

V-MAGE: A Game Evaluation Framework for Assessing Vision-Centric Capabilities in Multimodal Large Language Models

Xiangxi Zheng¹, Linjie Li², Zhengyuan Yang², Ping Yu¹,
Alex Jinpeng Wang^{3†}, Rui Yan^{4†}, Yuan Yao¹, Lijuan Wang²

¹Nanjing University, ²Microsoft Research,

³Central South University, ⁴Nanjing University of Science and Technology
zhengxx@nju.edu.cn, jinpengwang@csu.edu.cn, ruiyan@just.edu.cn

Abstract

Recent advancements in Multimodal Large Language Models (MLLMs) have demonstrated impressive capabilities in visual-text processing. However, existing static image-text benchmarks are insufficient for evaluating their dynamic perception and interactive reasoning abilities. We introduce **Vision-centric Multiple Abilities Game Evaluation (V-MAGE)**, a novel game-based evaluation framework designed to systematically assess MLLMs' visual reasoning in interactive, continuous-space environments. V-MAGE features five distinct video games comprising over 30 carefully constructed evaluation scenarios. These scenarios are set in free-form, visually complex environments that require models to interpret dynamic game states and make decisions based solely on visual input, thereby closely reflecting the conditions encountered by human players. To ensure robust and interpretable comparisons across models, V-MAGE employs a dynamic ELO-based ranking system that accounts for varying difficulty levels and task diversity. Benchmarking state-of-the-art MLLMs against human baselines reveals that while leading models approach human-level performance in simple tasks, their performance drops significantly in complex scenarios requiring advanced reasoning and task orchestration. This persistent performance gap highlights fundamental limitations in current MLLMs' ability to perform vision-grounded, interactive frame-by-frame control in simulated continuous-time environments. Through extensive analyses, we demonstrate the utility of V-MAGE in uncovering these limitations and providing actionable insights for improving the visual and reasoning capabilities of MLLMs in dynamic, interactive settings. Code is publicly available at <https://github.com/CSU-JPG/V-MAGE>.

1 Introduction

Building on the success of Large Language Models (LLMs) in text-based tasks (Bai et al., 2023; Cai et al., 2024; OpenAI, 2023a), researchers have extended their capabilities to visual-text multimodal tasks through Multimodal Large Language Models (MLLMs) (OpenAI, 2023b; Liu et al., 2023a; Team, 2023; Yang et al., 2023; Li et al., 2024; Wang et al., 2024; Bai et al., 2025a; Zhu et al., 2025). Various multimodal evaluation benchmarks, such as MME (Fu et al., 2023), MMBench (Liu et al., 2023b), SEED-Bench (Li et al., 2023) have driven improvements in MLLM performance. With improving model capabilities, researchers are shifting toward open-world, dynamic, multi-round tasks beyond static benchmarks with fixed image-text inputs, as these better reflect real-world interaction and reasoning challenges. Among the promising approaches for evaluating models in such dynamic settings, game-based evaluation has emerged as a promising alternative, offering a more natural and interactive assessment of a model's perception and reasoning abilities.

While progress has been made in game-based MLLM benchmarks, current approaches predominantly rely on text-based (Costarelli et al., 2024; Hu et al., 2024; Duan et al., 2024) or grid-based (Zhang et al., 2024; Wang et al., 2025a; Paglieri et al., 2024) games. In such settings, limited visual reasoning demands and static, fully textually renderable content restrict evaluation of spatial, temporal, and dynamic complexities crucial for real-world problem-solving. In contrast, the rich visual information inherent in video games presents a valuable opportunity to assess MLLMs' genuine visual reasoning capabilities, potentially addressing current methodological limits.

To address the lack of vision-centric video game benchmarks, we present **Vision-centric Multiple Abilities Game Evaluation (V-MAGE)**, which al-

[†]Corresponding authors: Alex Jinpeng Wang, Rui Yan

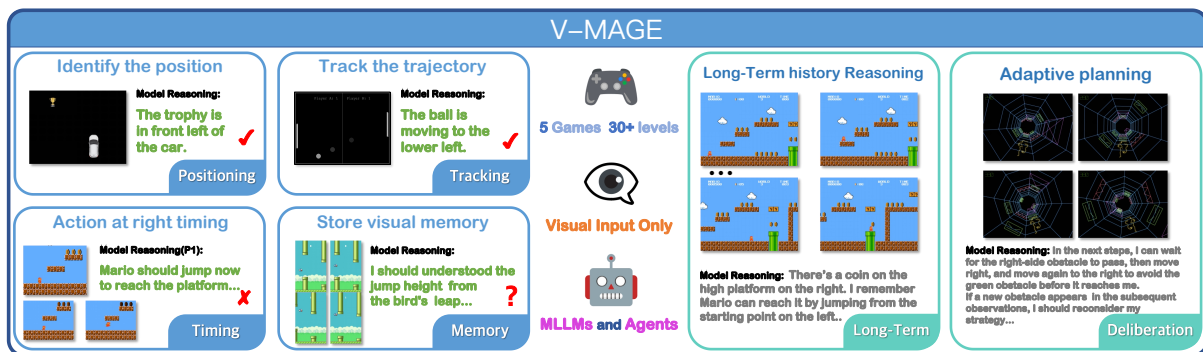


Figure 1: The overview of the V-MAGE benchmark, designed to evaluate vision-centric capabilities and higher-level reasoning of MLLMs across 5 free-form games with 30+ levels. V-MAGE assesses critical abilities in visual reasoning, providing a comprehensive evaluation of model performance in complex, dynamic environments.

lows a thorough assessment of diverse model and agent abilities within dynamic, interactive game environments and addresses key limitations in current game-based evaluations of MLLM capabilities.

With V-MAGE, we evaluate leading MLLMs across five interactive games across 30+ levels. Results highlight significant challenges posed by the dynamic visual interaction environment for MLLMs. The results reveal that current MLLMs, despite excelling in static benchmarks, lack perception, multi-step reasoning, and task orchestration for human-level play in dynamic games.

Our contributions are summarized as follows:

- We established V-MAGE, an interactive and visually rich evaluation framework focused on dynamic interaction and vision-centric reasoning. It also serves as a sandbox environment conducive to vision agent development.
- We evaluated various publicly available MLLMs with V-MAGE, measuring model performance with ELO scores and highlighting the significant gap between model performance and human-level proficiency on complex tasks.
- Through the evaluation results of V-MAGE, we further analyzed the reasons for the suboptimal performance of current MLLMs on video game tasks, including deficiencies in several fundamental visual capabilities, challenges in reasoning during prolonged interactions, and issues such as anchoring bias, among others.

2 Related work

MLLMs and Multimodal Agents. As LLMs (Qwen et al., 2025; Cai et al., 2024; OpenAI, 2023a) advance, MLLMs have emerged to handle multimodal tasks by integrating text and visual inputs (Zhu et al., 2025; Bai et al., 2025b;

Chen et al., 2024c; Liu et al., 2023a). Open-source models like InternVL and QwenVL are narrowing the gap (Chen et al., 2024b) with closed-source models such as GPT-4o (OpenAI, 2024), and Gemini (Team, 2023).

MLLMs are evolving into interactive multimodal agents, finding applications in areas such as robotics (Driess et al., 2023), virtual assistants (Brohan et al., 2023, 2022), GUI automation (Xu et al., 2024; Bonatti et al., 2024; Zhang et al., 2023), and game agents (Tan et al., 2024; Chen et al., 2024a). These domains require sequential reasoning, memory, and adaptability, which static benchmarks inadequately capture.

MLLM Benchmarks. Classic MLLM benchmarks have focused on tasks like Visual Question Answering (VQA) (Antol et al., 2015; Goyal et al., 2017; Li et al., 2018; Marino et al., 2019) and image captioning (Chen et al., 2015; Agrawal et al., 2019; Sidorov et al., 2020). More comprehensive benchmarks, such as MME (Fu et al., 2023), MM-Bench (Liu et al., 2023b), SEED-Bench (Li et al., 2023), MMMU (Yue et al., 2024), and MM-Vet (Yu et al., 2023, 2024), introduce broader assessments across multiple domains.

Most of these evaluations rely on structured multiple-choice and VQA-style tasks, limiting their ability to measure real-world problem-solving and interactive reasoning. Recent multimodal agent benchmarks like OSWorld (Xie et al., 2024), Windows Agent Arena (Bonatti et al., 2024), and COMMA (Ossowski et al., 2024) assess broader capabilities: open-ended tasks in real environments, OS interaction, and multi-agent collaboration.

Evaluating MLLMs in Games. Recent work (Tan et al., 2024; Chen et al., 2024a; Ruoss et al., 2024) has explored MLLMs in interactive gaming environments. Existing game-based

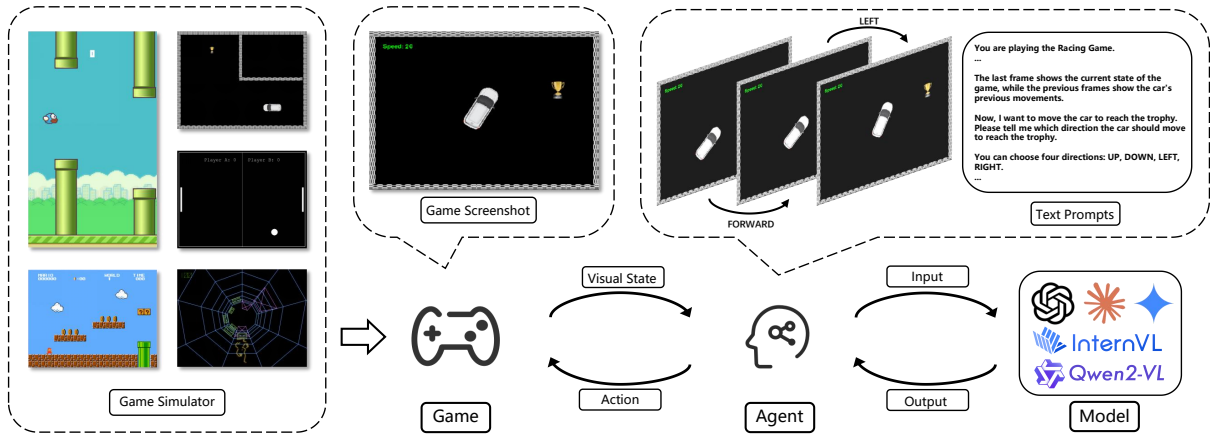


Figure 2: V-MAGE games and evaluation pipeline. V-MAGE employs five distinct games, each with several levels, to facilitate a decomposed evaluation of model performance. These games include **FlappyBird**, **Race**, **SuperMario**, **Pong** and **TempestRun**. During the evaluation process, the **Agent** module receives visual game state information directly from the **Game** module, primarily in the form of screenshots. The **Agent** module then structures these screenshots, combined with prompts containing the game rules, into the appropriate input format for MLLMs. Subsequently, the model’s output is processed by the **Agent** module to generate executable actions, which are then transmitted back to the **Game** module to update the environment state.

benchmarks span text-only settings (Costarelli et al., 2024; Hu et al., 2024; Duan et al., 2024), grid- or board-style visual settings (Zhang et al., 2024; Wang et al., 2025a; Paglieri et al., 2024), escape-room style multimodal reasoning environments (Wang et al., 2025b), and GUI/adventure-game agents with long story arcs (Ahn et al., 2025). These benchmarks are valuable and complementary, but they emphasize different capability profiles. In particular, many settings either remain text-dominant, expose substantial symbolic structure, or involve UI elements that can be mapped to discrete states. V-MAGE instead targets continuous-space, arcade-style environments in which key information lies in pixel-level dynamics such as motion, collision, relative position, and timing. This makes V-MAGE a complementary benchmark for evaluating vision-centric control and temporally grounded reasoning rather than a replacement for prior game benchmarks.

3 V-MAGE Benchmark

We present V-MAGE, a benchmark built on video game environments designed to evaluate the comprehensive performance of MLLMs, with a focus on vision-centric capabilities. Its defining features are as follows:

- **Vision Centric Gameplay.** Models receive only visual input, requiring pixel-level scene understanding, object tracking, and spatial-temporal reasoning. V-MAGE features continuous-space environments, allowing models to explore the almost infinite state space. Each game is designed

with different difficulty levels that target various skill dimensions.

- **Extensible Evaluation Framework.** V-MAGE extends beyond model evaluation to assess agentic skills that are out-of-scope for current MLLMs. Our game-agent-model three-module evaluation pipeline allows optimizations in both MLLMs and their agent strategies.
- **Adaptive ELO-based Ranking.** V-MAGE uses a dynamic ELO system to provide a unified and interpretable metric across diverse games and difficulty levels. Unlike raw scores varying in scale across tasks, ELO captures relative skill via win-loss dynamics between model performances on shared levels.

3.1 Evaluation Pipeline

V-MAGE separates the game environment from the MLLM, ensuring that all information is conveyed solely through visual input. Rather than evaluating raw inference latency, V-MAGE evaluates interactive frame-by-frame control in a simulated continuous-time environment. Concretely, when an action is requested, the game is paused while the model processes the visual inputs and generates its response. This frame-pausing mechanism decouples temporal reasoning from infrastructure-related delays and enables fair comparison across models with different serving conditions.

As depicted in Figure 2, the system operates through iterative action cycles composed of three sequential components. The Game Module serves as the environment interface, executing game logic,

capturing screenshots of the current state, and transmitting these visual frames to subsequent modules. The Agent Module integrates three critical data streams: (1) current visual inputs, (2) short-term temporal context from past observations, and (3) task-specific textual prompts such as game rules. This synthesized input is structured into a multimodal format compatible with the MLLM. The Model Execution Phase completes the cycle, wherein the MLLM generates an action command that undergoes semantic validation by the Agent Module before being relayed back to the Game Module for state updates.

To prioritize unbiased evaluation of core MLLM capabilities, V-MAGE’s architecture adopts a deliberately minimalistic design, avoiding auxiliary subsystems that might obscure model performance. The framework retains modular extensibility, allowing strategy modifications without altering core protocols, thereby ensuring benchmarking rigor while accommodating specialized needs.

V-MAGE incorporates five human-playable video games (Figure 2), each featuring 3 to 10 levels, culminating in over 30 distinct evaluation environments. In contrast to traditional grid-based evaluation setups, V-MAGE selects games based on specific principles. The games feature free-form or continuous-space visual environments, facilitating more nuanced and flexible model movement and interaction. Crucially, to effectively assess vision-centric capabilities, the game environments are designed to be **visually irreducible**. This characteristic ensures that the system state cannot be fully discretized or textually summarized without significant information loss, thereby necessitating continuous visual grounding throughout the reasoning process. Appendix B details game selection criteria and sources.

3.2 Games and Levels

Existing game-based benchmarks indicate that MLLMs frequently struggle to achieve meaningful scores at standard human-level difficulties in conventional game-based benchmarks (Zhang et al., 2024; Wang et al., 2025a).

This limits their discriminative power for fine-grained capability assessment and inter-model comparisons in complex tasks. To address this, V-MAGE introduces a multi-level assessment framework that evaluates models across various skill dimensions and provides granular performance diagnostics through difficulty-stratified tasks. Specif-

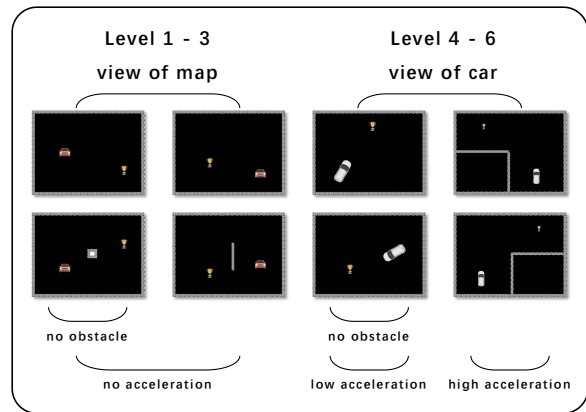


Figure 3: Race level design. Six levels progressively increase in difficulty while sharing the core objective: navigating a car to a trophy. Detailed Race level configurations are provided in Appendix Table 17.

ically, levels are designed for each game with gradually increasing complexity, varying control paradigms and perceptual challenges. For instance, Figure 3 illustrates the level design in Race. Detailed information on the level design for all games can be found in Appendix B.2.

3.3 Evaluation Metrics

V-MAGE employs a robust implementation of the ELO rating system to evaluate MLLMs, building on recent advancements in LLM benchmarking (Duan et al., 2024; Jiang et al., 2024). This competitive evaluation framework ensures reliable model rankings by leveraging competitive evaluation mechanisms, ensuring robustness through dynamically balanced interactions.

Games demonstrate diverse scales and difficulty thresholds. The ELO system mitigates this variability by providing a standardized metric for comparing model performance across heterogeneous environments. Moreover, ELO is inherently sensitive to performance consistency. A model that achieves a high average score through a mix of exceptional successes and frequent failures may still be ranked lower, as its instability would likely lead to more losses in direct, pairwise matchups. This allows the ELO system to reward reliable performance over erratic, high-variance gameplay, which simple score averaging might otherwise obscure.

Furthermore, it measures incremental advancement in games with non-linear scoring plateaus, where minor improvements vary by performance range. For instance, it differentiates progress from 80 to 85 and breakthroughs from 95 to 100.

We introduce an ELO-based ranking system to

Table 1: The comparison of V-MAGE with existing game-based evaluation benchmarks. *Text in V-MAGE only represents the instructions for game rules and output format.

Game Benchmarks	Game Type	Input	Reasoning Type	Level Design
GameBench(Costarelli et al., 2024)	Word	Text-Only	Text Reasoning	✗
GameArena(Hu et al., 2024)	Word	Text-Only	Text Reasoning	✗
GTBench(Duan et al., 2024)	Word	Text-Only	Text Reasoning	✗
ING-VP(Zhang et al., 2024)	Grid Based	Single-Image-Text	Visual Aid	✓
LVL-Playground(Wang et al., 2025a)	Grid Based	Single-Image-Text	Visual Aid	✓
BALROG(Pagliari et al., 2024)	Word / Grid Based	Single-Image-Text	Text / Visual Aid	✗
Orak(Park et al., 2025)	Video	Single-Image-Text	Text / Visual Aid	✗
V-MAGE	Video	Multi-Images-Text*	Vision-Centric Reasoning	✓

assess model performance by means of competitive pairwise comparisons. In each game level, models are randomly matched in pairs for up to 100 evaluation rounds. The outcomes are determined based on their gameplay scores and the rates of valid actions taken. Detailed mathematical formulations of the pairing mechanism, rating updates, and stabilization process are provided in Appendix C.

3.4 Comparison to Existing Works

Humans play dynamic games using visual perception and intuitive reasoning, a process largely uncaptured by most existing MLLM game benchmarks. Many of these rely on grid-based games (Wang et al., 2025a; Zhang et al., 2024) where states are textually representable. While such benchmarks assess text-based reasoning, similar to traditional LLM tasks (Paglieri et al., 2024), they offer limited insights into MLLMs’ visual intuitive reasoning. Models often bypass genuine visual perception here, acting as OCR converters, which hinders assessing and improving vision-centric abilities. (Park et al., 2025) also employed video games as a testbed, wherein visual information remained auxiliary, and the game state was accessed chiefly through text-based inputs.

In contrast to this paradigm, V-MAGE shifts the evaluation focus by embedding models in dynamic visual environments that fundamentally require temporally grounded perception and action based on visual input. V-MAGE deliberately adopts environments lacking rigid grid structures, where the states of characters and objects cannot be easily simplified into sparse, coordinate-based textual descriptions. This design compels models to continuously leverage the visual modality during reasoning, rather than discarding it after an initial conversion.

Furthermore, unlike benchmarks where decisions are made based on a single static frame, such as in many chess-like environments, V-MAGE requires sophisticated temporal reasoning across sequences of frames to make informed decisions, more closely mirroring human gameplay dynamics.

By shifting evaluation to more naturalistic and visually complex dynamic game environments, V-MAGE provides a more rigorous and representative test of MLLM capabilities, particularly in assessing their visual intuitive reasoning. A holistic comparison between V-MAGE and existing game benchmarks is presented in Table 1.

4 Experiments

As the baseline setting for the V-MAGE benchmark, we evaluate state-of-the-art MLLMs using full-precision models under a minimal naive agent strategy (Appendix A.1.2) to ensure a fair comparison. The naive agent uses the most recent k frames (typically $k = 3$) for reasoning, together with reasoning history, prior actions, and game rules. Detailed experimental settings and prompts can be found in Appendix A and Appendix B.4.

4.1 Main Result

Scores and Rankings. The evaluation results clearly demonstrate a performance gradient across models ranging from 7B to 70B+ parameters. This also highlights that the dynamic visual reasoning tasks we propose represent a universal challenge for current MLLMs. Rankings from ELO scores and the Average Ratio may occasionally differ. This discrepancy arises because the ELO system rewards performance consistency (penalizing unstable, high-variance results) and provides a more balanced, holistic assessment across games with

Table 2: Performance comparison across different games based on the ELO ranking system. The Random baseline refers to randomly selecting actions from the predefined action space during decision-making phases. Average performance ratio, abbreviated as **Avg. Ratio**, refers to the average percentage of the model’s score compared to the human baseline score.

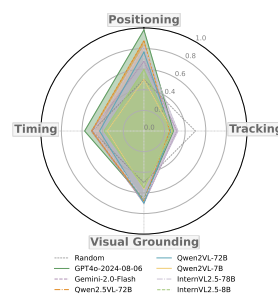
Model	Flappybird	Pong	Race	Supermario	Tempestrun	Avg. ELO Score	Avg. Ratio (%)
GPT-5-2025-08-07	1572	1939	1710	1584	1743	1710	43.4
Gemini-2.5-Pro	1526	1602	1660	1758	1474	1604	36.3
Claude-3.7-Sonnet	1560	1570	1633	1582	1369	1543	30.8
Gemini-2.5-Flash	1578	1524	1520	1531	1489	1528	23.8
GPT-4o	1557	1449	1581	1518	1527	1526	26.6
Gemini-2.0-Flash-Thinking	1517	1479	1503	1564	1516	1516	22.6
GPT-5.1-2025-11-13	1552	1514	1507	1449	1411	1486	20.1
Gemini-2.0-Flash	1494	1461	1437	1499	1530	1484	16.7
Qwen3-VL-235B-A22B-Instruct	1567	1441	1517	1556	1496	1515	24.3
Qwen2.5-VL-72B-Instruct	1556	1442	1506	1541	1530	1515	22.8
InternVL2.5-78B	1463	1462	1465	1543	1528	1492	19.2
Qwen2-VL-72B-Instruct	1426	1445	1442	1505	1547	1473	16.5
InternVL2.5-8B	1459	1448	1431	1373	1495	1441	12.9
Qwen2.5-VL-7B-Instruct	1457	1446	1423	1354	1517	1439	12.1
Qwen2-VL-7B-Instruct	1470	1447	1408	1362	1501	1438	11.4
Keye-VL-8B-Preview	1419	1444	1428	1381	1499	1434	13.1
Phi-4-multimodal-instruct	1404	1454	1420	1482	1385	1429	13.7
Random	1422	1434	1410	1417	1445	1426	11.0

varying score scales. In contrast, the Avg Ratio metric can be skewed when averaging across tasks with imbalanced performance levels. More detailed analyses are provided in Appendix A.2 and A.3.

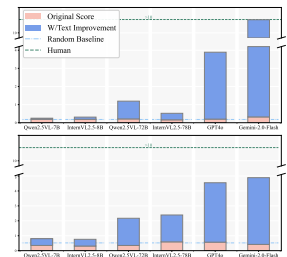
Significant Performance Gap Between MLLMs and Humans in Complex Scenes.

We recruited five novice human participants to play the games under the same evaluation interface as the MLLMs and used their average scores as the human baseline. Participants were not specifically trained on the V-MAGE games; they received the same high-level natural-language rules as the models, observed the same screenshot-based inputs, and played under the same frame-pausing protocol. Appendix Figure 7 compares leading MLLMs and human players across different levels.

These results should be interpreted as evidence about vision-driven decision-making under controlled dynamic tasks, rather than as a blanket equivalence between arbitrary game scores and general visual reasoning. Under the paused protocol and simplified low-difficulty settings, human novices remain near ceiling, while model performance drops sharply as tasks require more temporal understanding and strategic adaptation. Together with the text-state ablation in Figure 4b, this suggests that the gap reflects a combination of visual perception, temporal reasoning, and action-planning limitations rather than mere reaction-speed constraints.



(a) Capability maps of the underlying visual capabilities of each model.



(b) Model performance with vs. without text information on Pong Level 2 (top) and Flappy Bird Level 3 (bottom).

4.2 Further Analysis

Unit Tests for Core Visual Abilities. We devised a unit test for vision-centric abilities by extracting foundational levels from V-MAGE. Figure 4a presents the capability profiles of various models across four core visual competencies. Scores near or below baseline suggest little effective relevant reasoning, while higher scores indicate a greater likelihood of correct reasoning. For each capability, effective reasoning was evaluated by calculating the percentage of model scores that exceeded a random baseline score on the corresponding unit test levels (as defined in Appendix E.1).

As depicted, most models substantially outperform the random baseline in **Positioning** and **Visual Grounding**, indicating a degree of proficiency in single-frame image comprehension and basic

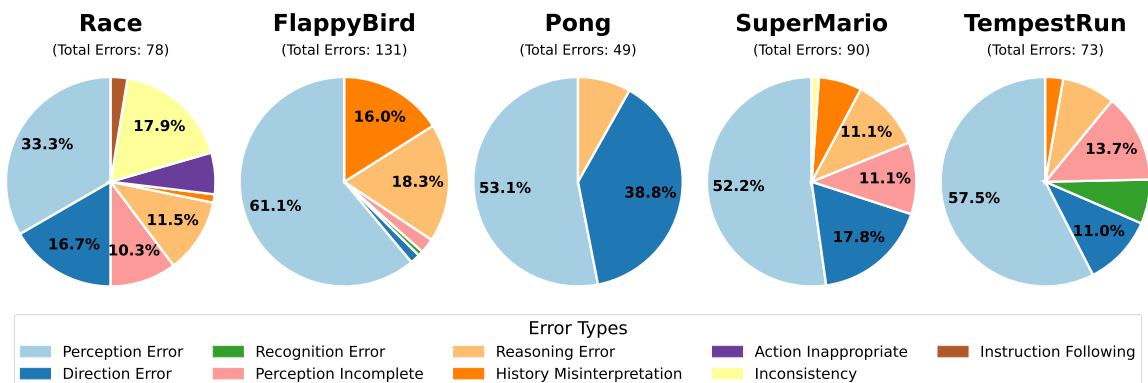


Figure 5: Error type probability distribution for GPT4o across 494 samples.

visual information perception. However, performance notably declines in **Tracking** and **Timing**, which require processing continuous frame information and executing precise spatiotemporal judgments. For the **Tracking** task, nearly all models fail to significantly surpass the random baseline.

Limitations Beyond Visual Perception. While visual perception constitutes a critical assessment dimension in V-MAGE’s game tasks, our experiments revealed additional limitations and deficiencies in other aspects. To validate this, we conducted supplementary experiments in simple levels providing textual game state descriptions, thereby bypassing the perception process.

As shown in Figure 4b, providing textual descriptions of the game state significantly improved the performance of most evaluated models, with this gain being particularly prominent in larger models such as Gemini and in games requiring precise state understanding like Pong. This notable performance increase when perception is bypassed strongly suggests that limitations in processing visual information are indeed a significant bottleneck for current MLLMs.

However, despite this substantial performance gain, the models’ scores still remained considerably

lower than the human baseline in most cases. This persistent gap indicates that while visual perception challenges are critical, the models’ limitations extend beyond merely “seeing” the state accurately. It highlights that significant bottlenecks also exist in the downstream processes responsible for robust interpretation of information (even when provided textually or perceived imperfectly), strategic planning, and effective action generation in complex and dynamic environments. Furthermore, the less pronounced performance improvement observed in smaller models (like Qwen2.5VL 7B) suggests that inherent limitations in their core reasoning capabilities may also act as a performance bottleneck. Check Appendix E.2 for more experimental details.

Anchoring Bias in Model Inference Processes.

When provided with historical information to aid reasoning in dynamic game progression, MLLMs often exhibit anchoring bias, particularly when processing similar consecutive frames. This bias manifests as an undue influence of prior inferences on current reasoning, hindering the accurate identification of subtle visual changes and unique frame details. Models tend to favor relying on historical textual descriptions over nuanced visual input, making them less sensitive to fine-grained visual

Table 3: Average number of rounds for each model to generate different responses.

Model	Race	FlappyBird	Pong	TempestRun	Avg.
Qwen2VL 7B	4.3	25.9	13.7	7.3	12.8
Qwen2.5VL 72B	2.3	19.3	2.6	5.3	7.4
InternVL2.5 8B	2.0	6.9	6.7	8.0	5.9
InternVL2.5 78B	6.8	16.0	2.0	3.0	7.0
GPT4o	1.0	1.6	1.0	1.0	1.1
PCC r (Avg. Rounds vs. ELO)	-0.57	-0.71	-0.87	-0.72	-0.72

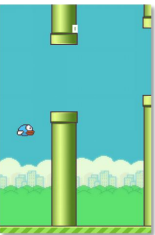

<p>FlappyBird</p>  <p>Level 2</p>	<p>GPT4o Description: The bird is slightly below the center of the pipe gap. The gap between the upper and lower pipe appears vertically aligned to allow safe passage if the bird is centered. Currently, the bird is a bit too low relative to the middle of the gap. Reasoning: Since the bird is positioned slightly below the center of the gap, and the goal is to move horizontally through the gap without touching either the top or bottom pipe, the bird needs to ascend slightly to align with the center of the gap. There is no gravity, so unless it is moved manually, the bird will stay at the same height. Action: UP ✓</p> <p>InternVL2.5-78B Description: The bird is currently positioned below the upper pipe and above the lower pipe, which means it is aligned well enough with the gap between the pipes. There is sufficient space for the bird to pass through without needing an adjustment. Reasoning: Since the bird is already correctly positioned within the gap between the pipes, taking an action like UP or DOWN would unnecessarily risk moving out of alignment with the gap. Therefore, the best course of action is to maintain its current position until it successfully passes through the gap. Action: KEEP</p> <p style="text-align: right; border: 1px solid black; padding: 2px;">Perception Error ✗</p>
<p>Race</p>  <p>Level 3</p>	<p>Gemini-2.0-Flash Description: The trophy is located to the right and slightly below the car. Reasoning: There is a checkered finish line obstacle situated vertically between the car and the trophy. Moving directly right would result in a collision with this obstacle and a game over. To reach the trophy, the car needs to move downwards first to avoid the finish line, and then move right. ✓</p> <p>Qwen2.5VL 72B Description: The trophy is positioned to the right and slightly below the car. There is a vertical line of blocks between the car and the trophy, which acts as an obstacle. Reasoning: To reach the trophy, the car needs to move to the right first to bypass the obstacle. Moving directly down or up would not help in reaching the trophy due to its current position relative to the car. Once the car moves to the right past the obstacle, it can then adjust its path if necessary to align with the trophy's position. Action: RIGHT ✗</p> <p style="text-align: right; border: 1px solid black; padding: 2px;">Reasoning Error ✗</p>

Figure 6: Case examples illustrating Perception Error and Reasoning Error in FlappyBird and Race. The FlappyBird example shows a *Perception Error* where the model misjudges the bird’s vertical position relative to the pipe gap. The Race example illustrates a *Reasoning Error* where the model fails to plan a path around an obstacle between the car and the trophy, resulting in a suboptimal action.

updates, consequently leading to unchanged reasoning content over extended game sequences.

As shown in Table 3, models vary significantly in their responsiveness; for instance, in FlappyBird, Qwen2.5VL 72B altered its reasoning only once every 19.3 responses on average, significantly less frequently than GPT-4o (1.6 responses). The Pearson correlation coefficients (**PCC** r) reveal a consistent negative correlation between the average rounds to change response and ELO score across games, with an average r of **-0.72**. This highlights a critical challenge in maintaining responsiveness to dynamic visual input and its direct impact on task success. To investigate the impact of pipeline settings (e.g., frame sampling and decision frequency) on anchoring bias, we conduct experiments, with results presented in Appendix E.4.

Analysis of Model Errors in V-MAGE. For GPT-4o’s complete inputs and responses across all game levels after one to two rounds of gameplay, we uniformly sampled 494 interaction sets for manual annotation and categorized the primary error types. The visualization results depicting the distribution of these errors are presented in Figure 5. The annotation method for error distribution and statistical details are provided in Appendix E.5.

The predominant error type was *perception error*, with *direction error* being particularly preva-

lent. In such cases, the model frequently misidentified directional cues in visual content, leading to incorrect inferences. Another major category was *reasoning error*, primarily involving logical flaws and decision-making failures, including misinterpretation of historical information (*history misinterpretation*) and suboptimal action selection (*action inappropriate*). Figure 6 presents examples illustrating perception and reasoning errors.

Additionally, *perception incomplete* errors were commonly observed, where the model failed to fully extract useful information from visual inputs, resulting in partial information acquisition. *Inconsistency* errors mainly occurred in scenarios permitting multiple valid solutions, where the model exhibited unstable decision-making by frequently revising its choices, ultimately leading to timeout failures due to excessive deliberation. Notably, *instruction following* errors were virtually absent, as GPT-4o consistently adhered to the provided prompts. Additional case study analyses are documented in Appendix G.

V-MAGE Poses Significant Challenges to MLLMs. Unlike conventional static VQA or text-reducible grid-based benchmarks, V-MAGE evaluates interactive frame-by-frame control in dynamic, vision-centric game environments under a paused continuous-time protocol. The framework exposes

persistent limitations in current MLLMs. Models struggle to integrate information across sequences of frames, which affects tracking, temporal reasoning, and trajectory understanding. This difficulty may contribute to anchoring bias, as models can over-rely on prior inferences and under-react to subtle visual changes. Furthermore, MLLMs exhibit fundamental deficiencies in planning, strategic decision-making, and action generation. These limitations persist even when initial visual processing challenges are mitigated, highlighting that deficiencies in the core reasoning process itself extend beyond perception.

5 Conclusion

This paper introduces V-MAGE, a pioneering game-based evaluation framework designed to assess the vision-centric capabilities of MLLMs in dynamic, interactive environments. Using over 30 levels across 5 games, we reveal significant limitations: models exhibit insufficient multi-image perception, leading to issues like anchoring bias, and demonstrate fundamental deficiencies in complex reasoning and strategic planning that persist even when perceptual challenges are mitigated. Highlighting the need for enhanced multi-frame processing and strategic planning, V-MAGE establishes a rigorous standard to drive the development of robust, human-like visual intelligence.

Limitations

While V-MAGE represents a significant step towards evaluating MLLMs in interactive, dynamic, and visually complex environments that closer resemble real-world tasks, the current benchmark is still constrained by the inherent scale and complexity of the included games. This is a current boundary imposed by balancing complexity with controllability for systematic evaluation. As MLLM capabilities continue to advance and hardware performance improves, we anticipate being able to incorporate larger and more complex game environments in future iterations. These future environments will be designed to offer a wider array of challenges, further pushing the limits of MLLM evaluation and narrowing the gap between simulated and real-world performance assessment.

Ethical Considerations

This research contributes to the field of multi-modal models by providing a novel and challenging

benchmark for evaluating vision-centric capabilities in dynamic environments. The primary positive impact is facilitating the diagnosis of limitations in current MLLMs and guiding future research towards developing more capable, robust, and potentially safer AI systems for real-world interaction. As our work focuses on foundational evaluation in simulated environments and does not involve the deployment of high-risk models or the collection of sensitive personal data, the potential for negative societal impacts is considered minimal and indirect at this stage. We believe that developing better evaluation tools is a crucial step towards building more reliable and trustworthy AI.

Acknowledgment

This work is supported by the Frontier Technologies R&D Program of Jiangsu (BF2024059) and the National Natural Science Foundation of China (Grant #62572229 and #62502544).

References

- Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. 2019. No-caps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957.
- Jaewoo Ahn, Junseo Kim, Heeseung Yun, Jaehyeon Son, Dongmin Park, Jaewoong Cho, and Gunhee Kim. 2025. Flashadventure: A benchmark for gui agents solving full story arcs in diverse adventure games. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 23365–23395.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *ICCV*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025a. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu,

- Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. [Qwen2.5-vl technical report](#). *Preprint*, arXiv:2502.13923.
- Rogério Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Bucker, Lawrence Jang, and Zack Hui. 2024. [Windows agent arena: Evaluating multi-modal os agents at scale](#). *Preprint*, arXiv:2409.08264.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alex Herzog, Jasmine Hsu, Brian Ichter, and 35 others. 2023. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *arXiv preprint arXiv:2307.15818*.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, and 32 others. 2022. Rt-1: Robotics transformer for real-world control at scale. In *arXiv preprint arXiv:2212.06817*.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, and 81 others. 2024. [Internlm2 technical report](#). *Preprint*, arXiv:2403.17297.
- Peng Chen, Pi Bu, Jun Song, Yuan Gao, and Bo Zheng. 2024a. [Can vlms play action role-playing games? take black myth wukong as a study case](#). *Preprint*, arXiv:2409.12889.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, and 1 others. 2024b. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences*, 67(12):220101.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, and 1 others. 2024c. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Anthony Costarelli, Mat Allen, Roman Hauksson, Grace Sodunke, Suhas Hariharan, Carlson Cheng, Wenjie Li, Joshua Clymer, and Arjun Yadav. 2024. [Gamebench: Evaluating strategic reasoning abilities of llm agents](#). *Preprint*, arXiv:2406.06613.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, and 3 others. 2023. Palm-e: An embodied multimodal language model. In *arXiv preprint arXiv:2303.03378*.
- Jinhao Duan, Renming Zhang, James Duffenderfer, Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin, Mohit Bansal, Tianlong Chen, and Kaidi Xu. 2024. [Gtbench: Uncovering the strategic reasoning limitations of llms via game-theoretic evaluations](#). *Preprint*, arXiv:2402.12348.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, and 1 others. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Lanxiang Hu, Qiyu Li, Anze Xie, Nan Jiang, Ion Stoica, Haojian Jin, and Hao Zhang. 2024. [Gamearena: Evaluating llm reasoning through live computer games](#). *Preprint*, arXiv:2412.06394.
- Dongfu Jiang, Max Ku, Tianle Li, Yuansheng Ni, Shizhuo Sun, Rongqi Fan, and Wenhui Chen. 2024. [Genai arena: An open evaluation platform for generative models](#). *Preprint*, arXiv:2406.04485.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023. [Seed-bench: Benchmarking multimodal llms with generative comprehension](#). *Preprint*, arXiv:2307.16125.
- Chunyuan Li, Zhe Gan, Zhengyuan Yang, Jianwei Yang, Linjie Li, Lijuan Wang, Jianfeng Gao, and 1 others. 2024. Multimodal foundation models: From specialists to general-purpose assistants. *Foundations and Trends® in Computer Graphics and Vision*, 16(1-2):1–214.
- Qing Li, Qingyi Tao, Shafiq Joty, Jianfei Cai, and Jiebo Luo. 2018. Vqa-e: Explaining, elaborating, and enhancing your answers for visual questions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 552–567.

- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, and 1 others. 2023b. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- OpenAI. 2023a. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- OpenAI. 2023b. [Gpt-4v\(ision\) system card](#).
- OpenAI. 2024. [Hello gpt-4o](#). *OpenAI Blog*.
- Timothy Ossowski, Jixuan Chen, Danyal Maqbool, Zefan Cai, Tyler Bradshaw, and Junjie Hu. 2024. Comma: A communicative multimodal multi-agent benchmark. *arXiv preprint arXiv:2410.07553*.
- Davide Paglieri, Bartłomiej Cupiał, Samuel Coward, Ulyana Piterbarg, Maciej Wolczyk, Akbir Khan, Eduardo Pignatelli, Łukasz Kuciński, Lerrel Pinto, Rob Fergus, Jakob Nicolaus Foerster, Jack Parker-Holder, and Tim Rocktäschel. 2024. [Balrog: Benchmarking agentic llm and vlm reasoning on games](#). *Preprint*, arXiv:2411.13543.
- Dongmin Park, Minkyu Kim, Beongjun Choi, Junhyuck Kim, Keon Lee, Jonghyun Lee, Inkyu Park, Byeong-Uk Lee, Jaeyoung Hwang, Jaewoo Ahn, Ameysa S. Mahabaleshwarkar, Bilal Kartal, Pritam Biswas, Yoshi Suhara, Kangwook Lee, and Jaewoong Cho. 2025. [Orak: A foundational benchmark for training and evaluating llm agents on diverse video games](#). ArXiv:2506.03610.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Anian Ruoss, Fabio Pardo, Harris Chan, Bonnie Li, Volodymyr Mnih, and Tim Genewein. 2024. Lmact: A benchmark for in-context imitation learning with long multimodal demonstrations. *arXiv preprint arXiv:2412.01441*.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. Textcaps: a dataset for image captioning with reading comprehension. In *ECCV*, pages 742–758.
- 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, and 9 others. 2024. Cradle: Empowering foundation agents towards general computer control. *arXiv preprint arXiv:2403.03186*.
- Google Gemini Team. 2023. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. [Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution](#). *Preprint*, arXiv:2409.12191.
- Xinyu Wang, Bohan Zhuang, and Qi Wu. 2025a. Are large vision language models good game players? In *International Conference on Learning Representations*.
- Ziyue Wang, Yurui Dong, Fuwen Luo, Minyuan Ruan, Zhili Cheng, Chi Chen, Peng Li, and Yang Liu. 2025b. [How do multimodal large language models handle complex multimodal reasoning? placing them in an extensible escape game](#). *Preprint*, arXiv:2503.10042.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. [Os-world: Benchmarking multimodal agents for open-ended tasks in real computer environments](#). *Preprint*, arXiv:2404.07972.
- Tianqi Xu, Linyao Chen, Dai-Jie Wu, Yanjun Chen, Zecheng Zhang, Xiang Yao, Zhiqiang Xie, Yongchao Chen, Shilong Liu, Bochen Qian, Philip Torr, Bernard Ghanem, and Guohao Li. 2024. [Crab: Cross-environment agent benchmark for multimodal language model agents](#). *Preprint*, arXiv:2407.01511.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. [Mm-vet: Evaluating large multimodal models for integrated capabilities](#). *Preprint*, arXiv:2308.02490.
- Weihao Yu, Zhengyuan Yang, Linfeng Ren, Linjie Li, Jianfeng Wang, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Lijuan Wang, and Xinchao Wang. 2024. [Mm-vet v2: A challenging benchmark to evaluate](#)

large multimodal models for integrated capabilities.
arXiv preprint arXiv:2408.00765.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.

Chi Zhang, Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2023. [Appagent: Multimodal agents as smartphone users](#). *Preprint*, arXiv:2312.13771.

Haoran Zhang, Hangyu Guo, Shuyue Guo, Meng Cao, Wenhao Huang, Jiaheng Liu, and Ge Zhang. 2024. [Ing-vp: Mllms cannot play easy vision-based games yet](#). *Preprint*, arXiv:2410.06555.

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, and 32 others. 2025. [Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models](#). *Preprint*, arXiv:2504.10479.

Appendix Overview

In the appendix, we provide the following contents:

Sec A: Introduces experimental settings and provides detailed information on the experiments, models, and their performance.

Sec B: Delineates the game selection methodology and sources, including all level designs and prompts.

Sec C: Presents specific details of ELO-based ranking system in V-MAGE.

Sec D: Provides ablation studies on pipeline settings(e.g., frame sampling strategy, resolution, etc.)

Sec E: Provides details of supplemental experimental analyses.

Sec F: Provides details on miscellaneous material, including a **statement of LLM usage**(Sec. F.1) and a discussion about broader impacts.

Sec G: Presents some case studies.

A Experiments Details

A.1 Evaluation Pipeline Details

V-MAGE employs a three-module architecture, as illustrated in Figure 2. The specific configurations and details within each module are as follows:

A.1.1 Game Module

The Game module encompasses the game simulators and their operational parameters. In V-MAGE, game simulators, representing environments developed with Pygame, are configured to present tasks that test specific visual and reasoning capabilities. To address potential confounding factors such as API latency and computational constraints, V-MAGE employs a frame-pausing mechanism during model inference. This ensures that the game environment remains static while the model processes visual inputs and generates actions, effectively decoupling timing evaluation from raw inference speed.

Regarding real-time execution and inference delays, the frame-pausing mechanism ensures fairness across models. While current models may not consistently achieve real-time inference due to API or GPU limitations, V-MAGE automatically pauses the game environment during model processing. This approach isolates the model’s temporal reasoning ability (strategic "when to act") from infrastructure-related delays, enabling a focused assessment of capabilities such as Timing. In V-MAGE, Timing is explicitly designed to evaluate

the model’s ability to choose the optimal timing of actions, not the system’s response speed.

The frames per second (FPS) for all our games is 30. In our standard benchmark setting, we use a sample rate of 3. This corresponds to the model making a decision approximately every 100 milliseconds (30 FPS / 3) in the game.

A.1.2 Agent Module

Researchers can modify the agent’s operational mode by adjusting the configuration file. This includes altering historical strategies, such as employing a longer history of steps or sparsely sampling more distant historical records to provide different contextual information to the models. However, for fairness in evaluation and due to resource constraints, we present results obtained using only the naive configuration as a baseline in our main findings. This allows for standardized comparisons across different MLLMs.

In the baseline agent strategy, we utilize historical information comprising three preceding screenshots and the model’s reasoning outputs associated with these screenshots, combined with the current game screenshot, to form the <history> and <image_current_step> components (as shown in Sec. B.4) that are incorporated into the model’s input framework.

The <history> component is structured as follows:

This screenshot is <step> steps before the current step of the game. After this frame, your reasoning message was <history_reasoning>. After the action was executed, the game info was <history_action_info>.

Upon obtaining the model’s output, the system parses the structured response format specified in the prompt to systematically extract both the cognitive reasoning log and action parameters. These extracted elements are subsequently stored and fed back to the game module through a standardized data pipeline. This process ensures consistent interpretation of the model’s decision-making trajectory while maintaining synchronization with the game state.

Taking an input from the actual evaluation process in the **Tempest Run** as an example:

Assume you are playing a PC game called 'Tempest Run'.

You need to control a character who moves through a three-dimensional space inside a futuristic tunnel filled with various obstacles and enemies. Your goal is to navigate through the tunnel, avoid or overcome obstacles, and run as far as possible. Avoid colliding with red spikes, purple walls, or failing to deal with green enemies.

Use the optimal combination of movements to progress through the tunnel smoothly and efficiently. Monitor the character's position relative to obstacles and react appropriately to avoid losing progress. Now, I will give you some history screenshots in the game for decision making.

<image>

This screenshot is three steps before the current step of the game. After this frame, your reasoning message was "'observation': 'The character is running forward in a narrow tunnel with blue lines forming the walls and ceiling. Red spikes are visible on the ground ahead, and green enemies are further down the tunnel. The path appears clear with no immediate obstacles or enemies to avoid.', 'reasoning': 'Maintaining the current trajectory with no action is optimal to continue moving straight towards the green enemy and prepare for a potential slide to eliminate it.', 'action': 'NONE'". After the action was executed, the game info was "Game is running."

...

<image>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show the character's previous movements.

Important notes: 1. Use JUMP to jump over red spikes on the ground. 2. Use SLIDE to duck and kick green enemies to eliminate them.

...

You should think step by step and respond with the following format, remember to respond with plain text without any special characters or symbols, DO NOT respond in markdown or Latex or any other format.

Response:

Observation: ... (Describe the character's current position and nearby obstacles or enemies.)

Reasoning: ... (Think step by step and explain how you choose the action.)

Action: ... (Choose ONE of the six actions to control the character. Do NOT add any other words.)

A.1.3 Model Module

The Model module is primarily responsible for model deployment and parameter control. In addition to closed-source models accessed via APIs, we deployed open-source models on an **8xV100 GPU Azure cluster** and an **8xA100 GPU Azure cluster** for our experiments, utilizing the **vLLM** library for efficient serving. For text output generation across all models, we standardized the decoding parameters by setting **top_p=0.9** and **temperature=0.8**.

A.2 Detailed Statistics

A.2.1 Visualization

Figure 7 compares the performance of leading MLLMs and human players across different game levels.

A.2.2 Score-based Performance

Cross-task result analysis reveals the limitations of parameter scaling: In Race Level 1 (with historical frame input), Qwen2VL showed a 429% improvement in score when scaling from 7B to 72B (from 10.43 to 55.19), reaching about 55% of the human baseline score. However, in more complex tasks such as Tempestrun Level 4, InternVL2.5-78B (199.78 points) only improved by 14.4% compared to its 8B version (174.58 points), still achieving only 25% of the human score (800 points). This suggests that parameter scaling cannot compensate for the inherent ability gap in complex dynamic tasks. The detailed scores are presented in Tables 4, 5, 6, 7, and 8.

A.2.3 Additional Indicators

Due to the dynamic game environments inherent in the V-MAGE evaluation process, certain levels may necessitate a considerable number of tokens during assessment. In this section, using the **Qwen2.5VL-72B** model deployed with **vLLM** as an illustrative example, we provide the statistically averaged frame counts (equal to the number of frames between two neighboring interactions multiplied by the number of interactions) and the average input and output token consumption per game round, serving as a reference. The detailed statistics are presented in Table 9, 10, 11, and 12.

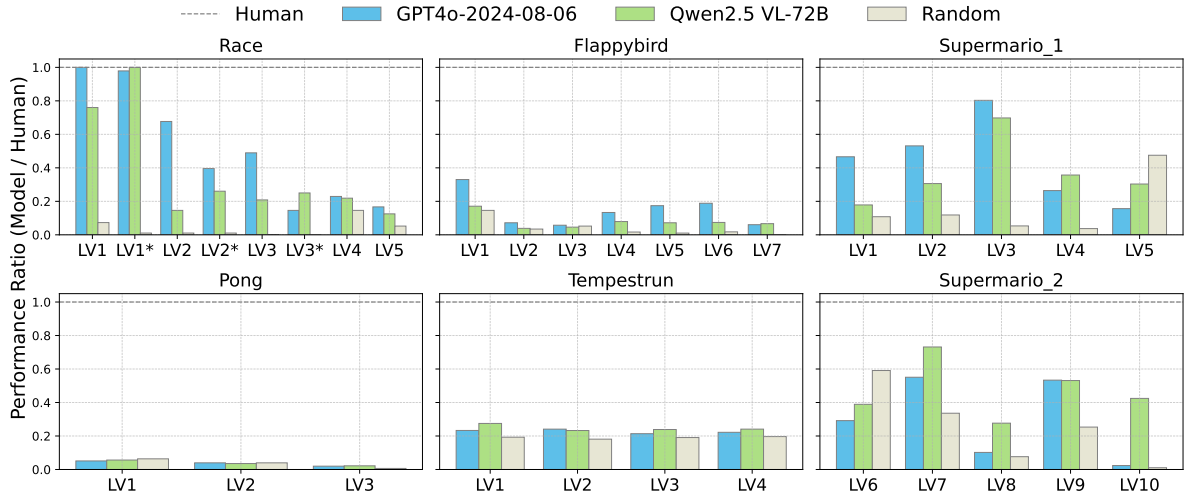


Figure 7: **The MLLM trails humans by a large margin in all six games.** The levels with an asterisk (*) represent ‘no history’. Detailed performance metrics for each model across individual game levels are provided in Appendix A.2 (Tables 4-8).

Table 4: Performance analysis based on average scores in Race

Level	GPT-4o	Gemini 2.0-flash	Qwen2.5VL 72B	Qwen2VL 7B	Qwen2VL 72B	InternVL2.5 8B	InternVL2.5 78B	Random	Human
Level1	99.99	35.41	76.01	10.43	55.19	28.12	64.56	7.30	100.00
Level1 No History	97.87	98.91	99.95	87.46	97.87	89.54	99.99	1.06	100.00
Level2	67.68	7.30	14.59	0.00	1.06	3.14	15.63	1.06	100.00
Level2 No History	39.57	22.92	26.04	1.06	23.96	5.22	26.04	1.06	100.00
Level3	48.94	5.22	20.84	4.18	7.30	6.26	11.47	0.02	100.00
Level3 No History	14.59	4.18	25.00	4.18	16.67	11.47	22.92	0.00	100.00
Level4	22.92	6.26	21.88	0.02	2.10	0.00	2.10	14.59	100.00
Level5	16.67	8.34	12.51	13.55	4.18	20.84	7.30	5.22	100.00

Table 5: Performance analysis based on average scores in Pong

Level	GPT-4o	Gemini 2.0-flash	Qwen2.5VL 72B	Qwen2VL 7B	Qwen2VL 72B	InternVL2.5 8B	InternVL2.5 78B	Random	Human
Level1	0.51	0.54	0.56	0.59	0.54	0.68	0.77	0.63	10.00
Level2	0.39	0.41	0.35	0.31	0.33	0.31	0.38	0.39	10.00
Level3	0.19	0.32	0.21	0.18	0.20	0.18	0.15	0.06	10.00

Table 6: Performance analysis based on average scores in Supermario

Level	GPT-4o	Gemini 2.0-flash	Qwen2.5VL 72B	Qwen2VL 7B	Qwen2VL 72B	InternVL2.5 8B	InternVL2.5 78B	Random	Human
Level10	18.77	108.30	339.57	12.51	29.16	14.61	80.19	8.36	800.00
Level1	372.85	109.41	142.76	33.41	216.67	69.83	203.12	86.50	800.00
Level2	424.92	127.17	244.78	102.12	338.47	102.12	186.48	94.83	800.00
Level3	802.99	429.10	697.91	188.54	565.46	286.44	610.26	53.19	1000.00
Level4	369.76	251.07	499.89	112.53	346.84	151.09	447.84	52.15	1400.00
Level5	125.08	258.33	242.72	232.29	192.75	209.41	433.23	380.13	800.00
Level6	233.36	325.96	311.36	324.96	296.79	267.70	344.74	472.78	800.00
Level7	440.66	527.96	585.21	161.48	490.48	220.86	491.52	268.74	800.00
Level8	91.75	211.43	248.96	51.13	179.20	76.09	168.74	68.79	900.00
Level9	693.56	594.67	690.46	162.62	508.24	201.12	756.02	329.19	1300.00

Table 7: Performance analysis based on average scores in Flappybird

Level	GPT-4o	Gemini 2.0-flash	Qwen2.5VL 72B	Qwen2VL 7B	Qwen2VL 72B	InternVL2.5 8B	InternVL2.5 78B	Random	Human
Level1	3.30	2.38	1.70	0.76	0.47	1.20	1.54	1.45	10.00
Level2	0.71	0.47	0.38	0.20	0.12	0.36	0.39	0.34	10.00
Level3	0.57	0.41	0.45	0.20	0.35	0.33	0.43	0.52	10.00
Level4	1.33	1.50	0.79	1.52	0.38	1.43	0.64	0.16	10.00
Level5	1.74	1.38	0.71	1.44	0.51	1.20	0.49	0.10	10.00
Level6	1.88	1.05	0.73	1.62	0.56	1.14	0.66	0.17	10.00
Level7	0.60	0.07	0.66	0.03	0.14	0.00	0.13	0.00	10.00

Table 8: Performance analysis based on average scores in Tempestrun

Level	GPT-4o	Gemini 2.0-flash	Qwen2.5VL 72B	Qwen2VL 7B	Qwen2VL 72B	InternVL2.5 8B	InternVL2.5 78B	Random	Human
Level1	466.25	478.35	549.98	446.92	519.22	444.71	475.22	385.72	2000.00
Level2	361.44	356.05	349.06	352.76	370.13	327.38	333.37	271.65	1500.00
Level3	213.73	197.91	238.74	208.75	220.21	197.71	216.64	190.71	1000.00
Level4	177.60	201.67	192.79	182.91	195.19	174.58	199.78	157.17	800.00

Depending on the differences in the models and the randomness of the games and reasoning, as well as other further experiments, the full research project may require **more** compute than the experiments reported here. The time of execution of the experiment depends on the network environment and computational power.

A.3 Inconsistency Between ELO and Performance Ratio Rankings

As shown in Table 2 in the main text, ELO and Performance Ratio sometimes do not align in rankings.

We examine **Keye-VL-8B-Preview** and **Qwen2.5-VL-7B-Instruct**, with **LLaVA-v1.6-Mistral-7B** as a control. The experimental results are presented in Table 13.

In Pong, Qwen shows higher ELO but lower average ratio. We analyzed level-wise scores and variances to explore this. The variance is calculated as: $\text{variance} = \frac{\sum_{i=1}^n (\text{score}_i - \text{score})^2}{n}$. The detailed level-wise results and variance statistics are reported in Table 14.

Keye’s higher variance across all Pong levels indicates unstable performance, where high-scoring outliers mask frequent weak results. In the ELO system, this instability leads to more losses against a consistent opponent, resulting in a lower rating despite a competitive average score.

Additionally, current models perform poorly on Pong, with ratios tightly clustered in the 0–10% range. When calculating the performance ratio by averaging across games, minor differences in Pong (3.5% vs. 3.9%) are overshadowed by larger gaps in other games (21% vs. 24%). The ELO system, in contrast, is based on the aggregate outcomes of all pairwise matchups. The ELO rating boost from a consistent pattern of wins in Pong is just as significant as from wins in any other game. This demonstrates that ELO is more robust in **fairly** assessing a model’s holistic capabilities across tasks with imbalanced performance levels.

We also observed that in terms of Response Format Accuracy, GPT-4o is slightly lower than Gemini model (by 0.04%), and InternVL2.5-78B is slightly lower than Qwen2-VL-72B (by 0.25%). This may also be an influencing factor.

B Games in V-MAGE

B.1 Principles and Standards for Game Selection

Simplified and unrealistic considerations.

While the simplified visuals in these games differ from real-world scenes, empirical evidence demonstrate that MLLMs comprehend core game semantics (objectives, rules, entities) despite stylistic simplifications. Performance limitations primarily emerge from perceptual inaccuracies (e.g., dynamic object tracking) and multi-step reasoning deficiencies rather than misinterpretation. V-MAGE therefore focuses more on **precise evaluation** than **visual realism** to drive targeted improvements in visual reasoning.

Selection criteria. The five games in V-MAGE share critical characteristics (e.g., non-textual state representation, free-form gameplay, and continuous-space environments) while offering diverse challenges.

Our current minimal set covers four **2D** game types, as summarized in Table 15.

The Linear Process implies that the game’s progression is, to some extent, enforced. In PongGame, the ball’s movement direction is determined by the game environment, requiring the model to move paddles on both sides to catch the ball, while in FlappyBird, the forward movement of the bird is compulsory, with the model controlling the height to navigate through pipes. OpenPlanning, in contrast, is relatively more open-ended. In RaceGame, the model can freely control the car’s movement and direction on a plane to reach a trophy. In SuperMario, the model can move and jump in a relatively open environment to collect rewards and earn points.

For **3D** environments, we selected **Tempest Run** for its streamlined visual elements.

V-MAGE’s flexible framework allows seamless integration of new PyGame-based environments. For instance, Tempest Run (one of our five games) was sourced from PyWeek[3], a community-driven game jam with thousands of open-source entries. This demonstrates our framework’s capacity to incorporate externally developed, human-designed games. We provide APIs to wrap new games into V-MAGE’s evaluation pipeline. This allows researchers to easily integrate additional games.

We will continue expanding the benchmark with more diverse titles that meet our selection criteria (e.g., Player vs Player (PVP) games) and will

Table 9: **SuperMario** Average Frames and Tokens Consumed

Metric	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7	Level 8	Level 9	Level 10	All
Average Frames	400	655.56	1000	641.2	234.00	300	300	148.06	504.35	950.45	5133.62
Average Prompt Tokens	150004.78	253457.38	379649.92	266125.09	100595.30	112241.57	111288.00	54231.22	192642.31	361208.32	1981443.89
Average Completion Tokens	10054.86	18437.53	25428.29	18242.65	7075.53	7639.30	7314.46	3710.36	13492.79	24197.61	135593.38

Table 10: **Race** Average Frames and Tokens Consumed

Metric	Level 1	Level 2	Level 3	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	All
	No History	No History	No History							
Average Frames	12.66	15.39	16.66	29.20	30.69	32.14	58.07	98.06	32.46	325.33
Average Prompt Tokens	1738.83	2309.85	2562.15	12317.22	14044.87	14934.32	31164.03	54346.81	17399.59	255136.23
Average Completion Tokens	275.35	531.41	595.91	693.11	937.66	1060.77	2243.55	3898.42	1517.55	20798.72

Table 11: **FlappyBird** Average Frames and Tokens Consumed

Metric	Level 1	Level 2	Level 3	Level 4	Level 5	Level 6	Level 7	All
Average Frames	224.73	133.34	76.49	153.11	153.87	152	143.28	1036.82
Average Prompt Tokens	98273.78	57332.59	32326.91	65853.39	66500.27	65322.27	56528.63	442137.84
Average Completion Tokens	9979.17	5772.39	3319.76	7142.32	7309.31	7082.73	6259.62	46865.30

Table 12: **TempestRun** and **PongGame** Average Frames and Tokens Consumed

Metric	Level 1	Level 2	Level 3	Level 4	All	Level 1	Level 2	Level 3	All
	Average Frames	173.58	92.70	38.98	28.72	333.98	221.79	83.98	47.00
Average Prompt Tokens	108291.56	57096.18	33218.84	22874.80	237820.07	136254.76	50056.53	26981.30	213292.59
Average Completion Tokens	7000.84	3799.28	2316.48	1619.33	15942.53	10998.40	4064.67	2208.39	17271.46

Table 13: Elo Scores and Average Performance Ratios (E/R) Across Games.

	Race(E/R)	SuperMario(E/R)	Pong(E/R)	FlappyBird(E/R)	TempestRun(E/R)
Qwen2.5-VL-7B-Instruct	1487/0.120	1459/0.239	1503/0.035	1431/0.030	1485/0.210
Keye-VL-8B-Preview	1487/0.118	1430/0.217	1495/0.039	1450/0.044	1513/0.239
LLaVA-v1.6-Mistral-7B	1462/0.051	1374/0.127	1494/0.035	1489/0.077	1379/0.169

open-source both the codebase and detailed documentation to facilitate community contributions.

B.2 Design and Implementation

As previously mentioned, V-MAGE enhances the diversity of the evaluation environment by expanding it through level design. Tables in this section detail the settings, rewards, and design objectives

for each game’s levels. For more comprehensive visual comparisons and prompt information, please refer to Appendix B.4.

Race Game is a skill-based driving game where the objective is to control a car through a maze-like track to reach the trophy while avoiding obstacles. The car is represented by a red or white vehicle with a visible front and back, while the trophy

Table 14: Pong Scores by Level (Avg: average score, Var: variance).

	L1 Avg	L1 Var	L2 Avg	L2 Var	L3 Avg	L3 Var
Qwen2.5-VL-7B-Instruct	0.48	0.50	0.37	0.25	0.20	0.18
Keye-VL-8B-Preview	0.68	0.67	0.26	0.33	0.23	0.36
LLaVA-v1.6-Mistral-7B	0.48	0.58	0.29	0.26	0.29	0.34

Table 15: 2D Game Taxonomy in V-MAGE

	XY-axis	XZ-axis
Linear Process	PongGame	FlappyBird
Open Planning	RaceGame	SuperMario

is shown as a golden cup icon. The surrounding white-lined boundaries represent walls, which the car must avoid. For the overall observation and action spaces of the game, including the task and reward definitions, please refer to Table 16.

Each level in Race has a different set of rules and challenges. As presented in Table 17, we manually designed six levels. Levels 1–3 use a *map-view perspective* ("map" view), where models adjust absolute coordinates. The four types of movement operations directly translate the vehicle on the map according to the direction of action. Conversely, Levels 4–6 shift to a *first-person perspective* ("car" view), where the observation is centered on the vehicle and movements are performed from the vehicle’s perspective, requiring interpretation of velocity vectors and acceleration constraints. Furthermore, acceleration is introduced in the high-difficulty levels, which further expands the observation space. This requires the model to extract more information from the visual input, including current speed and acceleration, in order to perform rational reasoning.

SuperMario is a two-dimensional side-scrolling platformer where the player controls the character Mario navigating through environments populated with various platforms, enemies, and obstacles. The goal is to traverse the level, collect coins, evade or defeat enemies, and reach the flagpole at the stage’s conclusion. Players must avoid falling off platforms, colliding with enemies, or being struck by obstacles. Successful gameplay in-

volves employing optimal movement combinations for smooth and efficient progression, alongside monitoring Mario’s position relative to environmental elements. The task and reward definitions are shown in Table 18.

To provide a comprehensive evaluation of MLLMs’ visual reasoning and planning capabilities, SuperMario features ten levels with configurations detailed in Table 19. These levels vary in enemy count, coin and coinbox quantities, maximum allowed rounds, and specific gameplay mechanics or focuses. Of these, level 10 serves as a standard human-difficulty benchmark, providing a 1:1 replica of the original Super Mario game’s World 1-1 stage.

Flappy Bird is a widely recognized side-scrolling mobile game serving as a common benchmark in reinforcement learning. The objective is to control a bird’s vertical movement to navigate through a continuous series of horizontal gaps within vertically oriented pipes. Successful traversal of a pipe pair increments the player’s score, while collision with any pipe or the ground constitutes a terminal state, ending the game. The game mechanic involves a constant downward gravitational pull, counteracted by discrete upward ‘flaps’ initiated by the player. The task and reward definitions are shown in Table 20.

Given the high difficulty of human-standard levels for MLLMs, we designed seven levels with progressive difficulty. Specifically, as presented in Table 21, levels 1-3 constitute a simplified game environment where the gravity factor is removed, and height is controlled via UP and DOWN actions to navigate through the pipes. Levels 4-6 are based on the standard difficulty but incorporate a ‘KEEP’ option, enabling the model to maintain the bird’s altitude through this action. Within the same difficulty tier, levels are differentiated by varying the bird’s forward speed and the pipe gap width. Level 7 represents the standard human game difficulty,

Table 16: Race Environment Details (* means potentially requires observation).

Observation space	Action Space	Task	Reward
Car Position Trophy Position Obstacle Position* Speed* Acceleration* Facing Angle*	UP, DOWN, LEFT, RIGHT	Move the car to reach the trophy	+100 Success +0 Timeout +0 Destroyed

Table 17: Race Level Configurations

Level	View	bstacle	Initial Direction	Acceleration	Max Rounds	Sample Frames
1	Map	No	-	No	100	1
2	Map	Yes	-	No	150	1
3	Map	Yes	-	No	150	1
4	Car	No	Vertical(up)	Low	150	3
5	Car	No	Horizontal(random)	Mid	150	3
6	Car	Yes	Vertical(up)	Mid	150	1

Table 18: SuperMario Environment Details.

Observation space	Action Space	Task	Reward
Mario Position Platforms Position Enemies Position Obstacles Position	UP, UP+LEFT, UP+RIGHT, LEFT, RIGHT, NONE	Collect coins and evade or defeat enemies	+100 for collecting a coin +100 for defeating a Goomba Penalties for falling or collisions

Table 19: SuperMario Level Configurations

Level	Enemy count	Coin Count	CoinBox Count	Max Rounds	Gameplay
1	0	6	2	400	Common
2	2	6	2	1000	Common
3	0	17	4	1000	Long History (Two ways)
4	2	17	4	1000	Long History (Two ways)
5	3	8	0	300	Left or Right
6	0	13	0	300	Left or Right
7	0	8	0	300	Left or Right
8	0	12	0	1000	Jump Only
9	5	8	0	1000	Jump and Enemy
10	12	0	9	5000	Classic W1-1

Table 20: Flappy Bird Environment Details (* means only available at certain levels).

Observation space	Action Space	Task	Reward
Bird Position	UP	Maneuver the bird to avoid hitting the pipes	+1 per pipe pair passed +0 Collision
Bird Velocity	NONE		
Next Pipe Distance	DOWN*		
Gap VerticalPosition	KEEP*		

retaining the original game settings.

Pong Game is a classic two-player adversarial game. The objective is to control the paddles on the left and right sides of the screen to return the ball, preventing it from passing one’s own paddle while simultaneously attempting to make the ball pass the opponent’s paddle. One point is awarded to the player for each successful return of the ball. The final score is the sum of both players’ scores. The task and reward definitions are shown in Table 22.

Considering the challenges MLLMs face in tracking and temporal tasks, we designed levels with varying difficulty. As shown in Table 23, difficulty for Levels 1-3 is differentiated by adjusting the paddle width and the speed of the ball. Within the same level, the initial position of the ball is randomized, but the relative difficulty remains consistent.

Tempest Run is a third-person perspective 3D runner game where the player controls a character moving within a futuristic tunnel filled with various obstacles and enemies. The objective is to navigate through the tunnel, avoiding or overcoming impediments, and to run as far as possible. Players must specifically avoid colliding with red spikes, purple walls, or failing to manage green

enemies. Successful gameplay requires employing optimal combinations of movements for smooth and efficient tunnel traversal, alongside monitoring the character’s position relative to obstacles and reacting appropriately. The task and reward definitions are shown in Table 24.

To evaluate MLLMs’ visual comprehension and reactive capabilities within a dynamic 3D environment, Tempest Run includes four levels of varying difficulty. As outlined in Table 25, Levels 1-4 are primarily differentiated by parameters including role speed, cell length (denoting the distance between environmental segments), and random rate (controlling obstacle spawning frequency). These parameters collectively influence the pace of barrier generation and the overall visual complexity of the tunnel environment, thereby varying the level of challenge. Within the same level, the positioning of environmental elements is randomized, while maintaining consistent relative difficulty.

B.3 Original Sources

Thanks to the open-source community, we are able to leverage existing game codebases to build our benchmark. The codebases used in this work are summarized in Table 26. Most of these projects are

Table 21: FlappyBird Level Configurations

Level	Gravity	Availability of "DOWN"	Availability of "KEEP"	Others
1 - 3	No	Yes	Yes	Distinguished by gap clearance and speed
4 - 6	Yes	No	Yes	
7	Yes	No	No	Human Standard

Table 22: Pong Game Environment Details.

Observation space	Action Space	Task	Reward
Left Paddle Position Right Paddle Position Ball Position Ball Trajectory	LEFTUP LEFTDOWN RIGHTUP RIGHTDOWN NONE	Track the ball’s trajectory and maneuver the left and right paddles to intercept and return the ball.	+1 per successful hit +0 if ball passes paddle

Table 23: Pong Game Level Configurations

Level	Paddle Width	Ball Speed	Ball Size	Others
1	Big	Slow	Big	Ball initial position randomly changes.
2	Mid	Mid	Mid	
3	Small	Fast	Small	

Table 24: Tempest Run Environment Details.

Observation space	Action Space	Task	Reward
Current Character State Nearby Obstacles Position Nearby Obstacles Type Visual Information Quantity	JUMP, LEFT, RIGHT, SLIDE, RISE, NONE	Perform corresponding actions to avoid or destroy obstacles.	Score increases with distance run.

licensed under the MIT License; for those without an explicit license, we have obtained formal permission from the original authors for their use in this research.

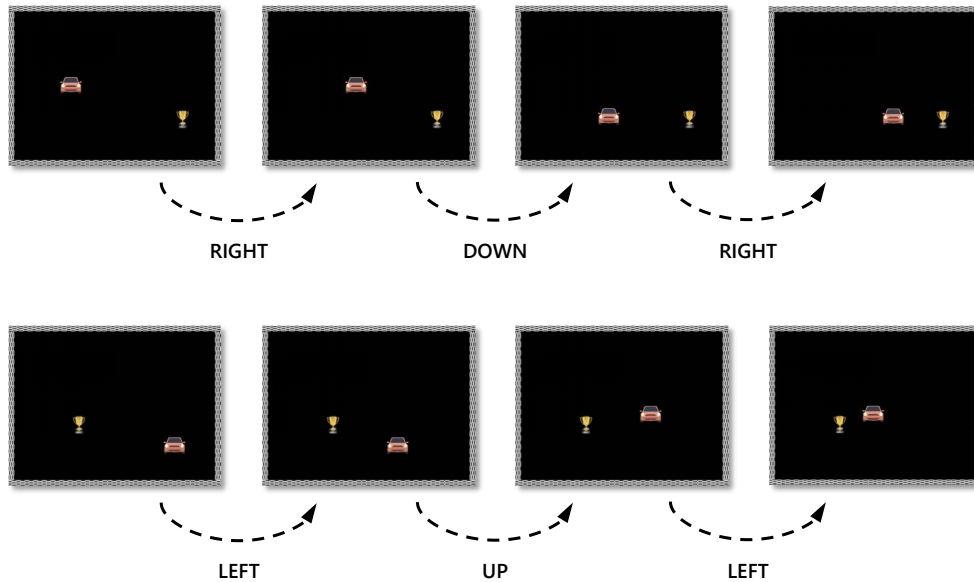
In most cases, the original codebases lacked comprehensive difficulty settings and level designs suitable for systematic evaluation. We therefore modified the default human-oriented game configurations to adapt them for benchmarking purposes, while meticulously designing a diverse set of challenging levels to ensure rigorous assessment.

B.4 Games and Prompts

All the games have been modified based on publicly available code. The detailed level designs and corresponding prompts for all games are illustrated in Figures 8–12, 13–16, 17–19, 20, and 21.

RaceGame

Level 1



Prompt:

You are playing the Racing Game.

<\$history\$>

<\$image_current_step\$>

These pictures depicts some game scenes with a trophy and a car.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

First, describe the relative position of the trophy with respect to the car currently.

Now, I want to move the car to reach the trophy. Please tell me which direction the car should move to reach the trophy.

You can choose four directions: UP, DOWN, LEFT, RIGHT.

You should think step by step and response with the following format, rememer to response in plain text format, not in markdown or any other format:

Description: ...

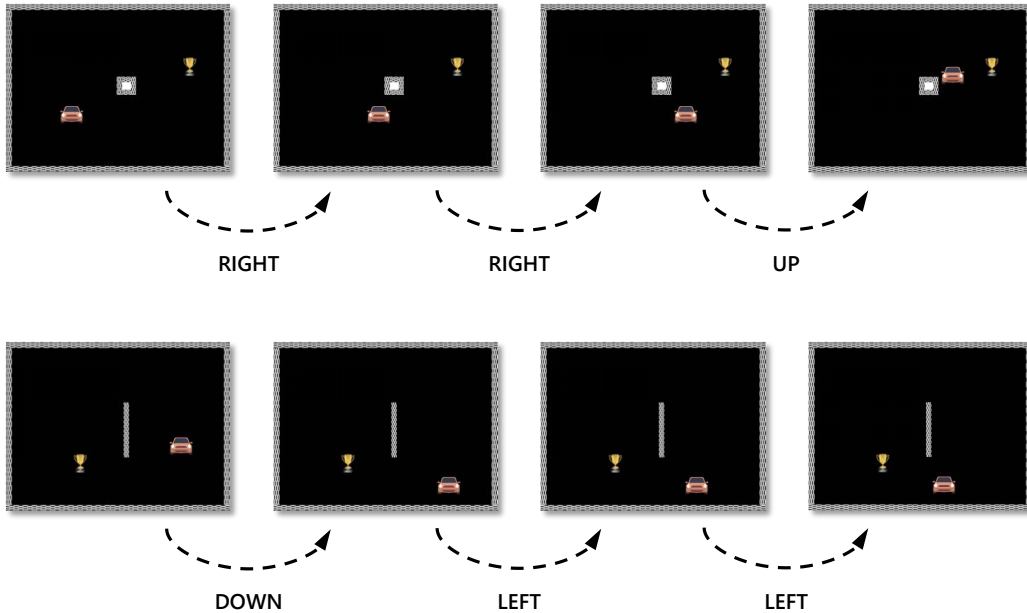
Reasoning: ...

Action: ... (one of the four directions)

Figure 8: **RaceGame Level 1: Level Design and Prompt Overview.** The images showcase the scene from Level 1, illustrating the level design and corresponding prompt. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty.

RaceGame

Level 2 - 3



Prompt:

You are playing the Racing Game.

<\$history\$>

<\$image_current_step\$>

These pictures depicts some game scenes with a trophy and a car.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

First, describe the relative position of the trophy with respect to the car currently.

Then, judge if there are obstacles between the car and the trophy. If there are, describe them and how they affect the car's movement.

Now, I want to move the car to reach the trophy. Please tell me which direction the car should move to reach the trophy. Remember, if the car collides with an obstacle, it will be game over.

You can choose four directions: UP, DOWN, LEFT, RIGHT.

You should think step by step and response with the following format, remember to response in plain text format, not in markdown or any other format:

Description: ...

Reasoning: ...

Action: ... (one of the four directions)

Figure 9: **RaceGame Level 2-3: Level Design and Prompt Overview.** The images showcase the scene from Level 2-3, illustrating the level design and corresponding prompt. Elements in the scene will randomly change their initial positions while maintaining consistent relative difficulty.

RaceGame

Level 1 - 3 No History



Prompt:

You are playing the Racing Game.

<\$image_current_step\$>

These pictures depicts some game scenes with a trophy and a car.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

First, describe the relative position of the trophy with respect to the car currently.

(Only for Level 2/3)

Then, judge if there are obstacles between the car and the trophy. If there are, describe them and how they affect the car's movement.

Now, I want to move the car to reach the trophy. Please tell me which direction the car should move to reach the trophy. Remember, if the car collides with an obstacle, it will be game over.

You can choose four directions: UP, DOWN, LEFT, RIGHT.

You should think step by step and response with the following format, rememer to response in plain text format, not in markdown or any other format:

Description: ...

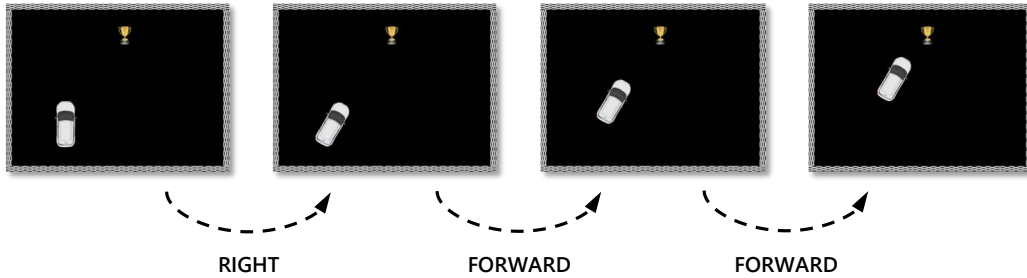
Reasoning: ...

Action: ... (one of the four directions)

Figure 10: **RaceGame Level 1-3 No History: Level Design and Prompt Overview.** The images showcase the scene from Level 1-3 No History, illustrating the level design and corresponding prompt. Elements in the scene will randomly change their initial positions while maintaining consistent relative difficulty. Same as the original levels except the input sequence has been changed to the single image.

RaceGame

Level 4



Prompt:

Assume you are playing a PC game called 'Racing Game'.

<\$history\$>

<\$image_current_step\$>

These pictures depicts some game scenes with a trophy and a car.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

First, describe the relative position of the trophy with respect to the car currently, and judge if there are obstacles between the car and the trophy. If there are, describe them and how they affect the car's movement.

Then, I want you to move the car to reach the trophy. Remember, if the car collides with an obstacle, it will be game over.

You should imagine you are the driver of the car, which means you are facing the same direction as the car. Then you should think about the car's movement from the driver's perspective.

To do this, you should judge if the car is moving, and if so, in which direction it is moving. You should also consider the car's current position and the direction it is facing.

You should think step by step and make some reasoning about the car's movement.

You can choose four actions: FORWARD, BACKWARD, LEFT, RIGHT.

Response with the following format, remember to response in plain text format, not in markdown or any other format:

Description: ...(including the relative position, obstacles, and the car's movement)

Reasoning: ...

Action: ... (one of the four actions)

Figure 11: **RaceGame Level 4: Level Design and Prompt Overview.** The images showcase the scene from Level 4, illustrating the level design and corresponding prompt. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty.

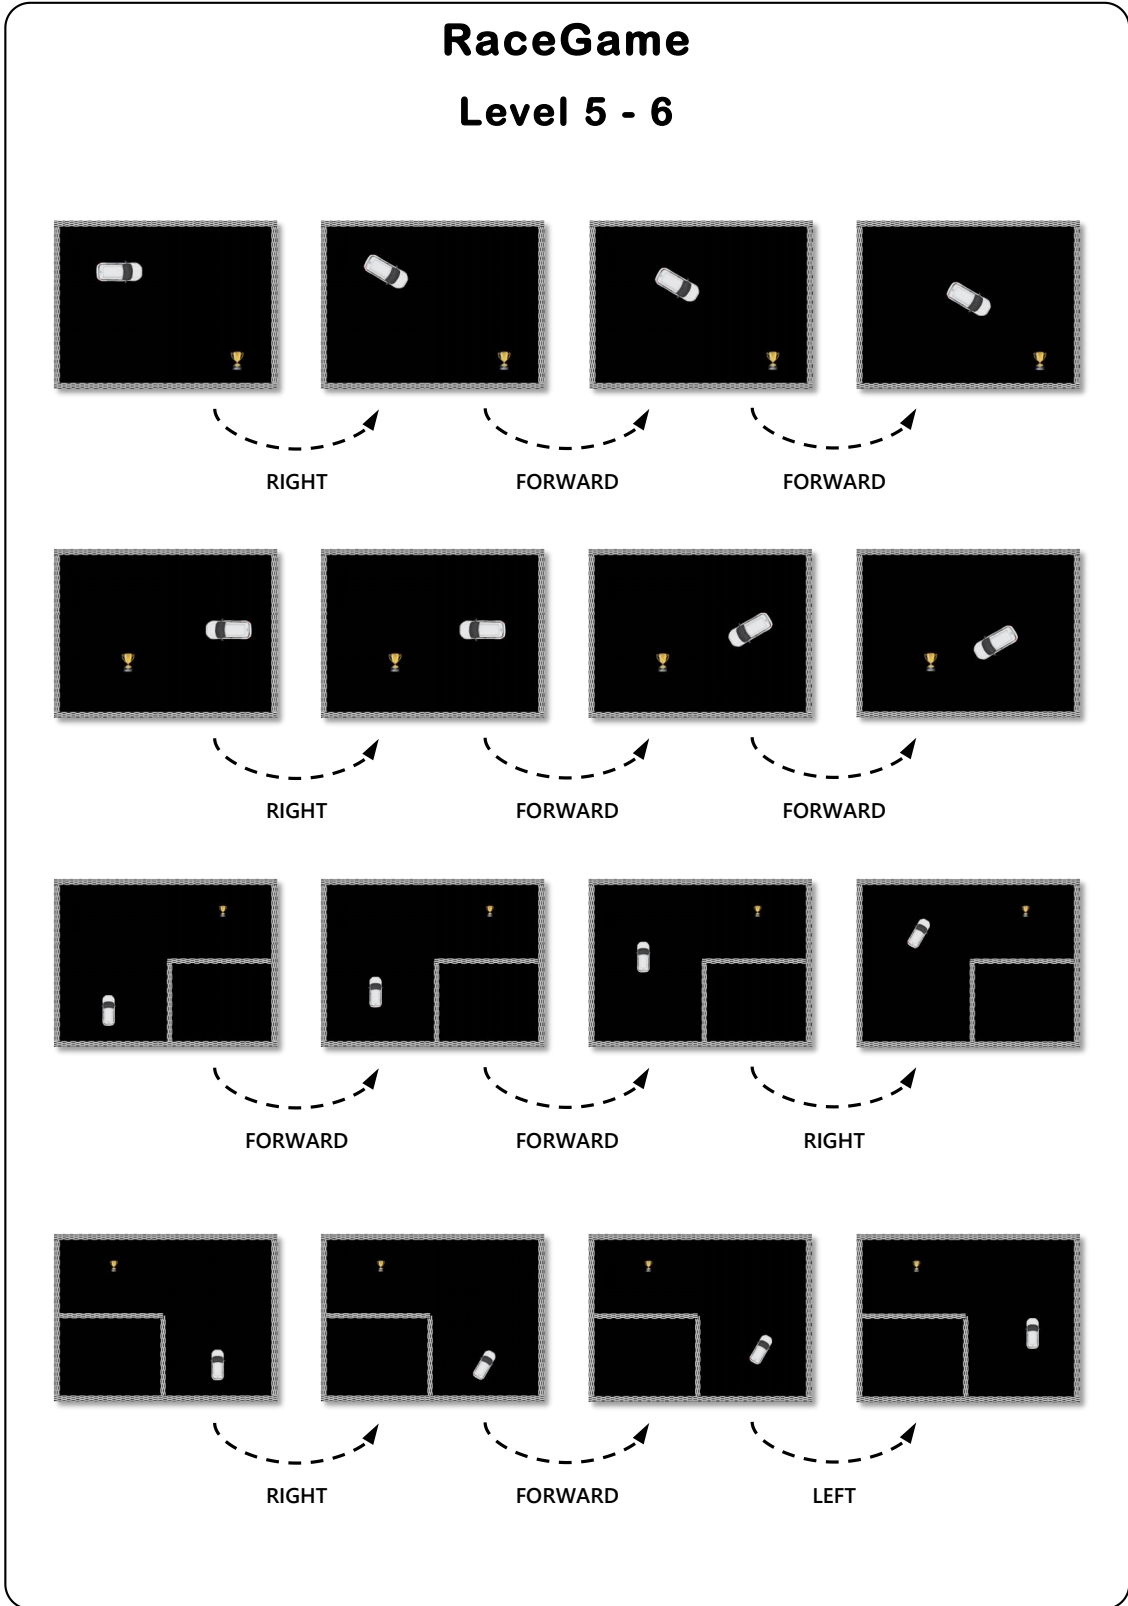


Figure 12: **RaceGame Level 5-6: Level Design and Prompt Overview.** The images showcase the scene from Level 5-6, illustrating the level design and corresponding prompt. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty. **The prompt is the same as in Level 4.**

SuperMario

Level 1



UP+RIGHT

RIGHT

UP+RIGHT

Prompt:

Assume you are playing a PC game called 'Super Mario'.

You need to control Mario, who can move in a two-dimensional space consisting of various platforms, enemies, and obstacles. Your goal is to navigate through the level, collect coins, avoid or defeat enemies, and reach the flagpole at the end of the stage. Avoid falling off the platforms, colliding with enemies, or being hit by obstacles.

Use the optimal combination of movements to progress through the level smoothly and efficiently. Monitor Mario's position relative to obstacles, platforms, and enemies.

<\$history\$>

<\$image_current_step\$>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show Mario's previous movements.

****Important notes:****

1. Mario can jump (actions involving UP) only if he is on the ground or on a solid surface like a platform or pipe.
2. If Mario is in mid-air, he can only use LEFT or RIGHT to adjust his position, or NONE to continue falling or moving with momentum.

You can make six types of actions to control Mario:

1. UP: Makes Mario jump upward (only available when Mario is on the ground or solid platforms).
2. LEFT: Moves Mario left.
3. RIGHT: Moves Mario right.
4. UP+LEFT: Makes Mario jump upward and left simultaneously (only available when on the ground or solid platforms).
5. UP+RIGHT: Makes Mario jump upward and right simultaneously (only available when on the ground or solid platforms).
6. NONE: No new action is performed; Mario continues to be affected by gravity (if airborne) or momentum from previous movements.

Note that DOWN has no effect and cannot be used, so you should never attempt to use it.

You should think step by step and respond with the following format, remember to respond with plain text without any special characters or symbols, DO NOT respond in markdown or Latex format.

Response:

Observation: ... (Describe Mario's current position, nearby platforms, enemies, and obstacles.)

Reasoning: ... (Think step by step and explain how you choose the action.)

Action: ... (Choose one of the six actions to control Mario. Do NOT add any other words.)

Figure 13: **SuperMario Level 1: Level Design and Prompt Overview.** The images showcase the scene from Level 1, illustrating the level design and corresponding prompt.

SuperMario

Level 2 - 4

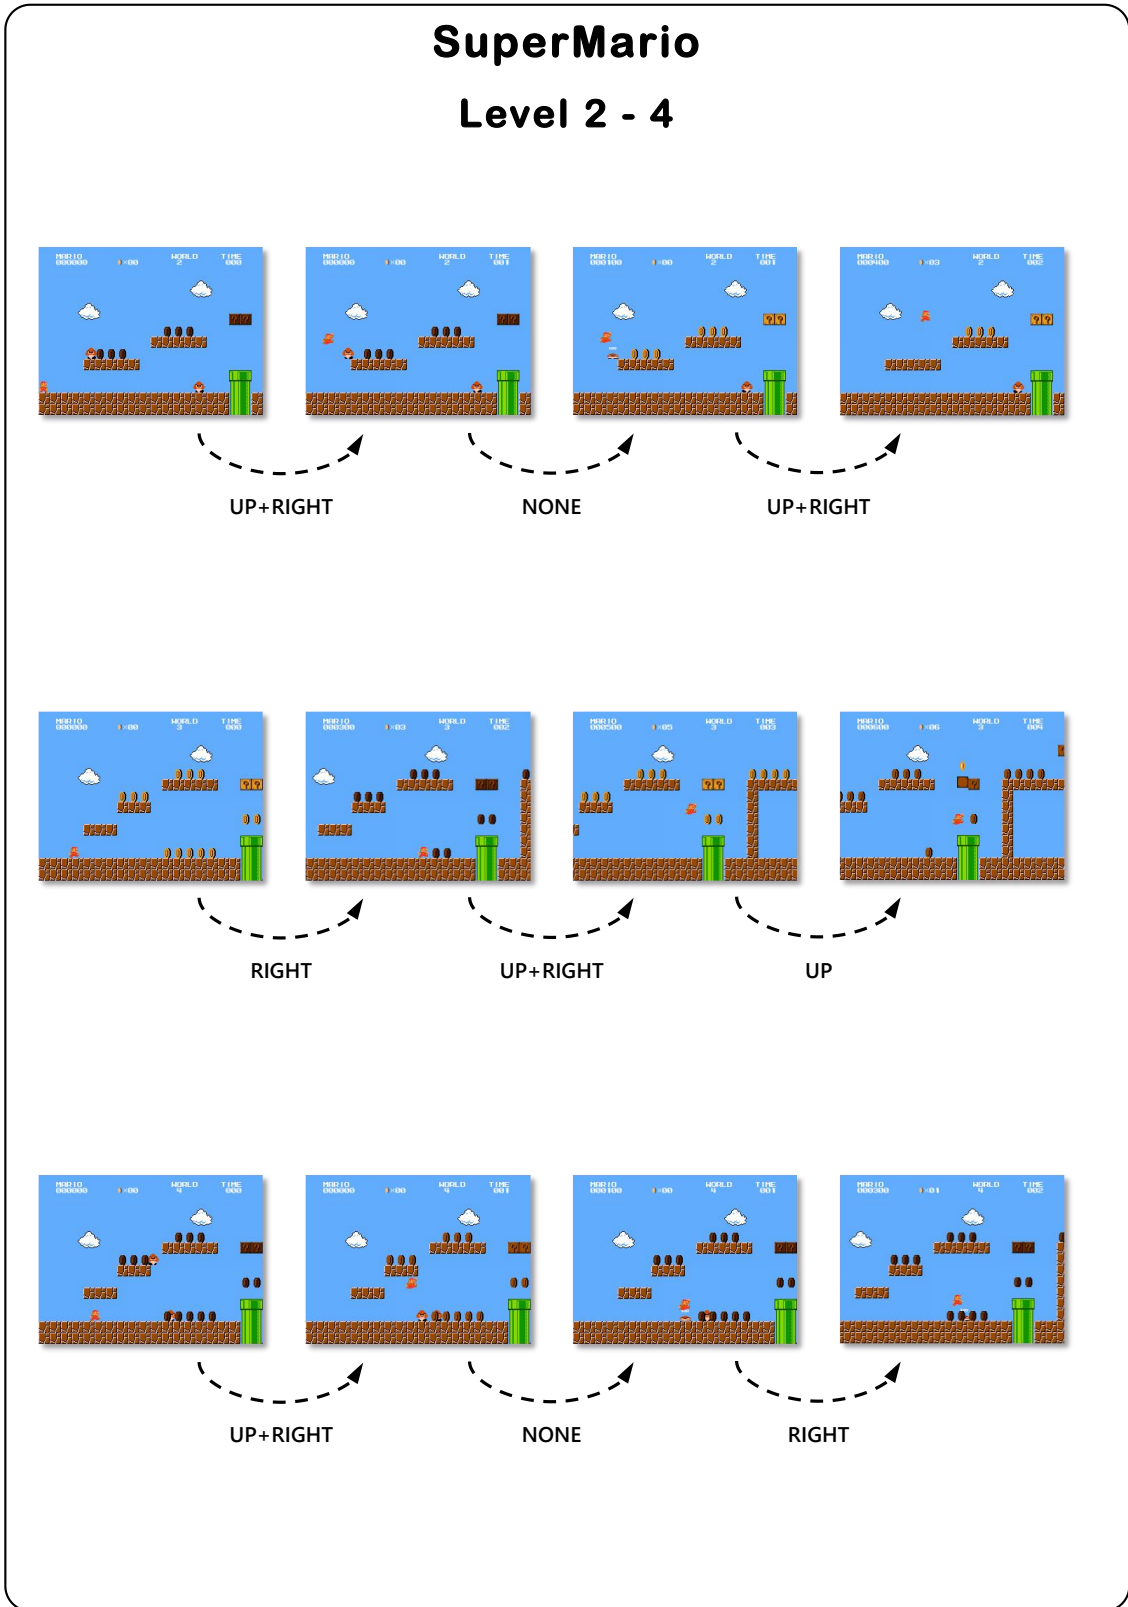


Figure 14: **SuperMario Level 2-4: Level Design and Prompt Overview.** The images showcase the scene from Level 2-4, illustrating the level design and corresponding prompt. **The prompt is the same as in Level 4.**

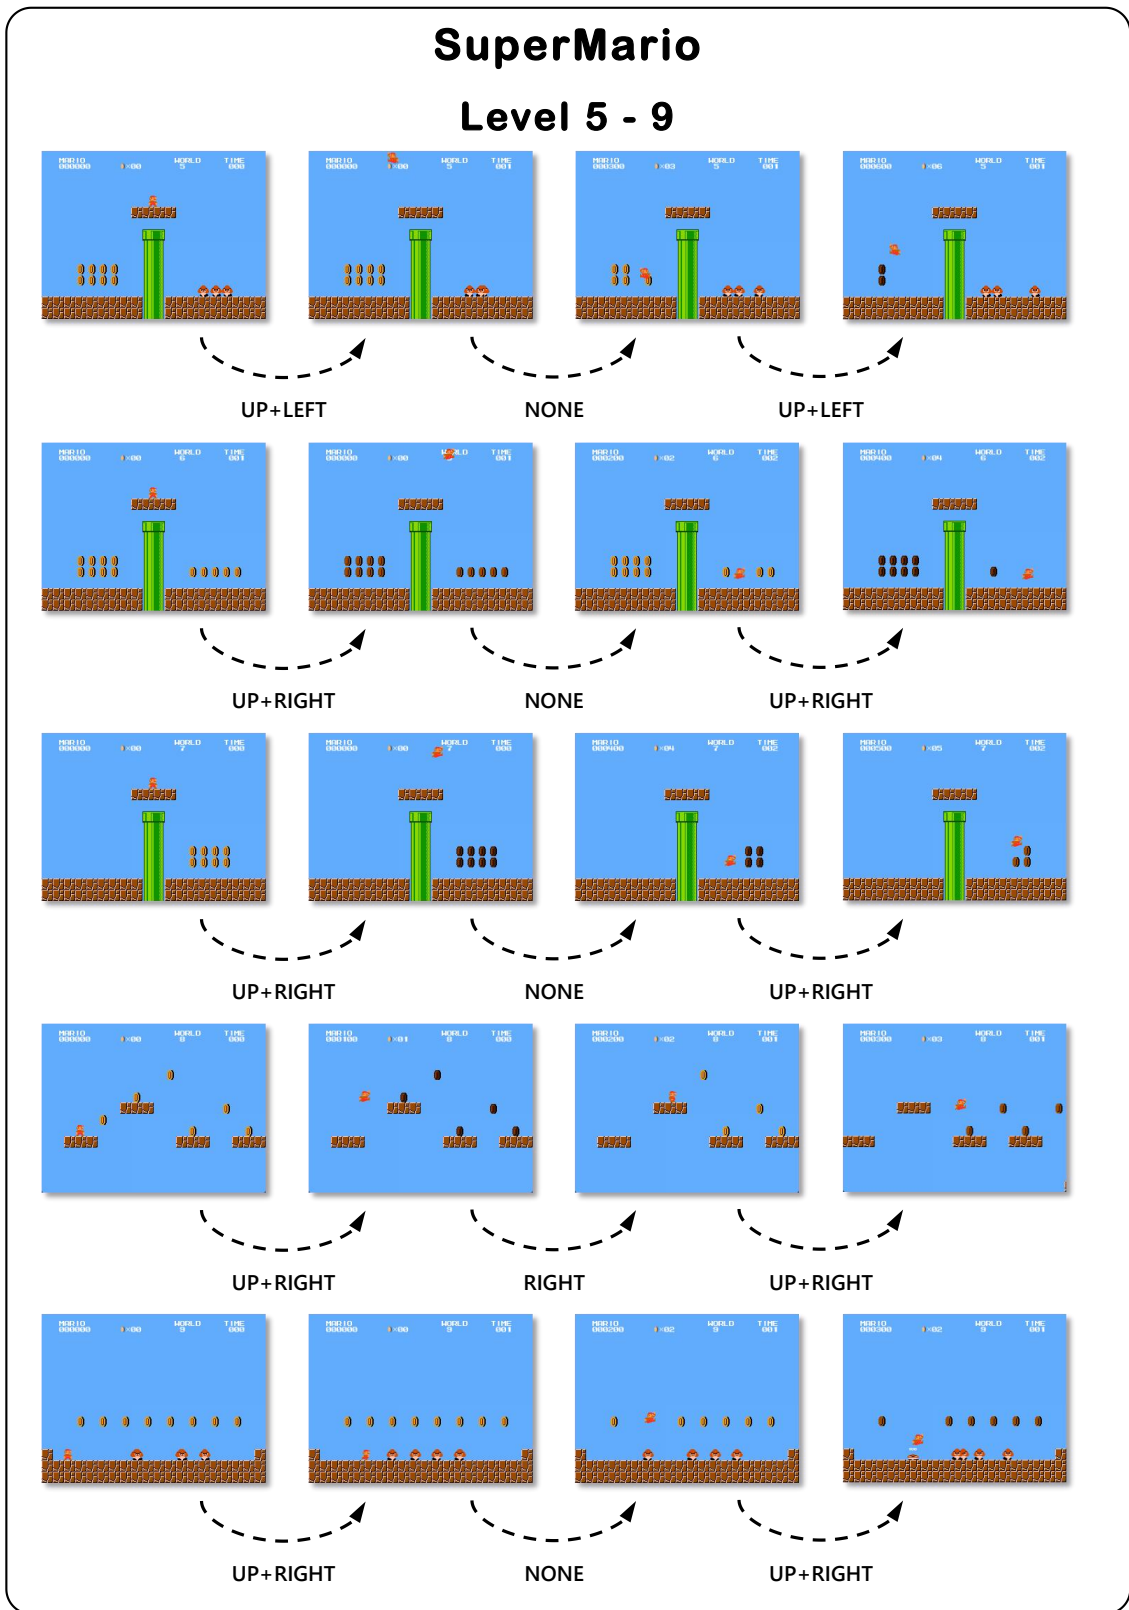


Figure 15: **SuperMario Level 5-9: Level Design and Prompt Overview.** The images showcase the scene from Level 5-9, illustrating the level design and corresponding prompt. **The prompt is the same as in Level 4.**

SuperMario

Level 10

This level is a 1:1 replica of the [World 1-1](https://www.mariowiki.com/World_1-1_(Super_Mario_Bros.)) scene from the original SuperMario game.

[https://www.mariowiki.com/World_1-1_\(Super_Mario_Bros.\)](https://www.mariowiki.com/World_1-1_(Super_Mario_Bros.))

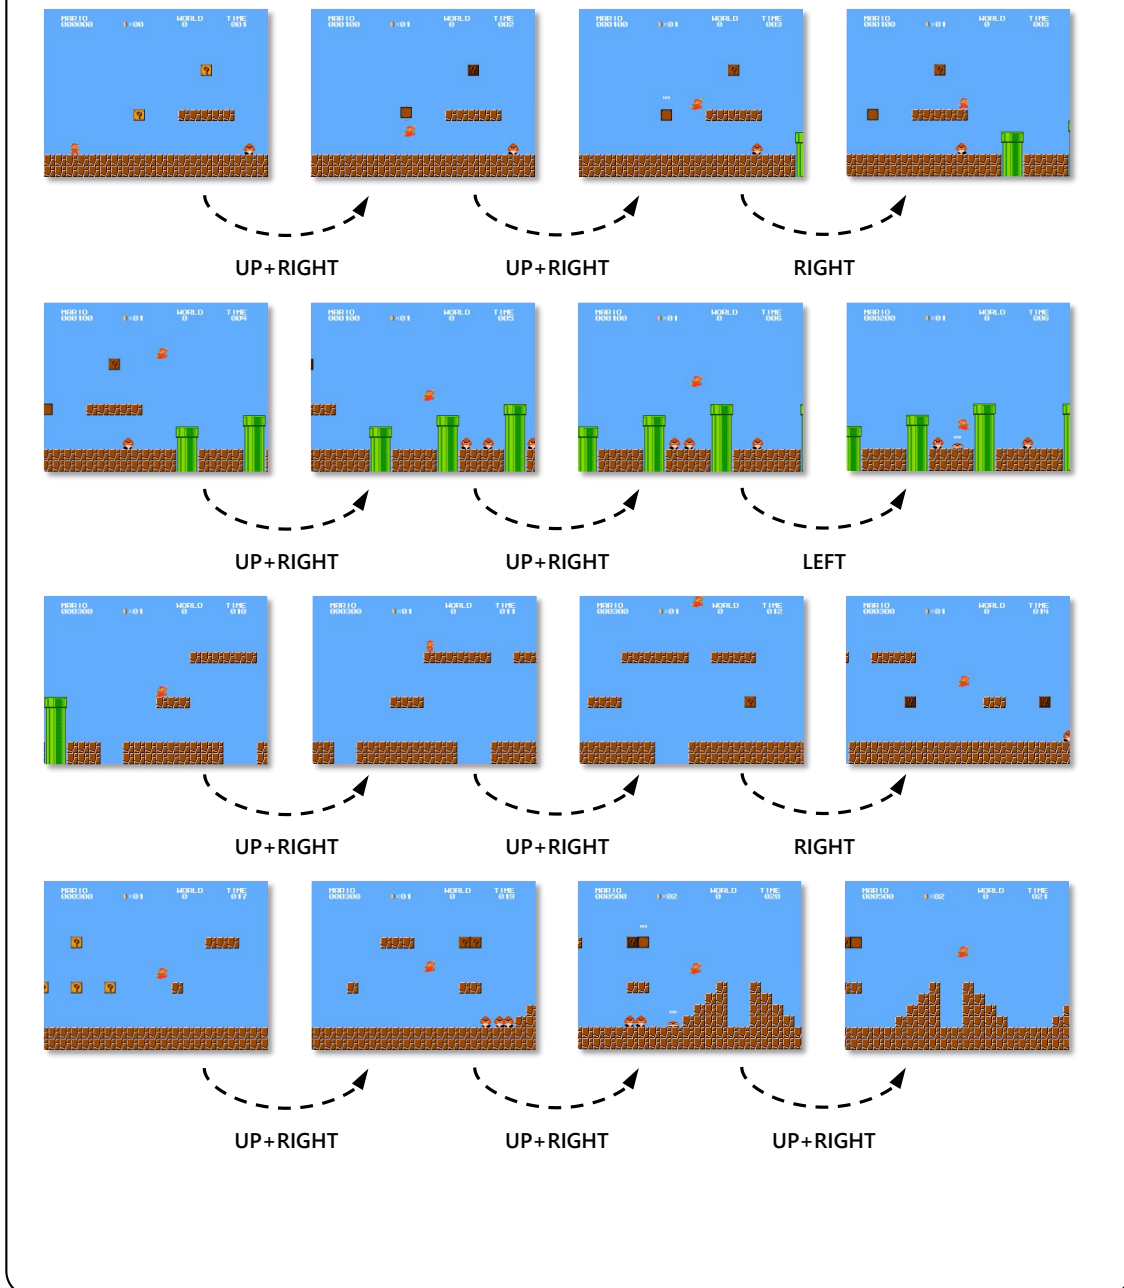
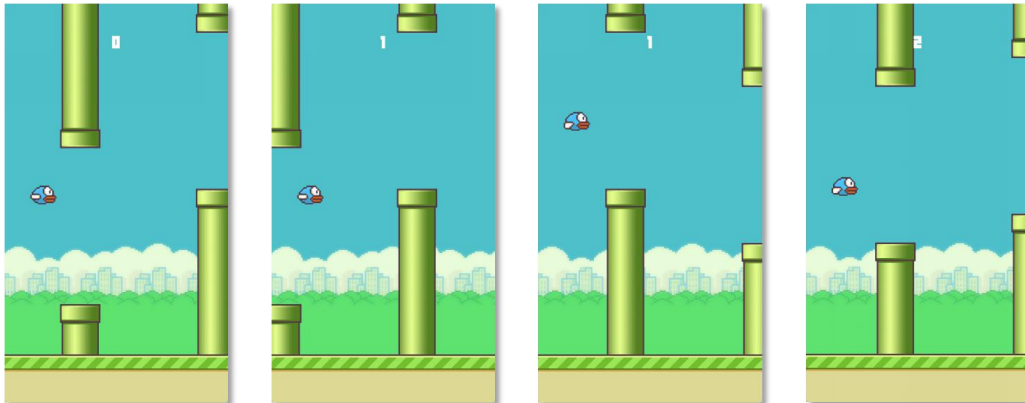


Figure 16: **SuperMario Level 10 (Standard Level): Level Design and Prompt Overview.** The images showcase the scene from Level 10, illustrating the level design and corresponding prompt. This is The standard level that matches the difficulty of the human game. **The prompt is the same as in Level 4.**

FlappyBird

Level 1 - 3



Prompt:

Assume you are playing a PC game called 'Flappy Bird'.

You need to control a bird that moves towards the right of the screen automatically. Your goal is to navigate through gaps between pipes without hitting them. Avoid hitting pipes or screen boundaries. Pass through as many gaps as possible to maximize your score.

Monitor the bird's position, the gap's height, and distance to the pipe.

In this version of the game, to make the game easier, there is no gravity, and the bird will NOT fall automatically. You can make the bird rise or fall by using the UP or DOWN action. Besides, you can also choose to keep the bird's current position without making it rise or fall.

In other words, you only need to determine the height of the bird can pass through the pipe at this time, if not, then through the UP or DOWN to adjust, when the height is appropriate, you can choose KEEP to let the bird through the pipe!

<\$history\$>

<\$image_current_step\$>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

You can make two types of actions to control the bird:

1. UP: Makes the bird rise a bit of distance.
2. DOWN: Makes the bird fall a bit of distance.
3. KEEP: The bird will keep the current position.

You should think step by step and response with the following format, remember to response the plain text without any special characters or symbols, DO NOT response in markdown or Latex format.

Observation: ... (describe the current position of the bird and the gap.)

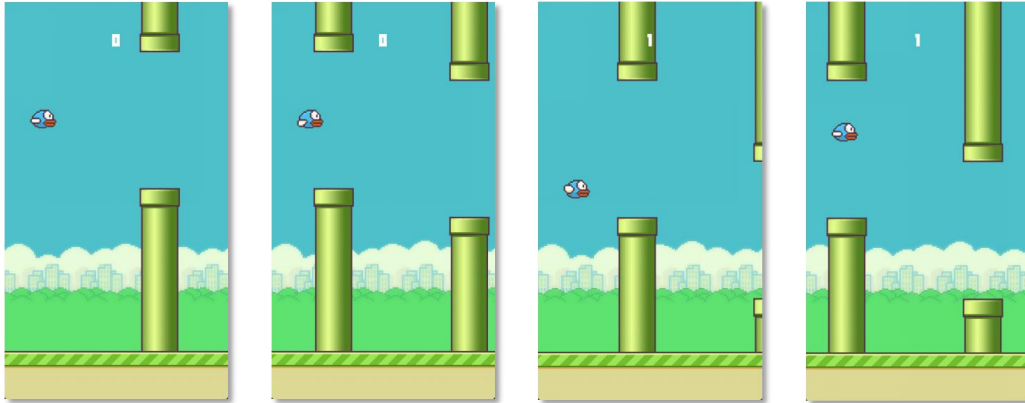
Reasoning: ... (think step by step and explain how you choose the action.)

Action: ... (choose one of the three actions (UP or NONE or KEEP) to control the bird.)

Figure 17: **FlappyBird Level 1-3: Level Design and Prompt Overview.** The images showcase the scene from Level 1, illustrating the level design and corresponding prompt. Levels are differentiated by the pipe gap width and the bird's forward speed. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty.

FlappyBird

Level 4 - 6



KEEP

NONE

UP

Prompt:

Assume you are playing a PC game called 'Flappy Bird'.

You need to control a bird that moves towards the right of the screen automatically. Your goal is to navigate through gaps between pipes without hitting them. Avoid hitting pipes or screen boundaries. Pass through as many gaps as possible to maximize your score.

Flap only when necessary to align with the next gap.
Monitor the bird's position, the gap's height, and distance to the pipe.
Adjust timing to maintain smooth movement and avoid collisions.

In this version of the game, to make the game easier, you can also choose to keep the bird's current position without making it rise or fall. This additional action can help you maintain the bird's height when needed. You can try to use this action to navigate through the gaps more effectively.

<\$history\$>

<\$image_current_step\$>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

You can make two types of actions to control the bird:

1. UP: Makes the bird rise.
2. NONE: The bird will fall a bit due to gravity.
3. KEEP: The bird will keep the current position.

You should think step by step and respond with the following format, remember to respond the plain text without any special characters or symbols, DO NOT respond in markdown or Latex format.

Response:

Observation: ... (describe the current position of the bird and the gap.)

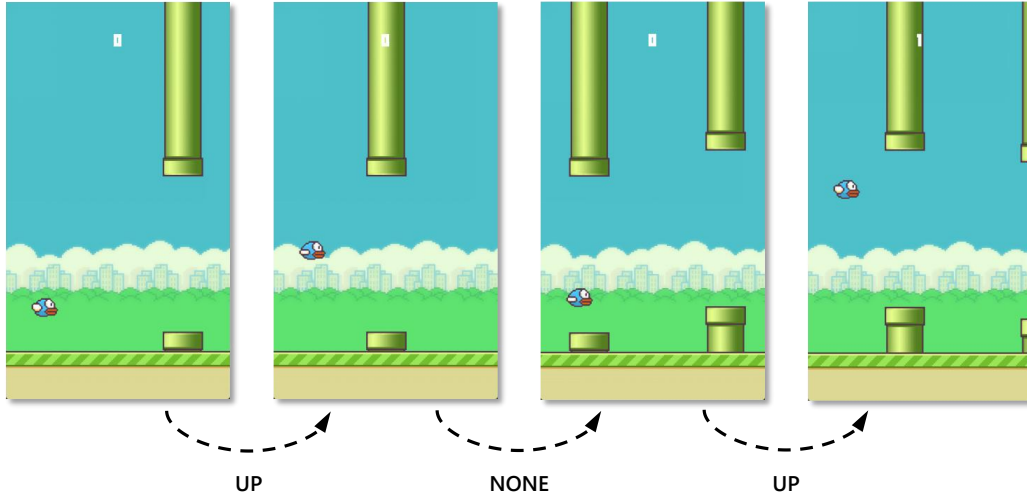
Reasoning: ... (think step by step and explain how you choose the action.)

Action: ... (choose one of the three actions (UP or NONE or KEEP) to control the bird.)

Figure 18: **FlappyBird Level 4-6: Level Design and Prompt Overview.** The images showcase the scene from Level 4, illustrating the level design and corresponding prompt. Levels are differentiated by the pipe gap width and the bird's forward speed. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty.

FlappyBird

Level 7



Prompt:

Assume you are playing a PC game called 'Flappy Bird'.

You need to control a bird that moves towards the right of the screen automatically. Your goal is to navigate through gaps between pipes without hitting them. Avoid hitting pipes or screen boundaries. Pass through as many gaps as possible to maximize your score.

Flap only when necessary to align with the next gap.

Monitor the bird's position, the gap's height, and distance to the pipe.

Adjust timing to maintain smooth movement and avoid collisions.

<\$history\$>

<\$image_current_step\$>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show the car's previous movements.

You can make two types of actions to control the bird:

1. UP: Makes the bird rise.
2. NONE: The bird may fall a bit due to gravity.

You should think step by step and respond with the following format, remember to respond the plain text without any special characters or symbols, DO NOT respond in markdown or Latex format.

Response:

Observation: ... (describe the current position of the bird and the gap.)

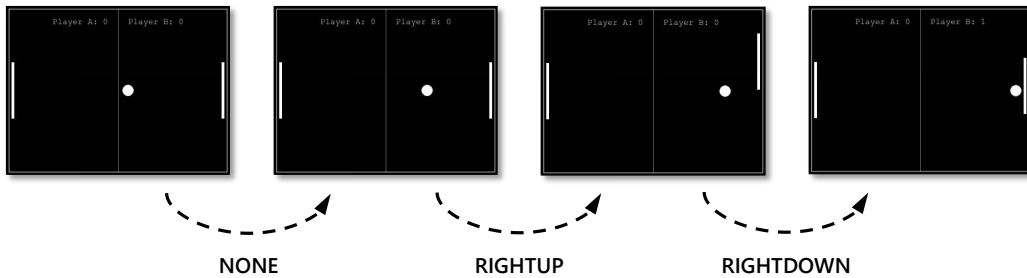
Reasoning: ... (think step by step and explain how you choose the action.)

Action: ... (choose one of the two actions (UP or NONE) to control the bird.)

Figure 19: **FlappyBird Level 7 (Standard Level): Level Design and Prompt Overview.** The images showcase the scene from Level 7, illustrating the level design and corresponding prompt. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty. This is The standard level that matches the difficulty of the human game.

PongGame

Level 1 - 3



Prompt:

Assume you are playing a PC game called 'Pong'.

You need to control two paddles on the left and right sides of the screen to hit a ball back and forth. Your goal is to prevent the ball from passing your paddle while trying to make it pass the opponent's paddle. Each time the ball is successfully hit by a paddle, that side gains one point. The final score is the sum of both sides' points.

Control the paddles' movements to align with the ball's trajectory and ensure successful hits. Adjust the paddles' positions smoothly to respond to the ball's movement and maintain control over the game.

<\$history\$>

<\$image_current_step\$>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show the paddles' and ball's previous movements.

You can make four types of actions to control the paddles:

1. LEFTUP: Moves the left paddle up.
2. LEFTDOWN: Moves the left paddle down.
3. RIGHTUP: Moves the right paddle up.
4. RIGHTDOWN: Moves the right paddle down.
5. NONE: No action.

You should think step by step and respond with the following format, remember to respond with plain text without any special characters or symbols, DO NOT respond in markdown or Latex format.

Observation: ... (describe the current positions of both paddles, the ball, and the ball's movement trajectory.)

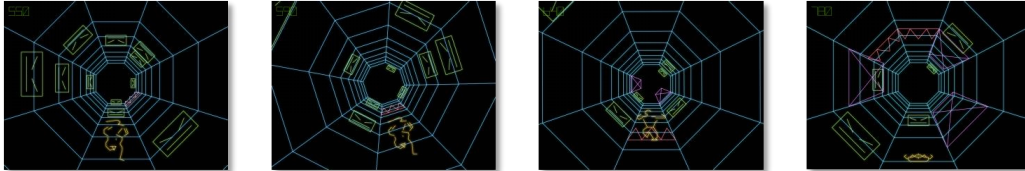
Reasoning: ... (think step by step and explain how you choose the action, including predictions about wall reflections and the ball's trajectory.)

Action: ... (choose one of the four actions: LEFTUP, LEFTDOWN, RIGHTUP, RIGHTDOWN, NONE. Do NOT add any other words.)

Figure 20: **PongGame Level 1-3: Level Design and Prompt Overview.** The images showcase the scene from Level 1, illustrating the level design and corresponding prompt. Levels are differentiated by the paddle width and the ping pong ball's speed. The ping pong ball in the same level will randomly change its initial position while maintaining consistent relative difficulty.

Tempest Run

Level 1 - 4



RIGHT

JUMP

SLIDE

Prompt:

Assume you are playing a PC game called 'Tempest Run'.

You need to control a character who moves through a three-dimensional space inside a futuristic tunnel filled with various obstacles and enemies. Your goal is to navigate through the tunnel, avoid or overcome obstacles, and run as far as possible. Avoid colliding with red spikes, purple walls, or failing to deal with green enemies.

Use the optimal combination of movements to progress through the tunnel smoothly and efficiently. Monitor the character's position relative to obstacles and react appropriately to avoid losing progress.

<\$history\$>

<\$image_current_step\$>

This screenshot represents the current step of the game.

The last frame shows the current state of the game, while the previous frames show the character's previous movements.

Important notes:

1. Use JUMP to jump over red spikes on the ground.
2. Use SLIDE to duck and kick green enemies to eliminate them.
3. Use LEFT or RIGHT to move around obstacles, such as purple walls or spikes.
4. Use RISE to return to a normal running position after a SLIDE.
5. NONE is a valid action to maintain the current state if no immediate action is needed.

You can make six types of actions to control the character:

1. JUMP: Makes the character jump upward, useful for avoiding ground obstacles like red spikes.
2. LEFT: Moves the character to the left.
3. RIGHT: Moves the character to the right.
4. SLIDE: Makes the character duck and slide forward, useful for dealing with green enemies or passing under certain obstacles.
5. RISE: Returns the character to a normal running position after sliding.
6. NONE: No new action is performed; the character maintains their current trajectory.

You should think step by step and respond with the following format, remember to respond with plain text without any special characters or symbols, DO NOT respond in markdown or Latex or any other format.

Response:

Observation: ... (Describe the character's current position and nearby obstacles or enemies.)

Reasoning: ... (Think step by step and explain how you choose the action.)

Action: ... (Choose ONE of the six actions to control the character. Do NOT add any other words.)

Figure 21: **Tempest Run Level 1-4: Level Design and Prompt Overview.** The images showcase the scene from Level 1, illustrating the level design and corresponding prompt. Levels are differentiated by the speed of barrier generation and the amount of visual information available. Elements in the same level will randomly change their initial positions while maintaining consistent relative difficulty.

Table 25: Tempest Run Level Configurations

Level	Role Speed	Cell Length	Random Rate	Others
1	Slow	Large	Low	Environmental elements initial positions randomly change.
2	Medium	Medium	Medium-Low	
3	Fast	Small	Medium-High	
4	Very Fast	Small	High	

Table 26: Game Codebase Sources

Game	Codebase
Race	https://github.com/tdostilio/Race_Game
FlappyBird	https://github.com/agneay/pygame-projects/tree/master/Flappy%20Bird
Pong	https://github.com/pyGuru123/Python-Games/tree/master/Pong
SuperMario	https://github.com/mx0c/super-mario-python
Tempest Run	https://github.com/davidpendergast/pygame-summer-team-jam

C Elo Performance Comparison Protocol Details

The core methodology for metrics evaluation in V-MAGE comprises two interconnected components: *performance comparison* and *statistical stabilization*.

Performance Comparison Protocol.

Each model begins with an initial Elo rating $R_m^{(0)} = 1500$, where $m \in \mathcal{M}$ represents the set of candidate models. We conducted 100 evaluation rounds for each game level ℓ where models were randomly paired in each round through a shuffle operation:

$$\mathcal{P}_t = \{(A_t, B_t) \mid A_t, B_t \stackrel{\text{rand}}{\sim} \mathcal{M}^\ell, A_t \neq B_t\} \quad (1)$$

where A_t and B_t denote the paired models in round t .

For paired models (A, B) , their game scores score_A and score_B are compared next. We first construct

$$f(m) = (\text{score}_m, \text{valid_rate}_m) \quad (2)$$

where valid_rate_m represents the proportion of actions output by model m in this game round that conform to the formatting requirements. The outcome $(S_A, S_B) \in \{(1, 0), (0, 1), (0.5, 0.5)\}$ is determined by:

$$(S_A, S_B) = \begin{cases} (1, 0) & \text{if } f(A) \succ f(B) \\ (0, 1) & \text{if } f(A) \prec f(B) \\ (0.5, 0.5) & \text{otherwise} \end{cases} \quad (3)$$

The rating update mechanism follows the classical Elo formulation with enhanced stability controls. For models A and B with pre-match ratings R_A and R_B , their expected performance is calculated as:

$$E_A = \frac{1}{1 + 10^{(R_B - R_A)/400}}, \quad E_B = \frac{1}{1 + 10^{(R_A - R_B)/400}} \quad (4)$$

where the denominator base 10 and scaling factor 400 establish a logarithmic relationship between rating differences and expected outcomes. The post-match ratings become:

$$\begin{aligned} R'_A &= R_A + K(S_A - E_A) \\ R'_B &= R_B + K(S_B - E_B) \end{aligned} \quad (5)$$

where K is a constant determining the sensitivity of the rating system (typically set to 32),

Stabilization through Randomized Iteration.

To ensure the robustness of rating updates, V-MAGE implements multi-pass stabilization protocol. All historical comparisons are aggregated into a win/loss pool:

$$\mathcal{W} = \bigcup_{g \in \mathcal{G}} \bigcup_{\ell \in \mathcal{L}_g} \bigcup_{t=1}^N (A_{glt}, B_{glt}, S_A^{glt}, S_B^{glt}) \quad (6)$$

which undergoes $T = 10^4$ independent shuffles. For each permutation $\pi_i(\mathcal{W})$, complete rating recalculation yields $R_m^{(i)}$. The final stabilized rating combines these trials:

$$\bar{R}_m = \frac{1}{T} \sum_{i=1}^T R_m^{(i)} \quad (7)$$

D Ablation Study On Pipeline Settings

D.1 Impact of History Sampling Configuration

We conducted supplementary experiments on the **Qwen2.5-VL-7B** and **Qwen2.5-VL-72B** models using various history strategies (including increasing the number of history steps and altering sampling methods). The results are presented in Tables 27 and 28. These scores were calculated as a

percentage of model scores versus human performance in a manner similar to Figure 7.

In the default setting of our main experiments, the history sampling configuration is one where decisions are made using information from the most recent **3steps**, and the game screen is sampled every **3frames**.

The experimental results show that simply increasing the length of the history window (e.g., from 3 to 8 steps) does not yield significant performance gains. This finding supports our core argument: the bottleneck for current MLLMs lies **not in the quantity** of historical information they receive, but in their ability to **understand and utilize** this dynamic visual information.

Therefore, we chose a 3-frame history as our baseline configuration. This provides the necessary temporal context while establishing a fair, simple, and effective standard for exposing the models' core deficiencies, without confounding the evaluation with complex agent strategies.

As mentioned in the main text, to investigate the impact of settings within the sampling strategies on anchoring bias, we also conducted relevant experiments, with the results presented in Appendix E.4.2.

D.2 Impact of Input Resolution

To systematically investigate the impact of input resolution on model performance, we conducted a new set of experiments, testing the **Qwen2.5-VL 7B** and **72B** models on four different resolutions.

The resolutions from 120 to 480 refer to images with heights of 120 to 480 pixels, respectively, with the width scaled according to the original aspect ratio. We selected these four resolutions to cover different levels of visual detail, from low to high. In the **default** setting of our main experiments, the model's input resolution was **360** pixels height. The scores were calculated as a percentage of model scores versus human performance in a manner similar to Figure 4. The results are presented in Table 29.

This data reveals a nuanced relationship: for the more capable 72B model, the overall performance trend improves with higher resolution, peaking at 480px. This suggests it can benefit from the finer details in higher-resolution images. However, for the smaller 7B model, performance peaks at our default setting of 360px and declines at the higher 480px resolution.

This indicates that the relationship between

model performance and input resolution is **not simply linear**. For less capable models, excessive resolution might introduce 'noise' that they struggle to filter effectively, thereby interfering with their decision-making process.

Table 27: Performance comparison of different history strategies for Qwen2.5-VL 7B and 72B models.

Game	Qwen2.5-VL-7B					Qwen2.5-VL-72B				
	3steps_2sample	3steps_5sample	3steps	5steps	8steps	3steps_2sample	3steps_5sample	3steps	5steps	8steps
race	11.20	11.20	12.60	11.20	12.40	30.00	29.00	29.60	32.60	33.60
supermario	20.10	22.10	22.60	22.80	21.20	34.50	33.90	42.10	36.40	39.80
pong	3.30	4.30	3.70	4.00	4.30	4.50	3.90	4.10	4.80	5.00
flappybird	6.70	11.20	3.40	2.10	5.00	17.70	13.10	8.10	13.30	13.60
tempestrun	18.80	17.80	21.10	18.80	17.80	22.00	21.10	24.80	22.70	23.70
average	12.02	13.32	12.68	11.78	12.14	21.74	20.20	21.74	21.96	23.14

Table 28: Performance comparison of different frame sampling strategies for Qwen2.5-VL 7B and 72B models.

Game	Qwen2.5-VL-7B				Qwen2.5-VL-72B			
	8frames	5frames	3frames	1frames	8frames	5frames	3frames	1frames
race	13.80	11.60	12.60	10.40	19.60	27.20	29.60	26.00
pong	4.00	3.90	3.70	4.50	4.80	5.90	4.10	7.60
flappybird	3.40	5.70	3.40	7.60	10.80	14.00	8.10	14.10
tempestrun	19.00	19.30	21.10	16.50	23.40	18.50	24.80	26.60
average	10.05	10.13	10.20	9.75	14.65	16.40	16.65	18.58

Table 29: Performance comparison of Qwen2.5-VL 7B and 72B models across different input resolutions (height in pixels). Scores are percentages relative to human performance.

Game	Qwen2.5-VL 7B				Qwen2.5-VL 72B			
	120 (7B)	240 (7B)	360 (7B)	480 (7B)	120 (72B)	240 (72B)	360 (72B)	480 (72B)
race	9.80	10.60	12.60	11.00	15.80	23.60	29.60	28.20
supermario	17.50	21.70	22.60	17.20	38.90	44.90	42.10	47.90
pong	4.00	3.90	3.70	3.60	3.50	3.50	4.10	3.60
flappybird	4.40	5.40	3.40	9.80	7.70	12.70	8.10	12.00
tempestrun	19.60	19.10	21.10	18.50	19.60	24.30	24.80	22.50
average	<i>11.06</i>	<i>12.14</i>	12.68	<i>12.02</i>	<i>17.10</i>	<i>21.80</i>	<i>21.74</i>	22.84

E Additional Experimental Details

E.1 Unit Tests For Core Visual Abilities Experiment

The unit testing framework conducts a systematic assessment of fundamental visual capabilities by drawing from the comprehensive V-MAGE bench-

mark. In each carefully designed level of a game, a random baseline score is first determined by averaging scores from random actions. Following this, the performance of each evaluated model on the said level is quantified by calculating the percentage of rounds where the model’s score outperforms this established random baseline. The specific game

Table 30: Basic visual capabilities and their corresponding simple game levels.

Visual Abilities	Game	Levels
Tracking	Pong	1, 2, 3
Positioning	Race	1, 1_no_history
Visual Grounding	TempestRun	1
Timing	FlappyBird	1, 2, 3

Table 31: Performance of MLLMs on Core Visual Ability Unit Tests (% Exceeding Random Baseline)

Model	Positioning	Tracking	Visual Grounding	Timing
Qwen2VL 7B	0.50	0.27	0.56	0.36
Qwen2VL 72B	0.76	0.26	0.70	0.43
Qwen2.5VL 72B	0.88	0.25	0.68	0.51
InternVL2.5 78B	0.82	0.33	0.66	0.49
InternVL2.5 8B	0.60	0.28	0.55	0.39
Gemini-2.0-Flash	0.68	0.32	0.70	0.51
GPT4o	0.98	0.29	0.66	0.58

levels used for assessing each ability are listed in Table 30.

As illustrated in Figure 4a, model performances across representative levels for four fundamental visual competencies reveal critical insights: In tracking tasks requiring cross-frame analysis, nearly all models underperform random baselines. This indicates that while current models achieve reasoning through caption-based approaches in single-frame tasks, they struggle to extract discriminative features in multi-frame scenarios requiring fine-grained spatiotemporal comparisons. The quantitative results for each model across the four core visual abilities are presented in Table 31.

It is important to interpret the results of these unit tests within their intended scope. Designed to assess fundamental visual competencies, these tests utilize a random baseline score as the primary reference point. While a model significantly outperforming this random baseline indicates a degree of relevant reasoning ability in that specific task dimension, it does not necessarily imply a high level of overall competence. The random baseline represents minimal performance, and even achieving scores far exceeding it on these foundational tests serves primarily to diagnose basic capabilities rather than validate advanced mastery required for complex gameplay.

E.2 Perceptual Skipping Experiment

To further investigate the interplay between visual perception and reasoning, we conducted supplementary experiments where textual descriptions of the game state were provided, effectively bypassing the visual perception module (see Table 32 for detailed results on Flappy Bird Level 3 and Pong Level 2).

The results indicate that alleviating the perceptual challenge generally improves performance, particularly for larger models like GPT-4o and the 72B/78B parameter models, supporting the hypothesis that visual perception is a significant bottleneck. However, even with this intervention, model scores remained substantially lower than the human baseline (>10), underscoring the presence of critical reasoning and planning deficiencies beyond visual perception, as discussed earlier.

Notably, the performance gains from text input were more pronounced for larger models, suggesting their enhanced capacity to leverage structured textual information for reasoning, whereas smaller models exhibited less consistent benefits or even performance degradation in some cases. This finding further highlights that while perception is a challenge, fundamental reasoning limitations persist across models and are not fully overcome even when provided with simplified, textual state representations.

Table 32: Model performance on simple levels with and without textual state information.

Model	Flappy Bird		Pong	
	w/o Text	w/ Text	w/o Text	w/ Text
Qwen2.5VL 7B	0.8	0.35	0.19	0.25
InternVL2.5 8B	0.31	0.76	0.19	0.31
Qwen2.5VL 72B	0.35	2.17	0.21	1.19
InternVL2.5 78B	0.59	2.39	0.16	0.52
GPT4o	0.57	4.55	0.20	3.89
Gemini-2.0-Flash	0.42	4.89	0.32	>10
random	0.52		0.18	
human	> 10		> 10	

Table 33: Action Efficiency on Race Levels 1–3. Lower Action Efficiency is better.

Model	Success Rate (L1 / L2 / L3)	Action Efficiency (L1 / L2 / L3)
Qwen3-VL-235B-A22B	85% / 37% / 40%	4.5 / 8.9 / 8.4
Qwen2.5-VL-72B	100% / 27% / 27%	4.6 / 7.9 / 8.1
Qwen2.5-VL-7B	55% / 9% / 7%	6.0 / 13.1 / 6.3
Qwen2-VL-7B	7% / 0% / 2%	3.0* / - / 5.0*

*Low success rates introduce survivorship bias: only the easiest successful seeds are counted.

E.3 Action Efficiency as a Fine-Grained Diagnostic

To better differentiate models that achieve similar success rates on easy or medium-difficulty levels, we introduce **Action Efficiency**, defined as the average number of decision steps required to reach a successful outcome (lower is better). We evaluate four Qwen-family models on Race Levels 1–3, where score ties are common.

We also find that incorporating Action Efficiency into ELO updates preserves the ranking order while increasing separation between efficient and inefficient policies. This supports using Action Efficiency as a secondary diagnostic when raw success rates alone are insufficiently discriminative.

E.4 Anchoring Bias Experiments

E.4.1 Details and Examples

Due to the possibility of models receiving identical visual inputs over multiple rounds in Super Mario (e.g., being stuck in a corner), we conducted a statistical analysis using lots of rounds of responses from each model across the other four games. This was done by iterating through the recorded interactions for each level and measuring the number of sequential rounds where the model’s output re-

sponse remained unchanged. The average of these durations across interactions within a game provides the metric presented in Table 34.

The results indicate that GPT-4o updates its responses more actively and frequently when the visuals change, while other models do so less frequently. This may suggest that GPT-4o is more sensitive to subtle visual updates, enabling it to make timely inferences and more accurately track game progress.

In the example of Figure 22, GPT-4o’s response during the final interaction is as follows:

Table 34: Average number of rounds for each model to generate different responses.

Model	Race	FlappyBird	Pong	TempestRun	Avg.
Qwen2VL 7B	4.3	25.9	13.7	7.3	12.8
Qwen2.5VL 72B	2.3	19.3	2.6	5.3	7.4
InternVL2.5 8B	2.0	6.9	6.7	8.0	5.9
InternVL2.5 78B	6.8	16.0	2.0	3.0	7.0
GPT4o	1.0	1.6	1.0	1.0	1.1
PCC r (Avg. Rounds vs. ELO)	-0.63	-0.86	-0.88	-0.64	-0.75



Figure 22: In the prior reasoning processes, GPT-4o accurately identified the car’s orientation relative to the trophy’s position. However, in the final reasoning instance, due to anchoring bias, the model misinterpreted previous historical information and incorrectly assumed that the car’s orientation was already directed toward the trophy. Consequently, despite correctly determining the direction, the model executed an erroneous action.

Description: The car is currently facing diagonally towards the **bottom-right** of the screen. The trophy is positioned near the car, slightly to the **right**. There are no visible obstacles between the car and the trophy, indicating a clear path. The car has been moving consistently forward in the previous steps, maintaining its alignment towards the trophy.

Reasoning: Given the car’s orientation towards the trophy and the absence of obstacles, the car’s current position and direction remain favorable for reaching the trophy. Continuing to move forward is the most efficient strategy to maintain the direct path to the trophy. The car’s current trajectory should allow it to reach the trophy without any adjustments.

Action: FORWARD

E.4.2 Impact of Frame Sampling and Decision Frequency

To further quantify the relationship between perceptual sensitivity and anchoring bias, we have conducted a deeper quantitative exploration of the relationship between anchoring bias and model performance. To more objectively measure a model’s reaction to dynamic changes in the game world, we introduced a new metric: **Average Response Game Frames** (abbreviated as ‘**avg**

frames’ in the results). This metric is calculated by: Avg. Response Frames = sampling interval * avg. rounds to generate a different response, and represents **how many game frames, on average, have elapsed before a model makes a substantive change in its reasoning**.

We performed a series of experiments with different sampling strategies, first testing the Qwen2.5VL-72B model. The results are shown in Table 35:

Stable ‘visual reaction threshold’ in strong models: From the avg frames metric, the Qwen2.5VL-72B model demonstrates remarkable consistency across different sampling strategies, with its average response time stabilizing within a narrow range of **20-30** game frames. This suggests that the model possesses a relatively constant intrinsic reaction threshold, where a certain amount of accumulated visual change triggers a shift in its reasoning.

Regarding task score (avg score): The 72B model’s performance clearly improves as the sampling interval decreases, with the highest score achieved at the highest decision frequency (1-frame interval). Under such high-frequency decision-making, the model can capture crucial task timings with the highest precision. As the decision frequency decreases, the opportunities for the model

Table 35: Performance of Qwen2.5VL-72B under different frame sampling strategies. The top section shows response frequency metrics, while the bottom shows game scores.

Game	8frames	5frames	3frames(default)	1frames
<i>Average Response Game Frames</i>				
race	5	1.8	2.3	10.8
pong	1.7	1.7	2.6	19.7
flappybird	1.5	2.3	19.3	64.2
tempestrun	1.8	14.2	5.3	23.2
avg request	2.5	5	7.4	29.5
avg frames	20	25	22.2	29.5
<i>Game Score</i>				
race	19.60	27.20	29.60	26.00
pong	4.80	5.90	4.10	7.60
flappybird	10.80	14.00	8.10	11.90
tempestrun	23.40	18.50	24.80	26.60
avg score	14.65 ↓2.00	16.40 ↓0.25	16.65 ↑0.00	18.03 ↑1.38

to take appropriate action at the right moment are reduced, thus may lead to a drop in performance.

Next, we compared the **Qwen2.5VL-7B** and **72B** models under the same sampling strategies. The results are presented in Table 36.

The correlation between the "visual reaction threshold" and "task score" remains clear when comparing across models. The 72B model exhibits a **lower** Average Response Game Frames (indicating higher perceptual sensitivity) and a **higher** task score, while the 7B model shows the opposite. This is consistent with the conclusions about anchoring bias: a more powerful model possesses greater sensitivity to dynamic visual perception, which forms the basis for more accurate decision-making in interactive tasks.

E.5 Analysis of Model Errors in V-MAGE

We have collected **2,351** prompt-response pairs generated by GPT-4o while completing all levels for 1 to 5 rounds. From these, **494** examples were randomly and uniformly sampled for manual error annotation. The frequency of occurrence for various error types is shown in Tables 37 and 38, corresponding to GPT-4o and Qwen3-VL-235B-A22B-Instruct, respectively.

In addition, complementing the analysis of GPT-4o, we also investigated the best-performing open-source model under our evaluation framework, namely Qwen3-VL-235B-A22B-Instruct. The results are detailed as follows.

The definitions of each error type are presented as follows:

- **no error**: There is no error in the response.
- **perception error**: Description misinterpreted elements.
- **direction error**: A type of perception error. Confused directions (e.g., LEFT/RIGHT)
- **recognition error**: A type of perception error. Failed to identify key objects/elements.
- **perception incomplete**: Description missed important elements in the scene.
- **reasoning error**: Flawed logic in the reasoning section for the chosen action.
- **history misinterpretation**: A type of reasoning error. Misunderstood the game history.
- **action inappropriate**: A type of reasoning error. The chosen Action is clearly wrong given the Observation/Reasoning.
- **inconsistency**: Inconsistent action plans in multiple response processes.
- **instruction following**: Failed to follow instructions in the prompt.

Table 36: Comparison of Qwen2.5VL-7B and 72B models across sampling strategies. The 72B model shows a consistently lower reaction threshold (avg frames) and higher scores.

Game	8frames	5frames	3frames	1frames
Qwen2.5VL-7B				
avg request	13.5	12.5	34.7	97.8
avg frames	108.0	62.5	104.1	97.8
avg score	10.1	10.1	10.2	9.8
Qwen2.5VL-72B				
avg request	2.5	5	7.4	29.5
avg frames	20.0 ↓88.0	25.0 ↓37.5	22.2 ↓81.9	29.5 ↓68.3
avg score	14.7 ↑4.6	16.4 ↑6.3	16.7 ↑6.5	18.6 ↑8.8

Table 37: Error count by error type and game environment in GPT4o

Error Type	FlappyBird	Pong	Race	SuperMario	TempestRun
no error	30	18	54	88	21
perception error	80	26	26	47	42
direction error	2	19	13	16	8
recognition error	1	0	0	0	5
perception incomplete	3	0	8	10	10
reasoning error	24	4	9	10	6
history misinterpretation	21	0	1	6	2
action inappropriate	0	0	5	0	0
inconsistency	0	0	14	1	0
instruction following	0	0	2	0	0

Table 38: Error count by error type and game environment in Qwen3-VL-235B-A22B-Instruct

Error Type	FlappyBird	Pong	Race	SuperMario	TempestRun
no error	44	20	52	90	22
perception error	35	15	14	25	77
direction error	3	18	12	17	9
recognition error	0	0	0	0	0
perception incomplete	10	0	6	8	12
reasoning error	3	5	10	0	5
history misinterpretation	41	1	11	15	3
action inappropriate	0	0	6	0	0
inconsistency	10	5	15	0	0
instruction following	0	0	0	0	0

F Miscellaneous Material

F.1 LLM Usage Statement

Our research methodology centered on the evaluation of various MLLMs. Models such as GPT-4o served as the subjects within our V-MAGE pipeline, generating the outputs that form the basis of our analysis and conclusions on MLLM performance. The role of these MLLMs was strictly limited to this evaluation phase. The conceptualization and implementation of the V-MAGE framework and its software were carried out entirely by the authors.

For the manuscript preparation, we employed LLMs for the sole purpose of improving grammar and polishing the language. All scientific contributions, including the research ideas, experimental design, and results interpretation, originate exclusively from the authors.

F.2 Crowdsourcing and Research with Human Subjects

Research involving human subjects in this paper was limited to inviting a small number of participants to perform tasks within the V-MAGE game environments for the purpose of establishing a human performance baseline. Participants were provided with the standard game rules and objectives, identical to the instructions given to the evaluated models (see Appendix B.4 for details on setup and prompts). Participation was entirely voluntary, and no compensation was provided. No sensitive or personally identifiable information was collected from participants.

F.3 Institutional Review Board (IRB) Approvals

The research involving human subjects in this project has always been conducted under the guidance and supervision of our institution’s Institutional Review Board (IRB) and in full compliance with its policies. To formally document this compliance for publication, our research protocol was reviewed by the IRB committee. The committee confirmed the study’s classification as ‘minimal risk’ and has approved our research protocol.

Furthermore, to address the concern about reproducibility from a practical standpoint, the game tasks we selected have clear objectives and straightforward rules. For human players who understand these rules, their scores are highly stable and approach the theoretical performance ceiling. We are therefore confident that the human baseline

can be easily reproduced and verified by other researchers.

G Case Study

28.

We present additional case studies and error analyses, as shown in Figures 23, 24, 25, 26, 27, and

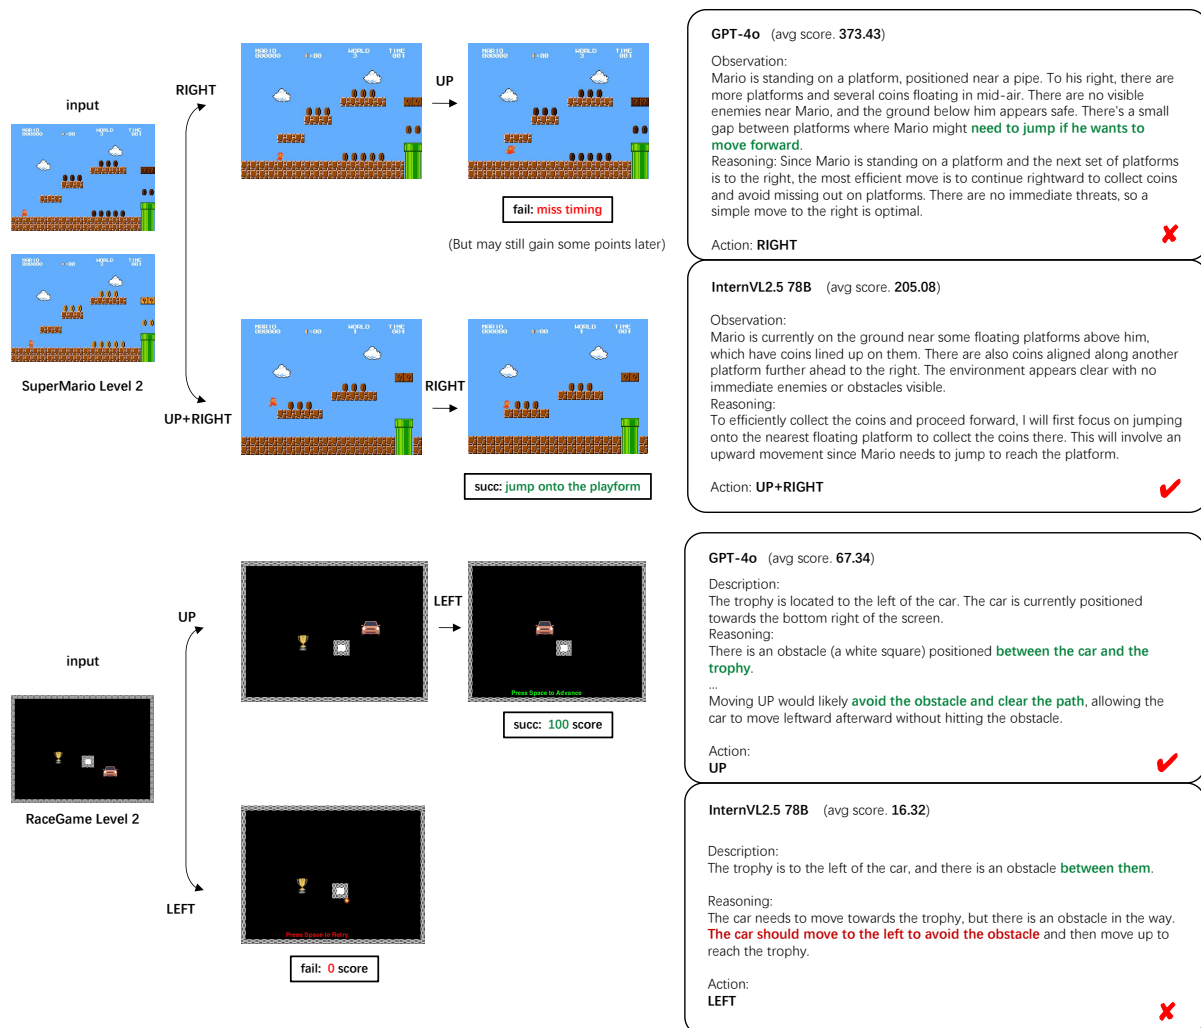
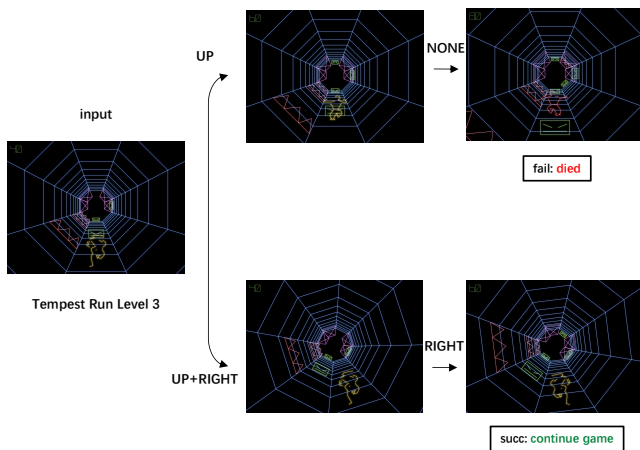


Figure 23: case studies group 1.



GPT-4o

Observation: The character is currently running in the tunnel. There is a red spike directly ahead on the ground, a purple wall to the left, and a green enemy on the path in front.

Reasoning: The red spike must be avoided by jumping over it. The green enemy can be dealt with by using the SLIDE action to kick it. The purple wall is on the left, so it is safe to move right.

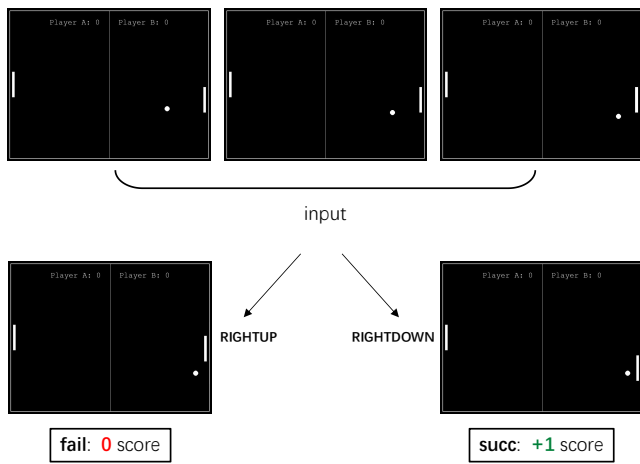
Action: **JUMP** ❌

Qwen2VL 72B

Observation: The character is currently positioned in the center of the tunnel. There are red spikes on the ground to the left and a green enemy directly ahead.

Reasoning: To avoid the red spikes, the character should move to the right. After moving right, the character will be in a position to slide under the green enemy.

Action: **RIGHT** ✅



Gemini-2.0-flash-exp

Observation: The left paddle is in the middle left of the screen. The right paddle is in the middle right of the screen. The ball is slightly below the center of the screen and moving towards the right paddle.

Reasoning: The ball is moving towards the right paddle. To prevent the ball from passing the right paddle, the right paddle needs to move down to intercept the ball. The left paddle is not in immediate danger, so no action is needed for the left paddle.

Action: **RIGHTDOWN** ✅

Qwen2VL 72B

Observation: The ball is moving towards the right paddle. The right paddle is currently positioned at the bottom of the screen, and the ball is above it.

Reasoning: To hit the ball, the right paddle needs to move up to align with the ball's trajectory. Moving the paddle up will increase the chance of successfully hitting the ball.

Action: **RIGHTUP** ❌

Figure 24: case studies group 2.

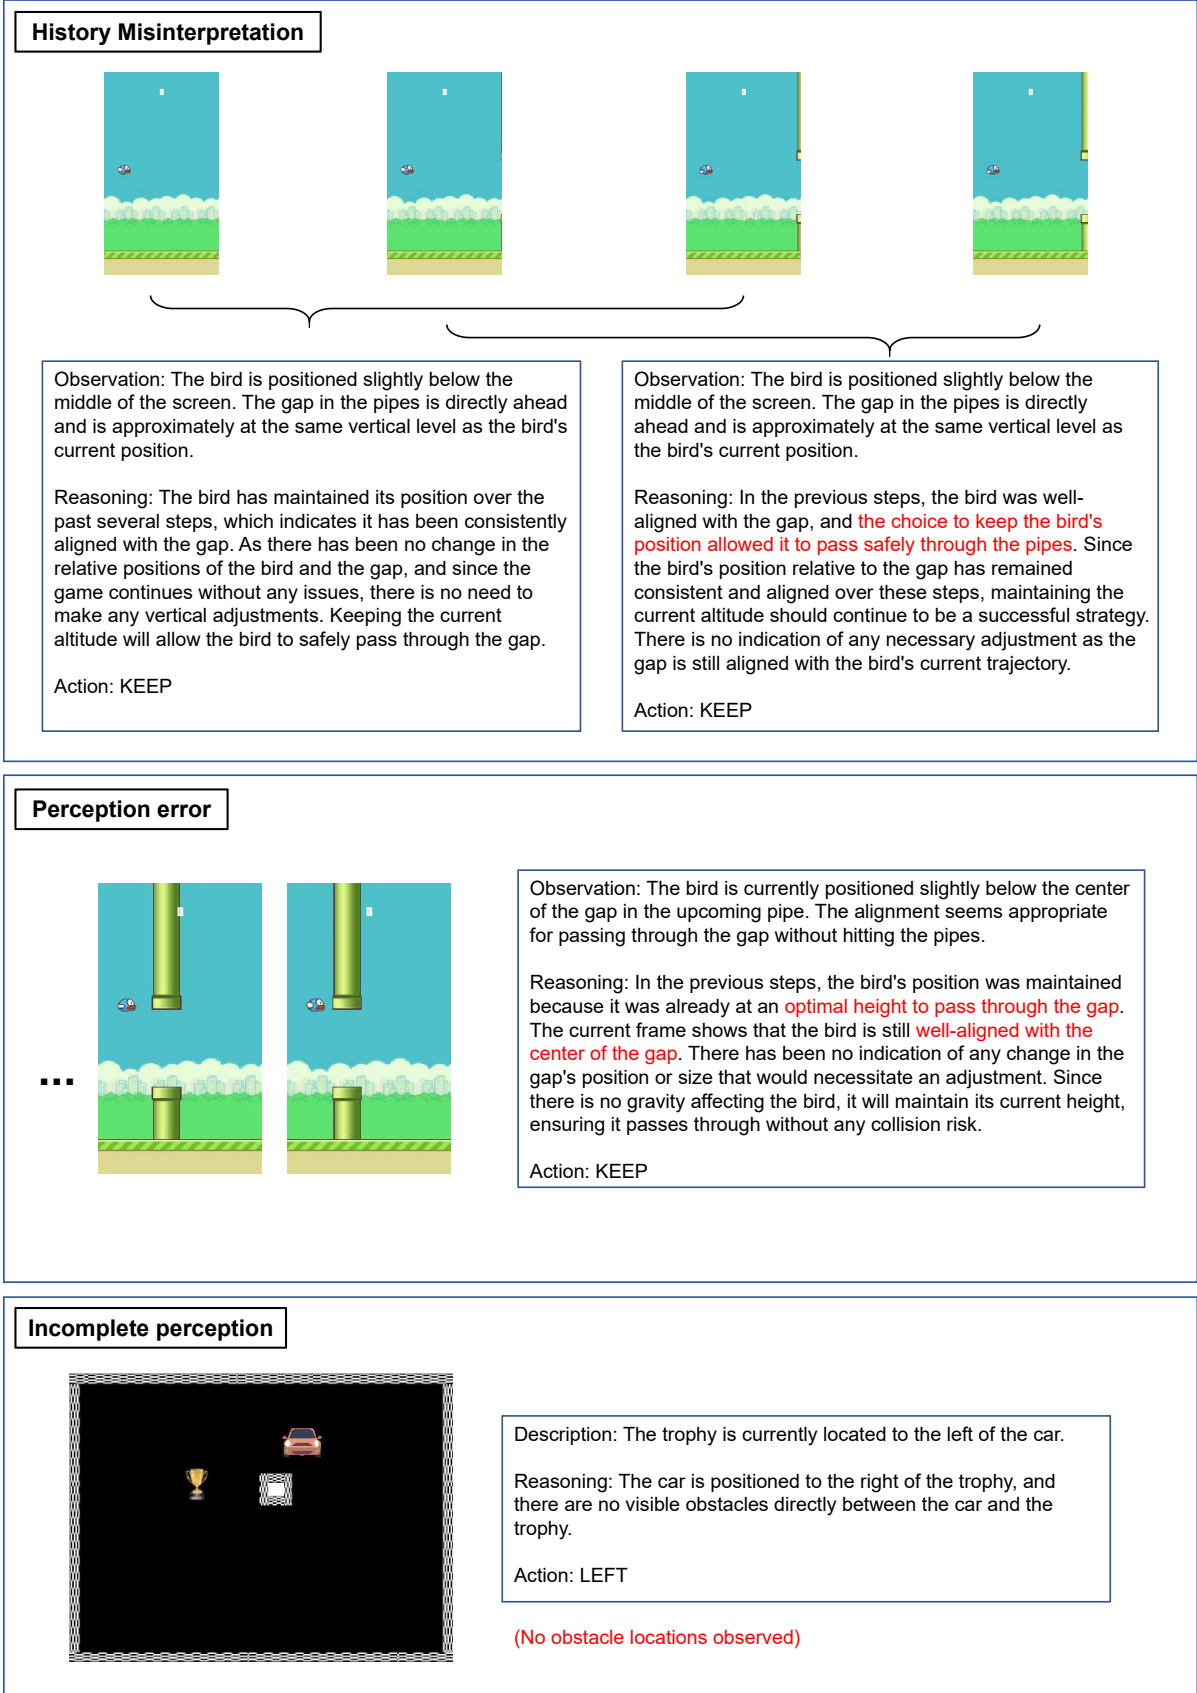


Figure 25: Error analysis in GPT4o cases.

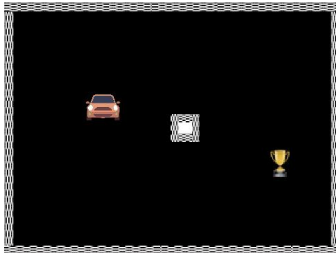
Inconsistency



Description: The trophy is currently positioned to the right of the car.

Reasoning: There is an obstacle, which appears to be a wall, located between the car and the trophy. The obstacle is directly in the path if the car attempts to move RIGHT towards the trophy. Therefore, the car should avoid moving directly towards the obstacle and instead find an alternative path.

Action: UP



Description: The trophy is located to the right and slightly below the car.

Reasoning: There is an obstacle (a square block) directly between the car and the trophy. Moving directly to the right would result in a collision with the obstacle. Therefore, the car must first navigate around the obstacle by moving downward, then move to the right to approach the trophy safely.

Action: DOWN

Action Inappropriate



Description: The trophy is to the right of the car.

Reasoning: The trophy is directly to the right of the car. There is a vertical obstacle (a wall) between the car and the trophy. The car needs to avoid colliding with the obstacle to reach the trophy.

Action: RIGHT

Perception error



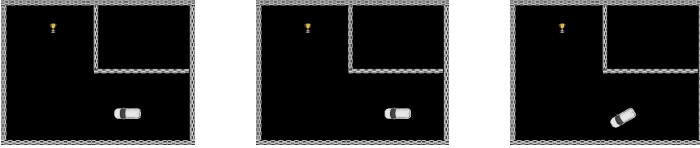
Description: The car is currently facing diagonally towards the bottom-right of the screen. The trophy is positioned in front of the car, slightly to the right. There are no visible obstacles between the car and the trophy, indicating a clear path. The car has been moving consistently forward in the previous steps, maintaining its alignment towards the trophy.

Reasoning: Given the car's orientation towards the trophy and the absence of obstacles, the car's current position and direction remain favorable for reaching the trophy. Continuing to move forward is the most efficient strategy to maintain the direct path to the trophy. The car's current trajectory should allow it to reach the trophy without any adjustments.

Action: FORWARD

Figure 26: Error analysis in GPT4o cases.

Direction Error




Description: The trophy is located to the top-left of the car. The car is currently facing diagonally upwards to the right side of the screen. There is an obstacle (a wall) directly to the left and slightly above the car, which separates it from the trophy. To reach the trophy, the car must navigate around this wall without colliding with any obstacles. The car is currently not moving but is positioned at an angle facing towards the top-right corner of the screen.

Reasoning: Since **the car is facing diagonally upwards to the right**, the immediate path to the trophy is obstructed by the wall directly to the left and slightly above. The car needs to move upwards and then turn left to navigate around the wall towards the trophy. From the driver's perspective, the car needs to first orient itself to face upwards more directly. The immediate action should be to adjust the car's direction by turning **LEFT**, which will help the car face upwards, allowing it to maneuver around the wall and approach the trophy.

Action: **LEFT**

Action Inappropriate




Description: The trophy is positioned in the **upper-right corner** of the frame, relative to the car's current position. The car is **facing directly towards the right side of the frame**. There is a vertical wall directly in front of the car, which extends upwards and prevents the car from moving forward directly towards the trophy. This wall creates a barrier between the car and the trophy, so the car cannot reach the trophy by moving straight ahead.

Reasoning: The car is positioned in such a way that if it moves forward, it will collide with the wall, resulting in game over. Since the car is facing right and the path to the trophy is blocked by the wall, the car needs to navigate around the obstacle to reach the trophy. The car must first turn right to begin moving around the wall. By turning right, the car can begin to reposition itself and align with a path that allows it to bypass the wall and proceed towards the trophy.

Action: **RIGHT**

Instruction Following

(In the prompt, we ask the model to output only one word in the Action, with no other content attached)



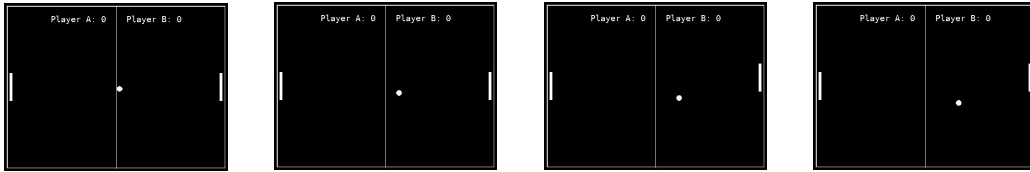
Description: The trophy is located to the right of the car.

Reasoning: There is a vertical obstacle (a white line) directly between the car and the trophy. The car cannot move directly to the right without colliding with this obstacle.

Action: **UP (to move the car up and around the obstacle)**

Figure 27: Error analysis in GPT4o cases.

Direction Error

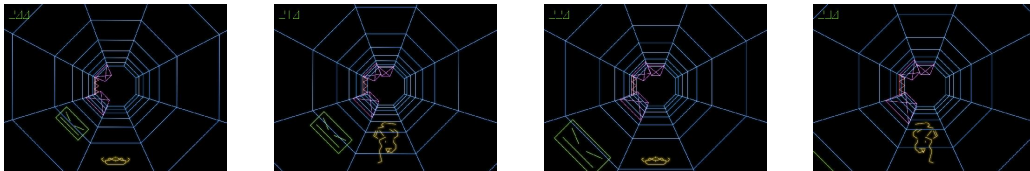


Observation: The left paddle is positioned in the middle of the left side. The right paddle is slightly above the center on the right side. The ball is moving towards the right side of the screen and is slightly above the center line, continuing its trajectory towards the right paddle.

Reasoning: The ball is heading towards the right paddle and is slightly **above** the center line. Given the ball's trajectory, it appears to be on a direct path towards the **upper part of the right paddle**. The right paddle is already positioned slightly above the center, which aligns well with the ball's current path. To ensure the ball is intercepted effectively, the right paddle should be moved up slightly to account for any upward drift in the ball's trajectory. The left paddle does not need any adjustment since the ball is moving away from it.

Action: RIGHTUP

Direction Error

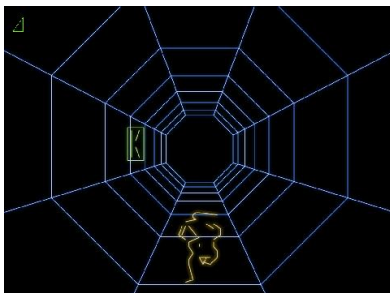


Observation: The character is currently in a normal running position at the bottom center of the tunnel. There is a purple wall approaching in the center of the tunnel.

Reasoning: The character is in a normal running position, and there is a purple wall **directly ahead**. To avoid the purple wall, the character should move around it by using the **LEFT** action to shift to the left. This will help the character to safely navigate through the tunnel without colliding with the wall.

Action: LEFT

Action Inappropriate



Observation: The character is currently running in the center of the tunnel, and there is a green enemy approaching in the **left section** of the tunnel ahead.

Reasoning: The green enemy is approaching on the left side, and to eliminate it, the character **needs to perform a slide**. Sliding will allow the character to duck and kick the green enemy. Since the character is currently running in the center, sliding will position them correctly to deal with the green enemy.

Action: SLIDE

Figure 28: Error analysis in GPT4o cases.