middle" difficulty. 364 The enemy in task 10 is the first boss-level monster that the player encounters in the game. For human novice players, the success rate for this task was 15.63%, while the VARP agent's average 366 success rate was 20%. Task 11 and 12 are classified as "very hard," so the success rates for both 367 human novices and the VARP agent were very low. Specifically, the VARP agent is limited by the 368 reasoning speed of VLMs, making it unable to input every game frame in real-time and only able to input keyframes at second-level intervals. In ARPGs, this can easily result in missing critical 369 information about enemy attacks. Therefore, task 11 and 12 are particularly challenging for the 370 agent. In terms of autonomous navigation, humans can easily find the final boss enemy of the level 371 within five minutes, but for VLMs, this is an almost impossible task. Without human guidance, the 372 success rate is 0%. Since the game provides no guidance or hints for navigation tasks and contains 373 many "invisible walls," VLMs lack the ability to perceive the correct path in the 3D scene without 374 human assistance. 375

To explore the performance boundary of the VLM-based VARP agent, we chose a very-hard task and used the health of the enemy "Guangzhi" in task 11 as a metric. Due to the constraints of GPT-40's maximum token count and inference time, we were unable to input all the video frames, which led to the VARP agent being unable to defeat the highly aggressive "Guangzhi". However, during
the combat, the VARP agent was capable of reducing the enemy's health by an average of about
40%, which may have reached the ability boundary of VLM for purely visual ARPG gameplay.

In summary, the VARP agent's performance on tasks 1 through 12 is already close to that of novice
 human players. However, in terms of 3D scene perception and prior knowledge, the VARP agent is
 still far inferior to humans.



Figure 3: Evaluation results on various VLMs and human.

4.5 COMPARISION EVALUATION

385

386

387

388

389 390

391

392 393

394

396

397

399

414

420

To ensure a fair comparison, we adapted Cradle (Tan et al., 2024) to the BMW game. Specifically, we
redesigned the predefined skills in Cradle to align with those of the VARP agent. We also revised the
prompts for each module in Cradle to focus on combat tasks for the BMW game. Additionally, we
used photo mode to pause the game. The VLM used in all experiments is GPT-40. The experimental
results are shown in Tab. 3.

405 In combat tasks, the success rate of VARP shows a significant improvement compared to Cradle, 406 especially in medium and difficult tasks such as tasks 9-10. During the experiments, we found 407 that Cradle tends to input a large amount of textual prompts into the VLM all at once during the 408 decision-making process. This often leads to decision-making errors by the VLM, such as attempt-409 ing to recover health when at full health, or forgetting to use spell skills when they are available. Furthermore, Cradle can only generate new actions based on textual prompts from game screen-410 shots. However, in ARPG games, complex actions are a combination of various atomic operations, 411 and there are no explicit textual prompts available for them. As a result, Cradle fails in medium and 412 difficult tasks due to its inability to learn new actions. 413

Table 3: Comparison results with Cradle.

Method	1	2	3	4	5	6	7	8	9	10
Cradle (Tan et al., 2024)	100%	80%	100%	60%	60%	100%	40%	100%	20%	0%
VARP	100%	100%	100%	100%	80%	100%	100%	100%	60%	40%

4.6 ABLATION STUDY

421 To evaluate the effectiveness of the self-optimizable action generation module(SOAG) and the de-422 composable task-specific auxiliary module(DTSA) in the action planning system, we conducted 423 experiments with each of these modules removed separately, calculating the success rate on the 424 benchmark. The VLM used in this part of the experiment is GPT-40-2024-05-13. As shown in 425 Fig. 4, without SOAG, the agent's performance significantly declines in the middle and hard tasks. 426 This is because the enemies in these tasks have high health points, resulting in prolonged battles. The 427 function of SOAG is to continuously learn the enemies' attack patterns, aiding players in dodging 428 and counterattacking. Therefore, in long-duration tasks like middle and hard tasks, the effectiveness 429 of SOAG becomes more apparent. On the other hand, DTSA aims to decompose large tasks into smaller ones, focusing more on precision. This approach helps prevent global errors caused by local 430 issues such as the forgetting and hallucination of the VLM. Hence, without DTSA, the agent tends 431 to fail in some easy tasks.

100 Full pipeline w/o DTSA 80 w/o SOAG Success Rate (%) 60 40 20 0 Defeat WolfScout Defeat WolfStalwart Defeat WolfSoldier Defeat Bullguard Defeat Erlang Defeat WolfSwornsword Defeat Croaky Defeat Crow Diviner Gather Open



4.7 CASE STUDIES

In this section, we will showcase some predefined actions and more newly generated actions with the corresponding game screenshots. For the VARP agent, these generated actions originate from two sources: one is the human-guided trajectory system, and the other is the SOAG in the action planning system. The VLM is GPT-40.

To validate the effectiveness of the human-guided trajectory system, we introduced human guidance 452 and conducted a case study on task 13, a task of very hard difficulty. The objective was for the VARP 453 agent to control the player character to move from the "Earth Temple" spawn point to the location 454 of the Bullguard enemy within 5 minutes. GPT-40 was chosen as the VLM. The experimental re-455 sults showed a success rate of 40%. This indicates that human guidance can significantly enhance 456 the decision-making accuracy of the agent. Fig. 5 shows the new action responsible for pathfinding 457 generated by the human-guided trajectory system during the execution of task 13. Additionally, 458 Fig. 5 depicts the new action generated by SOAG in response to the enemy, Bullguard, during the 459 initial phase of this combat task. The enemy's current action indicates an impending attack: "swinging the axe downwards three times consecutively." Therefore, the new action should be to dodge 460 consecutively more than three times before counterattacking. 461



476 477 478

Figure 5: Simple cases of new actions generated by human guidance and SOAG.

As the combat task continues, the VARP agent can use or learn more complex actions to cope with more challenging scenarios. As shown in Fig.6, the actions in the first and second rows are predefined functions. The VARP agent automatically detects whether to use these actions based on the input visual information. For example, if the immobilization spell skill can be used, the agent executes the "fight_immobilization_spell_skill" action. Similarly, if the player's health is low, it uses the "recover_health" action.

The actions in the third row are generated by the human-guided trajectory system. Human prior knowledge can effectively guide the agent to improve efficiency in navigation tasks.

442 443 444

445

446 447

448

449

450

451

432

433

434

435

436 437

438

439 440



Figure 6: Complex cases of some actions and corresponding game screenshots.

The new actions in the fourth and fifth rows are summarized by SOAG after each combat interac-514 tion between the player character and the enemy and stored in the action library. These actions are 515 specific to particular enemies and their attack patterns. For instance, in the fourth row, when the 516 agent observes that an enemy named Bullguard is raising up his weapon, it indicates that the enemy 517 is about to perform the action "chopping the axe downwards three times consecutively." The agent 518 can then find a specific counter-action in the action library. At the beginning of the combat, this 519 "fight_new_action_bullguard_raise_weapon" action is defined as dodging four times consecutively, followed by attacking five times, as shown in Fig.5. As the combat progresses, this action is opti-521 mized to counterattack during the intervals between dodges, significantly increasing the success rate 522 and efficiency in defeating the enemy, as illustrated in the fourth row of Fig.6. This demonstrates 523 that SOAG can continuously optimize the actions it generated.

524 525

513

5 CONCLUSION

526 527 528

In this study, we have explored the boundaries of current Vision Language Models (VLMs) in the 529 context of complex action role-playing games (ARPGs) using "Black Myth: Wukong" as our exper-530 imental platform. Our proposed framework, VARP, introduces a novel approach to game interaction 531 by leveraging visual-only inputs for action planning in ARPG environments. The VARP framework 532 demonstrates its potential by achieving an 90% success rate in basic and moderate combat scenarios, suggesting that VLMs can be effectively utilized in tasks traditionally dominated by reinforcement 534 learning. Our proposed benchmark can effectively evaluate the performance of visual-only agents 535 in the BMW game. Additionally, the human operation dataset we provide offers a valuable resource 536 for future research, enabling the study of human-like gameplay and action decision-making in visually complex environments. Our findings underscore the promise of multimodal agents in enhancing generalization and performance in action-oriented tasks within video games. Moving forward, the 538 insights gained from this research could pave the way for more sophisticated agent designs that can handle a broader range of challenges in ARPGs and beyond.

540	REFERENCES
541	KEI EKENCES

549

555

559

561

562

566

567

568

569

577

- 542 analoganddigital. Dqn_play_sekiro. https://github.com/analoganddigital/DQN_ 543 play_sekiro, 2021. Accessed: 2024-09-17.
- 544 Anthropic. The claude 3 model family: Opus, sonnet, haiku, 2024.
- Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environ ment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279, 2013.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.
- Turing's Cat. Ai-wukong: Rl-based arpg gamebot. https://github.com/
 Turing-Project/RL-ARPG-Agent, 2024.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and
 Yu Su. Mind2web: Towards a generalist agent for the web, 2023. URL https://arxiv.
 org/abs/2306.06070.
 - Benjamin Ellis, Jonathan Cook, Skander Moalla, Mikayel Samvelyan, Mingfei Sun, Anuj Mahajan, Jakob Foerster, and Shimon Whiteson. Smacv2: An improved benchmark for cooperative multiagent reinforcement learning. Advances in Neural Information Processing Systems, 36, 2024.
- 563 fange, 2024. URL https://www.bilibili.com/video/BV1n74y1m7U3/.
- 564 565 Google. Gemini, 2024.
 - Izzeddin Gur, Hiroki Furuta, Austin Huang, Mustafa Safdari, Yutaka Matsuo, Douglas Eck, and Aleksandra Faust. A real-world webagent with planning, long context understanding, and program synthesis, 2024. URL https://arxiv.org/abs/2307.12856.
- Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan,
 and Dong Yu. Webvoyager: Building an end-to-end web agent with large multimodal models,
 2024. URL https://arxiv.org/abs/2401.13919.
- 573 Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia
 574 Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. Human575 level performance in 3d multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865, 2019.
- Karol Kurach, Anton Raichuk, Piotr Stańczyk, Michał Zając, Olivier Bachem, Lasse Espeholt, Carlos Riquelme, Damien Vincent, Marcin Michalski, Olivier Bousquet, et al. Google research football: A novel reinforcement learning environment. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 4501–4510, 2020.
- Junkai Li, Siyu Wang, Meng Zhang, Weitao Li, Yunghwei Lai, Xinhui Kang, Weizhi Ma, and Yang
 Liu. Agent hospital: A simulacrum of hospital with evolvable medical agents. *arXiv preprint arXiv:2405.02957*, 2024.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. Grounding dino: Marrying dino with grounded pre-training for open-set object detection, 2024. URL https://arxiv.org/abs/2303.05499.
 - OpenAI. Gpt-40. https://www.openai.com/, 2024.
- Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and
 Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings* of the 36th annual acm symposium on user interface software and technology, pp. 1–22, 2023.

- Siyuan Qi, Shuo Chen, Yexin Li, Xiangyu Kong, Junqi Wang, Bangcheng Yang, Pring Wong, Yifan Zhong, Xiaoyuan Zhang, Zhaowei Zhang, et al. Civrealm: A learning and reasoning odyssey in civilization for decision-making agents. *arXiv preprint arXiv:2401.10568*, 2024.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, et al. Chatdev: Communicative agents for software development. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15174–15186, 2024.
- Mikayel Samvelyan, Tabish Rashid, Christian Schroeder De Witt, Gregory Farquhar, Nantas
 Nardelli, Tim GJ Rudner, Chia-Man Hung, Philip HS Torr, Jakob Foerster, and Shimon Whiteson.
 The starcraft multi-agent challenge. *arXiv preprint arXiv:1902.04043*, 2019.
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning, 2023. URL https://arxiv.org/abs/2303.11366.
- 609 Claude 3.5 Sonnet. New and improved embedding model, 2022.

616

631

632

633

- Weihao Tan, Wentao Zhang, Xinrun Xu, Haochong Xia, Ziluo Ding, Boyu Li, Bohan Zhou, Junpeng Yue, Jiechuan Jiang, Yewen Li, Ruyi An, Molei Qin, Chuqiao Zong, Longtao Zheng, Yujie Wu, Xiaoqiang Chai, Yifei Bi, Tianbao Xie, Pengjie Gu, Xiyun Li, Ceyao Zhang, Long Tian, Chaojie Wang, Xinrun Wang, Börje F. Karlsson, Bo An, Shuicheng Yan, and Zongqing Lu. Cradle: Empowering foundation agents towards general computer control, 2024. URL https://arxiv.org/abs/2403.03186.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models, 2023a. URL https://arxiv.org/abs/2305.16291.
- Junyang Wang, Haiyang Xu, Jiabo Ye, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. Mobile-agent: Autonomous multi-modal mobile device agent with visual perception, 2024a. URL https://arxiv.org/abs/2401.16158.
- Zihao Wang, Shaofei Cai, Anji Liu, Yonggang Jin, Jinbing Hou, Bowei Zhang, Haowei Lin, Zhaofeng He, Zilong Zheng, Yaodong Yang, Xiaojian Ma, and Yitao Liang. Jarvis-1: Openworld multi-task agents with memory-augmented multimodal language models, 2023b. URL https://arxiv.org/abs/2311.05997.
- Zihao Wang, Shaofei Cai, Guanzhou Chen, Anji Liu, Xiaojian Ma, and Yitao Liang. Describe,
 explain, plan and select: Interactive planning with large language models enables open-world
 multi-task agents, 2024b. URL https://arxiv.org/abs/2302.01560.
 - Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. URL https://arxiv.org/abs/2201.11903.
- Peter R Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian, Thomas J Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, et al. Outracing champion gran turismo drivers with deep reinforcement learning. *Nature*, 602(7896):223– 228, 2022.
- Hui Yang, Sifu Yue, and Yunzhong He. Auto-gpt for online decision making: Benchmarks and
 additional opinions. *arXiv preprint arXiv:2306.02224*, 2023.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
 React: Synergizing reasoning and acting in language models, 2023. URL https://arxiv.org/abs/2210.03629.
- Yuexiang Zhai, Hao Bai, Zipeng Lin, Jiayi Pan, Shengbang Tong, Yifei Zhou, Alane Suhr, Saining Xie, Yann LeCun, Yi Ma, and Sergey Levine. Fine-tuning large vision-language models as decision-making agents via reinforcement learning, 2024. URL https://arxiv.org/abs/2405.10292.

- Chi Zhang, Penglin Cai, Yuhui Fu, Haoqi Yuan, and Zongqing Lu. Creative agents: Empowering agents with imagination for creative tasks, 2023a. URL https://arxiv.org/abs/2312.02519.
- Chi Zhang, Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang
 Yu. Appagent: Multimodal agents as smartphone users, 2023b. URL https://arxiv.org/ abs/2312.13771.
- Wangchunshu Zhou, Yixin Ou, Shengwei Ding, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen, Shuai Wang, Xiaohua Xu, Ningyu Zhang, Huajun Chen, and Yuchen Eleanor Jiang. Symbolic learning enables self-evolving agents. 2024. URL https://arxiv.org/abs/2406. 18532.
- 659 660 661

662

663 664

665

666

667 668

669

670

671

A APPENDIX

- A.1 OVERVIEW
 - Clarifications and Limitations (§A.2)
 - Additional Related Work (§A.3)
 - Additional Dataset Collection (§A.4)
 - Additional Performance Evaluation (§A.5)
 - Ethical Consideration (§A.6)
 - Demo Video (§A.7)
- 672 673 A.2 CLARIFICATIONS AND LIMITATIONS

674 Thank you for reading our paper. We would like to begin by clarifying and explaining some potential 675 concerns that may arise due to the extreme popularity of the BMW game. We want to emphasize that 676 our framework has broad applicability and will subsequently be generalized to include more games 677 and other scenarios, not just limited to the BMW game. In this paper, we explored the potential 678 of using VLMs to execute action combos in game tasks, particularly focusing on how it achieves 679 victory against medium-powered monsters by leveraging the advantages of both action planning and visual trajectory modules. Additionally, we provided a human operation dataset, which presents pos-680 sibilities for integrating technologies such as multi-modal retrieval-augmented generation, imitation 681 learning, and reinforcement learning. 682

683 We must also candidly acknowledge some limitations in our research, specifically: 1) Task Defini-684 tions: As LLM- and VLM-based agents are still evolving, the current task definitions are somewhat 685 simplistic. 2) Game Scenarios: Our research has only been tested within the BMW game and has not yet been extended to other scenarios. 3) Dataset Size: We have a limited amount of data, and 686 in the future work, we plan to recruit more volunteers to collect higher quality data to enhance the 687 depth of our research. 4) Model Capabilities: As shown in the performance evaluation section, there 688 is still room for improvement in existing VLMs, including speed and accuracy. Therefore, it would 689 be interesting to train an ARPG-specific VLM, such as VideoGameBunny¹. 690

Finally, we sincerely welcome your new ideas and feedback regarding this work, or even contribute
your game records. Please feel free to reach out to us, and we look forward to exploring together,
ultimately making VLMs play games as well as humans.

694 695

696

701

- A.3 ADDITIONAL RELATED WORK
- A.3.1 RL-BASED AGENTS IN ARPGS

Reinforcement learning (RL) has shown significant improvements in video games(Zhai et al., 2024;
Bellemare et al., 2013; Kurach et al., 2020; Berner et al., 2019; Ellis et al., 2024; Samvelyan et al., 2019; Jaderberg et al., 2019; Qi et al., 2024; Wurman et al., 2022) especially action role-playing

¹https://videogamebunny.github.io/

games (ARPG). DQN-play-sekiro (analoganddigital, 2021) employs the deep Q-network (DQN) al-gorithm to train AI to automate gameplay in "Sekiro: Shadows Die Twice." This project observes the game visuals and makes decisions based on the current state, gradually mastering the game strategy to defeat boss-level enemies. Additionally, AI achieves interactive learning based on reinforcement learning in "Black Myth: Wukong" by recognizing game images and scripting simulated keyboard input signals(Cat, 2024; fange, 2024). This method uses successful dodging as positive feedback while being attacked by monsters as negative feedback, prompting the AI to optimize its decision-making process.

These creative works not only showcase the potential of AI in complex gaming environments but also provide effective means for game testing and automated gameplay. However, agents trained solely using RL methods can only be applied to a limited range of specific tasks. For new tasks, the agent needs to be retrained. Therefore, agents based on this method have poor generalization capabilities.

A.4 ADDITIONAL DATASET COLLECTION

We collected a total of approximately 1,000 valid data samples, each representing a video segment
of a human completing a task along with the corresponding mouse and keyboard operation records.
Among these, 4.0% represent task 1, 12.5% represent task 2, and so on. The specific information is
shown in Fig. 7.



A.5 ADDITIONAL PERFORMANCE EVALUATION

Based on Sec. 3.4 of the main text, we also recorded the average time and average inference count for the VARP agent without human guidance to complete each task. Each inference count represents the generation of an executable action, with each combat action containing an average of 8.6 atomic operations. Additionally, we recorded the number of atomic operations performed by human for each task. By dividing this number by 8.6, we estimated the inference count of human in combat tasks. As shown in Tab. 4, compared to humans, the VARP agent has a much lower inference count in task 1, task 9, and task 10. This indicates that humans tend to perform a large number of redundant operations when completing more difficult or time-consuming tasks in ARPGs, which is not conducive to task completion. In contrast, the actions generated by the VARP agent in these tasks are relatively more refined and concise.

58	Task ID	GPT-40 (time)	Claude (time)	Gemini (time)	GPT-40 (count)	Claude (count)	Gemini (count)	Human(count)
59	1	16.09	19.12	17.14	71.6	88	77	98.7
60	2	0.53	0.65	0.63	3.8	5	4.4	2.3
00	3	0.18	0.23	0.18	1.4	1.6	1.4	1.7
51	4	0.57	0.68	0.64	4.6	5.2	4.6	3.0
20	5	0.69	0.77	0.68	5.5	6.25	5.4	3.5
2	6	0.27	0.25	0.11	2	1.8	1.2	1.3
3	7	0.81	0.78	0.69	5.4	5.7	5.75	3.0
./	8	0.41	0.42	0.38	3.8	3.2	2.8	2.6
)- <u>+</u>	9	1.24	1.19	1.19	8.3	9	8.5	16.7
5	10	2.20	-	2.06	13.5	-	13	36.6
66					•			

Table 4: Additional evaluation results of the average time (in minutes) and average inference count.

A.6 ETHICAL CONSIDERATION

Our method can automatically play ARPGs, which may lead to game cheating and false advertising. This can have a significant negative impact on society. Therefore, it is crucial to consider meth-ods that can reliably distinguish between genuine and forged content. We strongly condemn the unauthorized and malicious use of this technology and emphasize the need to consider ethical issues when using our method.

A.7 DEMO VIDEO

> We have provided a detailed demo video to demonstrate the effectiveness of our VARP agent. Please refer to https://varp-agent.github.io/.