

# GlitchBench: Can large multimodal models detect video game glitches?

Mohammad Reza Taesiri<sup>1</sup>, Tianjun Feng<sup>1</sup>, Anh Nguyen<sup>2</sup>, Cor-Paul Bezemer<sup>1</sup>

<sup>1</sup>University of Alberta, {mtaesiri, robbie020428, bezemer}@ualberta.ca

<sup>2</sup>Auburn University, anh.ng8@gmail.com

## Abstract

Large multimodal models (LMMs) have evolved from large language models (LLMs) to integrate multiple input modalities, such as visual inputs. This integration augments the capacity of LLMs for tasks requiring visual comprehension and reasoning. However, the extent and limitations of their enhanced abilities are not fully understood, especially when it comes to real-world tasks. To address this gap, we introduce **GlitchBench**, a novel benchmark derived from video game quality assurance tasks, to test and evaluate the reasoning capabilities of LMMs. Our benchmark is curated from a variety of unusual and glitched scenarios from video games and aims to challenge both the visual and linguistic reasoning powers of LMMs in detecting and interpreting out-of-the-ordinary events. We evaluate multiple state-of-the-art LMMs, and we show that **GlitchBench** presents a new challenge for these models. Code and data are available at: <https://glitchbench.github.io/>

## 1. Introduction

The video game industry boasts an estimated annual revenue of USD 217 billion [57] with a total of 3.2 billion gamers worldwide in 2022 [1]. Automatically detecting in-game glitches is, therefore, a highly demanding task, but that remains a long-standing challenge [12, 39, 51, 55, 56, 65, 66, 72, 83]. A *glitch* is an unexpected frame that occurs within a game due to either an unforeseen software bug, player actions, or unanticipated interactions between game elements and does *not* result in a program crash. From a computer vision perspective, glitch detection involves recognizing an extremely wide spectrum of long-tail video frames, from rendering (e.g., 3D objects with missing textures; Fig. 2f), unrealistic physics (e.g., two people sitting in an invisible car; Fig. 2b), to semantic errors (e.g., indoor rain; Fig. 1). Unlike software bugs that can be caught by examining the code alone, glitches are more non-trivial to detect because they are mostly the result of a one-time event that unexpectedly happens in-game.

A holy grail of game quality assurance is to build a

GPT-4V fails to focus on the rain inside the room



**Question:** What is wrong with this image?

**Ground truth:** It is raining inside the building

**GPT-4V:** There’s nothing inherently “wrong” with this image; however, if you’re referring to inconsistencies or peculiarities that might not align with the established Batman canon, the red Batman suit is unusual since Batman is traditionally known in black or very dark colors. The use of red could be an artistic choice for the game, an alternate suit available within the game’s universe, or indicative of a different character who takes on the Batman mantle.

Figure 1. The image depicts a screenshot in which it rains inside a room. While the rain should be what is wrong with the image, GPT-4V fails to reason correctly and instead focuses on the color of Batman’s costume. Note that the ground truth is never presented as part of the prompt in our study.

general glitch detector that works for any game of any genre and mechanics. We set the first step toward this goal by building **GlitchBench**, an evaluation benchmark of 593 glitches, leveraging the public’s crowd knowledge from the game community’s reports on [reddit.com/r/GamePhysics](https://www.reddit.com/r/GamePhysics/). The glitches span across 205 games of various genres. Each glitch has a video clip, a representative frame, a one-line description, and a reference to a corresponding Reddit thread where gamers discussed the error.

Large image-text, multimodal models (LMMs), such as GPT-4V [2], are improving at an unprecedentedly fast pace. They excel in many existing tasks, including object detection [44, 75], multi-step reasoning [4, 5, 10, 35], and detailed image captioning [2, 38, 42, 52, 76]. Testing LMMs on **GlitchBench** may yield important findings not only to the game industry but also to the Artificial Intelligence (AI) community because glitch detection requires a combination of knowledge and understanding of image aesthetics, computer graphics, physics and commonsense reasoning (skills that are often tested individually in a benchmark [8]).

In this paper, we evaluate how well LMMs perform in detecting glitches from a single frame. Our main findings and contributions include:

1. We introduce **GlitchBench** that contains 593 images of video game glitches from 205 games for evaluating LMMs (Sec. 3).
2. We evaluate 11 state-of-the-art LMMs, including GPT-4V [2] and LLaVA [42] on our benchmark and in comparison with the performance on 6 other common benchmarks (Sec. 4).
3. LMMs are better at detecting glitches that violate simple physical laws (*e.g.*, a car flying in the air) than other more subtle glitches (*e.g.*, human limbs in an implausible pose; Fig. 6).
4. The state-of-the-art model on **GlitchBench** is GPT-4V with 43.4% accuracy. In the extensive captioning setup, we estimated the upper limits of models, and GPT-4V can achieve an accuracy of 64.9%, which is almost twice that of LLaVA, the second-best model (30.5%).
5. In sum, there exists a headroom of 30–35% on **GlitchBench** for future LMM models to improve, presenting an interesting challenge to the AI community.

## 2. Related Work

### 2.1. Multimodal, image-text datasets

Recently, there has been rapid development of large multimodal models that can process multiple modalities, including visual and textual inputs. Existing datasets that come with human-generated image captions, such as COCO Caption [13], Nocaps [3], CapFilt: [36] and Flickr30k [53], can serve as a simple way to evaluate language models. By providing the image, we can ask a model to describe it and then compare the generated caption with the ground truth [42, 43, 77]. Image captioning is a narrow domain and can be extended into visual question answering (VQA) by asking questions related to an image. Datasets like GQA [27], OK-VQA [49], VQAv2 [22], and Vizwiz [23] contain image-question pairs to probe the visual reasoning and understanding of LMMs.

Building upon simple VQAs, several benchmarks aim to increase the complexity of tasks over different dimen-

sions. TextVQA [63], OCR-VQA [50] and TextCap[62] propose questions about the text shown in the image. ScienceQA [47] and MathVista [48] focus on scientific topics and charts, while VCR [80] and Sherlock [80] focus on commonsense reasoning. Moreover, AI2D [26] is directed at questions concerning scientific diagrams, and IconQA [46] targets the comprehension of abstract diagrams. Each of these benchmarks is designed to push the boundaries of VQA systems by introducing specialized content that necessitates advanced reasoning and understanding.

There are also comprehensive evaluation frameworks that assess multimodal language models across a wider spectrum of capabilities. These evaluations extend beyond visual and textual reasoning to encompass a variety of skills such as generation, question answering, adherence to instructions, and the application of commonsense logic. Notable among these are SEED-Bench [33], MME [19], MM-Bench [45], MM-Vet [79], VisIT-Bench [8], which collectively serve to provide a robust measure of a model’s proficiency in handling tasks that integrate multiple modalities.

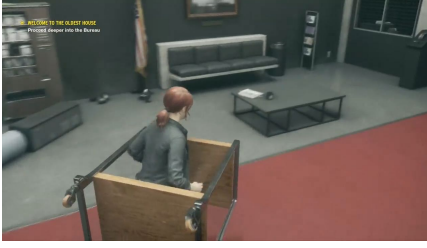
Unlike traditional datasets that contain queries about elements present in the image, our approach is novel in directing models to discern the atypical aspects, *i.e.*, glitches, with no linguistic hints provided. We show an image to the model and ask it to report unusual aspects of it. Such questions require a more integrated approach to visual and linguistic processing within an LMM to formulate a response.

### 2.2. Vision-language Stress Testing

Out-of-distribution (OOD) datasets have become a cornerstone for evaluating the capabilities and progress of machine learning models. In standard image classification, in particular the ImageNet [59] dataset, the introduction of datasets [24, 25, 25, 64] has underscored the importance of robustness and generalization in model evaluation. As we move from simple image classification tasks to more complex multimodal tasks, there is an increasing need for similar OOD datasets that can comprehensively test the generalization abilities of LMMs.

There are several studies that stress test various aspects of vision from different angles, such as compositional and spatial reasoning [20, 28, 29, 67], objects placed out of context and implausible scenes [9, 14, 84], and the exploitation of language and vision priors [18, 40].

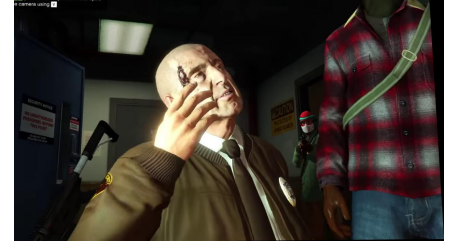
The closest benchmark to ours is Whoops [9], which is designed to challenge commonsense knowledge and reasoning in LMMs. However, our dataset differs in several ways: (1) The tasks in our dataset come from real-world tasks, specifically video game quality assurance, and are not purely intended to test models. (2) Whoops requires cultural and background knowledge to answer; for example, *A panda bear is catching salmon fish* is unusual since



(a) A person stuck in a piece of furniture



(b) Two people driving an invisible car



(c) A rifle floating in the air



(d) A person is floating in the air



(e) The gun in the hand is missing



(f) The table cover has a placeholder texture

Figure 2. Sample images from the **GlitchBench** showing glitches in various games with distinct styles. Samples (a)–(e) are captured from online videos, while sample (f) is generated inside the Unity game engine.

pandas subsist almost entirely on bamboo. In contrast, our dataset contains samples that contradict basic commonsense and the physical rules of the world. (3) Finally, images in Whoops are generated using Image-to-text models; they are clear without artifacts, centered in the image, and do not stress the visual side of the image, focusing on the context. In contrast, for our dataset, models need to fully scan the image to identify its unusual aspects, and there are many distracting elements present in the image, challenging them to focus on the correct part of the image.

### 2.3. Empirical Analysis of Recent LMMs

With the release of recent proprietary LLMs, such as GPT-4V and Bard [21], some studies attempt to evaluate and report the performance of these models on various benchmarks and tasks [16, 54, 73]. The main goal of these studies is to provide a comprehensive evaluation of the models across various well-established tasks and some narrow domains [71, 74]. The main difference between our work and these studies is that we have proposed a general stress-testing benchmarking dataset to measure the generalization power of various LLMs, both proprietary and open source, on a specific task related to the game industry.

## 3. GlitchBench

In this section, we describe the creation process of **GlitchBench**, a benchmark aimed at stress-testing visual perception and commonsense reasoning in LMMs, motivated by real-world game quality assurance tasks.

During development, video games go through many

stages of testing to reach certain quality standards before release. However, even after release, they can still exhibit unusual in-game events, or glitches. Glitches, often viewed as annoying bugs, can also possess a humorous and entertaining aspect. Players frequently report glitches across various social media platforms, particularly on Reddit and YouTube. A critical aspect of understanding glitches is the requirement of commonsense knowledge about the basic physical laws of the game’s universe, making them a suitable and practical candidate for testing machine learning models. Figure 2 shows six samples from **GlitchBench**.

### 3.1. Constructing the Dataset

In this section, we describe the creation process of our dataset. **GlitchBench** contains two parts: (1) 513 samples shared by players of video games, *i.e.*, frames collected from online sources, and (2) 75 synthetic samples.

**Samples shared by players of video games:** To construct our dataset, we sampled 1,000 videos from the GamePhysics [66] dataset. This dataset consists of videos from a subreddit with the same name, containing gameplay video clips with unusual events and glitches.

Next, we conducted a manual review process to filter videos based on two criteria: (1) the presence of a glitch in the video, and (2) the potentiality to detect the glitch from a single frame. The second criterion is essential because certain glitches, such as those involving rapid shaking or changes in size over time, cannot be detected from a still image alone.



After applying these filters, we extracted one frame from each remaining video, resulting in a collection of 650 samples. Our final round of manual reviews revealed two potential issues: (1) some glitches are not detectable from the extracted image and require more context to understand, and (2) some images contain the faces of gamers who streamed the content on an online platform (which could cause the LMM to identify these faces as what is wrong with the images). After removing videos that contain one of these issues, our final dataset contains 513 images.

**Generating synthetic samples with Unity:** To enhance our dataset, we supplemented samples from the GamePhysics dataset with 75 synthetic examples created inside the Unity game engine. These samples were specifically designed to mimic a subset of common development-stage bugs [39, 55, 65] that are not readily available in online social media platforms and, hence, to diminish the survivor bias effect. These flaws are often fixed before the public release of a game through the quality assurance process of a game development company and are therefore not often posted on social media.

Our synthetic sample generation process involves the injection of three categories of glitches into each scene: (1) placeholder textures, (2) object mesh distortions, and (3) low-resolution textures.

**Glitch-free Images:** Our main is on glitch images, as they are more challenging to capture and collect. However, to establish a baseline for comparison, we also included a set of glitch-free images. To accomplish this, we randomly selected gameplay walkthroughs from various games on YouTube. From these walkthroughs, we extracted a random subset of frames, resulting in the compilation of a dataset consisting of 330 frames sourced from a diverse array of games. The ground truth labels for these glitch-free images is *There is nothing wrong with this image*.

### 3.2. Labeling the Dataset

For all images, we provide a short description of the glitch present in the image. Our goal is to label the images briefly, highlighting only the unusual elements in simple language. For instance, if an image depicts a character with a contorted physique, the label would simply state, *“character has an unnatural body position”*.

It is important to highlight that some images can be described in many different ways. Diverse phrases such as *“falling from the sky”*, *“suspended in mid-air”*, or *“jumping in the air”* might all refer to a single event. Instead of handling such cases in the labeling process, in the evaluation process, we incorporate a language model to diminish the effect of this (see Sec. 4.1).

### 3.3. Categorizing the Glitch Types in the Images

In this section, we provide a high-level categorization of glitches in our dataset. While there have been some attempts to provide a taxonomy of video game bugs [32, 69], these taxonomies do not provide descriptions that are adequate to automate the bug categorization.

We propose a novel human-AI team-based method to build a categorization based on the descriptions of the images. This process is a collaborative effort between GPT-4 and humans, where GPT-4 suggests initial categories, and then humans refine these suggestions by providing feedback or asking the model to re-evaluate its output, harnessing the reflective ability of GPT-4 [61]. Finally, we manually bridge the resulting categories to those proposed by Lewis et al. [32] based on the semantics and instances of the glitches in our dataset.

**Process:** We prompt GPT-4 with all the glitch descriptions in our dataset and ask it to generate a categorization based on the descriptions and semantics of the glitches. In each subsequent iteration, we provide feedback in one of two ways: (1) we ask GPT-4 to review its previous answer through reflection, or (2) we explicitly instruct the model to merge two categories that are semantically similar. We stop when the model no longer changes its answer through reflection or when we can no longer merge categories.

In the last step, to assign each image to a category, we prompt GPT-4 with the description of the glitch and the final categories and ask it to assign each image to one of them. The final categories, the number of instances, examples for each category, and the parent category proposed by Lewis et al. [32] are outlined in Table 1.

## 4. Experiments

### 4.1. Experimental Setup

**Formulating Questions:** We designed **GlitchBench** as a free-format response benchmark, in contrast with traditional LMM benchmarks that utilize Yes/No or multiple-choice formats [19, 33]. We ask models to describe the unusual aspects of an image by answering three questions:

(Q1) *What is unusual about this image?*

(Q2) *What is wrong with this image?*

(Q3) *Describe the image in detail*

Note that we do not explicitly use the word *glitch* in the question, and we use simple language similar to what a layperson would use. During the inference, we allow models to come up with their own reasoning, and after the model generates the full response, we record it for further evaluation and comparison with the ground truth.

The rationale for free-format answers is that including an ‘unusual’ event description among choices hints to the LMM, letting it answer while disregarding visual aspects.

Table 1. Categorization of video game glitches in **GlitchBench**. Numbers highlighted in   show the number of images in each category. Categories highlighted in   show the corresponding categories proposed by Lewis et al. [32].

<b>Physics, Collision, and Spawn</b>	Images: 422
<span style="background-color: #e6f2ff;">(Non-Temporal → Invalid position)</span>	
<ol style="list-style-type: none"> <li>1. Objects and characters floating or stuck in the air.</li> <li>2. Characters or objects clipping through solid objects like walls, floors, or ground.</li> <li>3. Vehicles or characters falling under the game map.</li> </ol>	
<b>Animation and Pose</b>	Images: 75
<span style="background-color: #e6f2ff;">(Non-Temporal → Invalid graphical representation)</span>	
<ol style="list-style-type: none"> <li>1. Unusual or unnatural body positions.</li> <li>2. Characters in a T-pose or with distorted body parts.</li> <li>3. Incorrect animations for certain actions.</li> </ol>	
<b>Rendering and Texture</b>	Images: 67
<span style="background-color: #e6f2ff;">(Non-Temporal → Invalid graphical representation)</span>	
<ol style="list-style-type: none"> <li>1. Mesh stretches or objects with distorted shapes.</li> <li>2. Missing textures or objects displaying a “default” placeholder texture.</li> <li>3. Objects with low-resolution.</li> </ol>	
<b>Camera, User Interface and Lighting</b>	Images: 26
<span style="background-color: #e6f2ff;">(Non-Temporal → Invalid value change)</span>	
<ol style="list-style-type: none"> <li>1. Camera issues such as clipping inside objects or improper character views.</li> <li>2. In-game menus displaying incorrect elements.</li> <li>3. Shadows or lighting effects that do not match the environment.</li> </ol>	

We included question Q3 to assess whether the models can accurately report any glitches or unusual elements within the image in extensive captioning. Essentially, this question serves as a visual perception test, evaluating whether the models can identify and describe unusual aspects of the image in a more relaxed condition. For example, in the sample shown in Fig. 1, we test the model to see if it can identify the presence of rain in the room. In this case, it indicates that it is raining outside.

**Evaluation:** Following recent successes [8, 41, 78, 82] we employ a language model as a judge to evaluate the model’s responses. We use Llama-2-70B-Chat [68] to compare the model-generated text with the ground truth and determine whether the text conveys the same meaning or mentions the event highlighted by the ground truth (see Fig. 3).

We report the accuracy of each model on each tested question and present the average performance for **Q1** and **Q2** as the final benchmark result. **Q3** serves as the visual

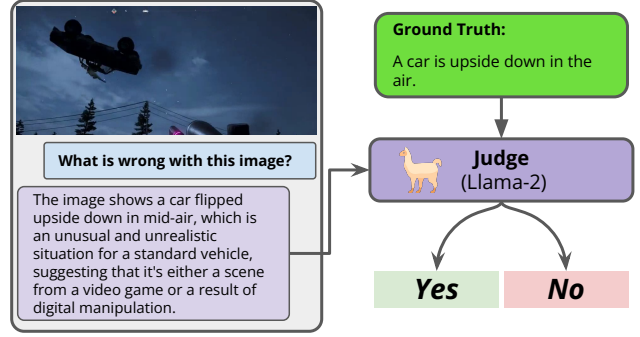


Figure 3. To evaluate a model’s response, we ask a judge (the Llama-2-70b-Chat model) to compare it semantically with the ground truth.

perception test, and we report the performance of the models on it separately.

To assess Llama-2’s judgment and determine if it can effectively serve as an evaluator, we manually reviewed a subset of responses for each model. For each model, we manually labeled 20 samples, with a total of 220 samples.

**Models:** In total, we evaluated 11 LMMs, including GPT-4V [2], and 10 open source models: LLaVA-1.5 (7B and 13B) [42], SPHINX (7B and 13B) [38], InstructBLIP (7B and 13B) [17], Qwen-VL-Chat (10B) [6], MiniGPT-v2 (7B) [11], OtterHD [34], and Fuyo (8B) [7]. We used the default temperature and top-p configurations provided with the model and API. We increased `max_token` to get full responses from models. (See Sec. A1 for details).

## 4.2. Quantitative Results

Table 2 shows the performance of all the tested models for the three questions. The Average performance on **Q1** and **Q2** is the main result of our benchmark. GPT-4V is the best-performing model, achieving 57.2% (**Q1**) and 29.5% (**Q2**) and an average of 43.4%. Next, LLaVA-1.5-13B achieves an average of 35.5% and is the best performing open-source model. These findings show **GlitchBench** is challenging for even state-of-the-art models.

The performance of GPT-4V on glitch-free images is higher than on glitchy images, with an average accuracy reaching 91.6%, suggesting that glitch-free images are much easier to handle.

Models exhibit different performance depending on the questions being asked, but all except for the SPHINX family show better performance when prompted with **Q1**. Nevertheless, the gap in performance varies, with GPT-4V showing the largest gap of 27.7pp (57.2 vs. 29.5). These results highlight that different prompts steer the behavior of LMMs differently and suggest that multi-step reasoning [31, 70] could also helpful with LMMs.

Table 2. Performance of various LMMs on **GlitchBench**. Numbers highlighted in   represent the average results of Q1 and Q2, which are the main results of the benchmark. Numbers related to Q3 serve as a visual perception test to measure the ability of models to report glitches in a relaxed manner. Numbers highlighted in   show the maximum agreement achievable with ground truth as perceived by Llama-2’s judgment (%). Numbers highlighted in   represent the results obtained from GPT-4V on glitch-free images.

Question	GPT-4V [2]		LLaVA-1.5 [42]		SPHINX [38]		InstructBLIP [17]		OtterHD [34]	Qwen -VL [6]	MiniGPT -v2 [11]	Fuyu [7]
	n/a	n/a	7B	13B	7B	13B	7B	13B	8B	10B	7B	8B
<b>Q1.</b> What is unusual about this image?	<span style="background-color: #c8e6c9;">88.2</span>	57.2	35.2	36.3	19.2	25.3	25.3	21.9	24.8	21.2	19.1	8.6
<b>Q2.</b> What is wrong with this image?	<span style="background-color: #c8e6c9;">95.5</span>	29.5	23.9	34.7	30.9	30.5	13.8	8.9	23.3	9.3	17.9	8.4
<b>Average</b>	<span style="background-color: #c8e6c9;">91.6</span>	43.4	29.6	35.5	25.0	27.9	19.6	15.4	24.0	15.2	18.5	8.5
<b>Q3.</b> Describe the image in detail.	-	64.9	28.0	30.5	17.5	21.9	16.0	11.8	21.6	14.0	16.0	7.6
<b>Maximum Agreement</b>	<span style="background-color: #c8e6c9;">95.5</span>	64.9	35.2	36.3	30.9	30.5	25.3	21.9	24.8	21.2	19.1	8.6

Our results also highlight that higher resolutions improve the performance. In particular, SPHINX-13B, which operates at a higher resolution than SPHINX-7B ( $448 \times 448$  vs.  $224 \times 224$ ), on average performs +2.9pp (27.9 vs. 25.0) better than the base model. Similarly, OtterHD, which employs Fuyu as the base model with enhanced flexibility and support for higher image resolutions, outperforms Fuyu on average by +15.5pp (24.0 vs. 8.5).

Asking LMMs to extensively caption the image using **Q3** only triggers GPT-4V to produce a very verbose response. In many cases, GPT-4V describes many details in the image and can touch upon the unusual aspects of the image. In this setup, GPT-4V can achieve 64.9%, which is an increase of +7.7pp over **Q1** and +21.5pp better than the benchmark results. This gap suggests that GPT-4V can *see* many details in the image, but it cannot easily focus on the unusual aspects of the image, indicating a gap in its reasoning capabilities across different modalities.

**Human evaluation:** Table 3 shows the results of the comparison between Llama-2 judgments and human evaluations, with the level of agreement for each model measured by Cohen’s Kappa [15]. Cohen’s Kappa demonstrates varying levels of concordance for each model. GPT-4V (0.80), InstructBLIP-7B (0.83), and Qwen-VL (1.00) exhibit substantial to perfect agreement. In contrast, OtterHD (0.50) had fair agreement, and Fuyu (-0.09) shows less than chance agreement, suggesting significant discrepancies. Overall, on all models except for Fuyu, we found above moderate agreement between Llama-2 and human judgment, while on six models, this agreement is substantial.

**Performance breakdown by category of glitches:** Figure 4 shows the breakdown of the performance of all tested models across the four studied glitch categories. GPT-4V is the best-performing model across all categories, with the exception of the *Rendering and Texture* category, where LLaVA-1.5-13B slightly outperforms it by 2.3pp (41.0% vs.

Table 3. Evaluating a subset of responses for comparing Llama-2 with human judgments: Llama-2 and humans exhibit moderate to substantial agreement on all models except for Fuyu.

Model	Llama-2	Human	$\kappa$
GPT-4V	60.0	50.0	0.80
LLaVA-1.5-13B	25.0	20.0	0.57
LLaVA-1.5-7B	35.0	15.0	0.49
Long-SPHINX	25.0	35.0	0.53
SPHINX	30.0	25.0	0.63
InstructBLIP-13B	20.0	10.0	0.62
InstructBLIP-7B	20.0	15.0	0.83
MiniGPT-v2	10.0	5.0	0.64
Qwen-VL	20.0	20.0	1.00
OtterHD	25.0	10.0	0.50
Fuyu	20.0	5.0	-0.09
$\mu \pm \sigma$	$26.4 \pm 12.8, 19.1 \pm 13.5$		0.64

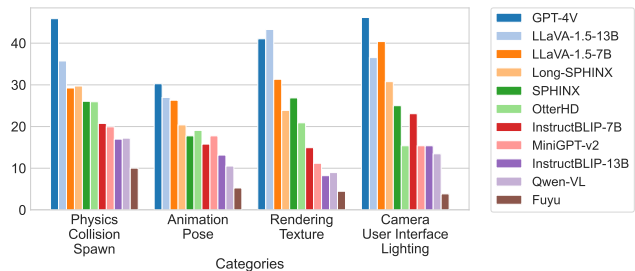


Figure 4. The performance of all tested models on different categories of images in **GlitchBench**.

43.3%). Overall, the *Animation and Pose* category consistently proves to be the most challenging. This category contains images of characters in unusual poses, distorted body joints, or twisted bodies. Fig. 6 shows a sample from this category with the character’s arm in an unnatural position.

### 4.3. Qualitative Observations and Analysis

#### Failing to reason about unusual aspects of the image:

We observed that in several cases, particularly in open-source models, the model reports phrases such as “*the problem with this image is that it is computer-generated*” or “*this is not an actual scene but a scene from a video game*”, along with similar phrases conveying the same meaning. These phrases suggest that, despite the model’s ability to *see* the content of the image, the language component of the model completely fails to reason about the content of the image.

Another observation is that InstructBLIP-13B often responds with “*nothing*” or similar phrases and completely fails to reason about the image. This is the reason why the smaller InstructBLIP-7B can achieve higher accuracy on **GlitchBench**. (See Sec. A3.1 for samples.)

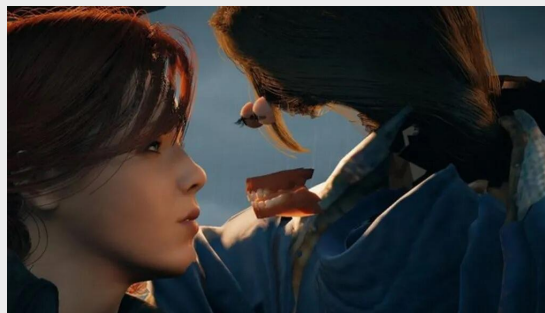
**GPT-4V struggles with faces:** GPT-4V is the best-performing model, yet it struggles with characters’ faces, as shown in Fig. 5. We found several issues when processing glitches related to faces, and in the majority of cases, GPT-4V fails to detect the glitch and sometimes hallucinates about characters wearing costumes ( Fig. A2), where there are basically no discernible facial features. On the other hand, smaller open-source models can sometimes detect glitches where GPT-4V fails, but they cannot describe the glitch clearly. We hypothesize that this might be due to the privacy features of GPT-4V, preventing it from seeing the face clearly. (See Sec. A3.2 for more samples.)

#### Poor performance in the *Animation and Pose* category:

Fig. 4 shows that *Animation and Pose* is the hardest category of glitches. During manual analysis, we found that LMMs struggle to detect unnatural body and limb configurations and incorrect animations being displayed. For instance, Fig. 6 shows an example of an unnatural arm position for a basketball player that GPT-4V cannot detect. This category can be further divided into three subcategories: (1) Heavily distorted body, *e.g.*, when a character’s body is overstretched or expanded; (2) Nuances in body configuration, such as contorted or twisted limbs and hands; and (3) Characters playing animations with invisible props (*e.g.*, a missing gun in the hand, as shown in Fig. 2e).

**Prevalent hallucination in open source models:** *Hallucination* typically refers to situations in which the model’s generated text contains information not present in the image [16, 58, 81, 85]. We noticed that open-source models often hallucinate extra objects or content in the image, *e.g.*, we found that Fuyu’s responses almost always contain hallucinations (see Sec. A3.4). The hallucination can contain additional characters or entirely new objects. These extra elements sometimes mislead Llama-2 into accepting an incorrect response as correct (see Sec. A3.5).

GPT-4V fails to recognize the absence of facial skin



**Question:** What is wrong with this image?

**Ground truth:** The character’s face mesh is not rendered correctly.

**GPT-4V:** The image appears to be from a video game or some form of CGI (computer-generated imagery) with very high-quality rendering. The characters depicted are very lifelike, suggesting a significant level of detail and investment in graphics. However, there’s nothing inherently “wrong” with the image in terms of content; it seems to portray a close-up moment between two characters.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. While it acknowledges the high quality of the graphics and the lifelike nature of the characters, it does not address the issue of the character’s face mesh not being rendered correctly, which is the main focus of the ground truth.

Figure 5. One of the several cases in which GPT-4V fails to detect a problem with facial features.

**Detecting some glitches requires paying attention to small details:** Different glitches in our dataset require varying levels of visual attention. For example, when a car is flying in the air, it usually occupies a large portion of the pixels on the screen, and models can easily pick up on such issues. This type of glitch is easier to catch, and GPT-4V, in particular, is very accurate at detecting it.

Some types of glitches require close attention to detail, such as clipping issues with clothing, where parts of the clothes intersect with the face or body of a character. While GPT-4V is generally the best model for detecting clipping, it is still not perfect. In some cases, GPT-4V misses the clipping, and in other cases, it hallucinates about clipping. (See Sec. A3.6 for samples.)



GPT-4V fails to detect unusual body configuration



**Question:** What is unusual about this image?

**Ground truth:** The left hand of the basketball player is in an unusual position and is stretched backward.

**GPT-4V:** The image appears to be a screenshot from a basketball video game, possibly from the NBA 2K series. The unusual aspects of this image include: One of the players has a bright orange beard that stands out from the rest. Another player seems to be kissing or blowing at the basketball, which is not a typical basketball move or gesture. The characters and graphics look computer-generated, which indicates it's from a video game rather than a real-life game. The interface at the bottom of the image shows various controls and buttons, which further suggests that it's from a video game.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth focuses on the unusual position of the basketball player's left hand, while the generated description highlights three unusual aspects of the image, none of which relate to the player's hand position.

Figure 6. The image shows a basketball player with an unnatural hand position. GPT-4V fails to focus on small details such as body configuration and is unable to report this issue.

## 5. Discussion and Limitation

### Comparing **GlitchBench** with other benchmarks:

The performance of various models across different benchmarks is presented in Tab. 4. It becomes evident that GPT-4V shows different performance against open-source models compared to **GlitchBench**. For instance, on VQAv2, LLaVA-1.5 and QWEN-VL score +5.8pp (80.0% vs 74.2%) and +5.3pp (79.5% vs 74.2%) higher than GPT-4V, respectively. However, on **GlitchBench**, they lag behind by -9.9pp (33.4% vs. 43.5%) and -28pp (15.4% vs. 43.4%).

Table 4. Comparing **GlitchBench** with other visual benchmarks — the bold numbers show the best model per benchmark (%)

Model/Task	Glitch (Ours)	VQAv2 [22]	OKVQA [60]	AI2D [30]	SEED [33]	POPE [37]	MMB [45]
GPT-4V	<b>43.4</b>	74.2	60.6	64.5	-	-	-
LLaVA	33.5	<b>80.0</b>	-	-	70.7	-	<b>67.7</b>
SPHINX	27.9	-	-	-	71.6	<b>90.8</b>	67.1
InstructBLIP	19.6	62.1	-	-	-	78.9	36.0
MiniGPT	18.5	-	57.0	-	-	-	-
QWEN-VL	15.4	79.5	58.6	62.3	58.2	-	60.6
OtterHD	15.2	-	-	-	-	86.1	58.5
Fuyu	8.5	77.4	<b>63.1</b>	<b>73.7</b>	-	-	-

The most notable gap is seen in Fuyu’s performance against GPT-4V: while Fuyu exceeds on both OKVQA and AI2D, it significantly lags behind on **GlitchBench** with only 8.5% compared to GPT-4V’s 43.4%.

In sum, across multiple existing LMM benchmarks, open-source models can perform on par with or even surpass GPT-4V. However, their performance on **GlitchBench**, which is derived from a real-world task in game quality assurance, falls significantly short of GPT-4V. In other words, the performance of models in real-world settings does not correlate well with existing benchmarks. This discrepancy partly comes from the design choices typical of LMM benchmarks, as they often opt for YES/NO or multiple-choice formats [19, 33, 45]. These formats allow models to find shortcuts for scoring high without necessarily generalizing well to other tasks.

**Limitation:** We constructed our dataset by randomly sampling videos and observed a prevalence of video games with an open-world genre on the Reddit website. Consequently, during our sampling process, video games from this genre, characterized by their distinct mechanics, were more frequently represented compared to other types.

## 6. Conclusion

We introduce **GlitchBench**, a new challenging benchmark for evaluating multimodal models on the video game glitch detection task. Detecting glitches requires various levels of reasoning skills, such as an understanding of the laws of physics and commonsense, making it well-suited for testing the generalization capabilities of large multimodal models. Comparing models’ performance on various multimodal benchmarks and **GlitchBench** reveals a disparity: good performance on prior benchmarks does not guarantee good performance on real-world tasks that demand extensive reasoning abilities. We show that **GlitchBench**, derived from real-world video game quality assurance, presents a new challenge for the AI community and is a valuable addition to existing multimodal benchmarks.



## References

- [1] Report: Epic games business breakdown & founding story. <https://research.contrary.com/reports/epic-games>. (Accessed on 11/15/2023). 1
- [2] OpenAI's GPT-4V(ision), 2023. 2, 5, 6
- [3] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019. 2
- [4] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 2
- [5] Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*, 2023. 2
- [6] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 5, 6
- [7] Rohan Bavishi, Erich Elsen, Curtis Hawthorne, Maxwell Nye, Augustus Odena, Arushi Somani, and Saḡnak Taşırılar. Introducing our multimodal models, 2023. 5, 6
- [8] Yonatan Bitton, Hritik Bansal, Jack Hessel, Rulin Shao, Wanrong Zhu, Anas Awadalla, Josh Gardner, Rohan Taori, and Ludwig Schimdt. Visit-bench: A benchmark for vision-language instruction following inspired by real-world use. *arXiv preprint arXiv:2308.06595*, 2023. 2, 5
- [9] Nitzan Bitton-Guetta, Yonatan Bitton, Jack Hessel, Ludwig Schimdt, Yuval Elovici, Gabriel Stanovsky, and Roy Schwartz. Breaking common sense: Whoops! a vision-and-language benchmark of synthetic and compositional images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2616–2627, 2023. 2
- [10] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023. 2
- [11] Jun Chen, Deyao Zhu1 Xiaoqian Shen1 Xiang Li, Zechun Liu2 Pengchuan Zhang, Raghuraman Krishnamoorthi2 Vikas Chandra2 Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: Large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. 5, 6
- [12] Ke Chen, Yufei Li, Yingfeng Chen, Changjie Fan, Zhipeng Hu, and Wei Yang. Glib: towards automated test oracle for graphically-rich applications. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1093–1104, 2021. 1
- [13] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 2
- [14] Myung Jin Choi, Antonio Torralba, and Alan S. Willsky. Context models and out-of-context objects. *Pattern Recognition Letters*, 33(7):853–862, 2012. Special Issue on Awards from ICPR 2010. 2
- [15] Jacob Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1): 37, 1960. 6
- [16] Chenhang Cui, Yiyang Zhou, Xinyu Yang, Shirley Wu, Linjun Zhang, James Zou, and Huaxiu Yao. Holistic analysis of hallucination in gpt-4v (ision): Bias and interference challenges. *arXiv preprint arXiv:2311.03287*, 2023. 3, 7
- [17] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. 5, 6
- [18] Stella Frank, Emanuele Bugliarello, and Desmond Elliott. Vision-and-language or vision-for-language? on cross-modal influence in multimodal transformers. *arXiv preprint arXiv:2109.04448*, 2021. 2
- [19] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xianwu Zheng, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 2, 4, 8
- [20] Tejas Gokhale, Hamid Palangi, Besmira Nushi, Vibhav Vineet, Eric Horvitz, Ece Kamar, Chitta Baral, and Yezhou Yang. Benchmarking spatial relationships in text-to-image generation. *arXiv preprint arXiv:2212.10015*, 2022. 2
- [21] Google. Bard, 2023. 3
- [22] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 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, 2017. 2, 8
- [23] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617, 2018. 2
- [24] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019. 2
- [25] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 2
- [26] Tuomo Hiippala, Malihe Alikhani, Jonas Haverinen, Timo Kalliokoski, Evanfiya Logacheva, Serafina Orekhova, Aino

- Tuomainen, Matthew Stone, and John A Bateman. Ai2d-rst: A multimodal corpus of 1000 primary school science diagrams. *Language Resources and Evaluation*, 55:661–688, 2021. 2
- [27] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 2
- [28] Huaizu Jiang, Xiaojian Ma, Weili Nie, Zhiding Yu, Yuke Zhu, and Anima Anandkumar. Bongard-hoi: Benchmarking few-shot visual reasoning for human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19056–19065, 2022. 2
- [29] Amita Kamath, Jack Hessel, and Kai-Wei Chang. What’s “up” with vision-language models? investigating their struggle with spatial reasoning. *arXiv preprint arXiv:2310.19785*, 2023. 2
- [30] Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 235–251. Springer, 2016. 8
- [31] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022. 5
- [32] Chris Lewis, Jim Whitehead, and Noah Wardrip-Fruin. What went wrong: a taxonomy of video game bugs. In *Proceedings of the fifth international conference on the foundations of digital games*, pages 108–115, 2010. 4, 5
- [33] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023. 2, 4, 8
- [34] Bo Li, Peiyuan Zhang, Jingkan Yang, Yuanhan Zhang, Fanyi Pu, and Ziwei Liu. Otterhd: A high-resolution multi-modality model. 2023. 5, 6
- [35] Chunyuan Li. Large multimodal models: Notes on cvpr 2023 tutorial. *arXiv preprint arXiv:2306.14895*, 2023. 2
- [36] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 2
- [37] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023. 8
- [38] Ziyi Lin, Chris Liu, Renrui Zhang, Peng Gao, Longtian Qiu, Han Xiao, Han Qiu, Chen Lin, Wenqi Shao, Keqin Chen, Jiaming Han, Siyuan Huang, Yichi Zhang, Xuming He, Hongsheng Li, and Yu Qiao. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models, 2023. 2, 5, 6
- [39] Carlos Ling, Konrad Tollmar, and Linus Gisslén. Using deep convolutional neural networks to detect rendered glitches in video games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 66–73, 2020. 1, 4
- [40] Fuxiao Liu, Tianrui Guan, Zongxia Li, Lichang Chen, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusion-bench: You see what you think? or you think what you see? an image-context reasoning benchmark challenging for gpt-4v (ision), llava-1.5, and other multi-modality models. *arXiv preprint arXiv:2310.14566*, 2023. 2
- [41] Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Aligning large multi-modal model with robust instruction tuning. *arXiv preprint arXiv:2306.14565*, 2023. 5
- [42] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 2, 5, 6
- [43] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 2
- [44] Shilong Liu, Hao Cheng, Haotian Liu, Hao Zhang, Feng Li, Tianhe Ren, Xueyan Zou, Jianwei Yang, Hang Su, Jun Zhu, et al. Llava-plus: Learning to use tools for creating multi-modal agents. *arXiv preprint arXiv:2311.05437*, 2023. 2
- [45] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023. 2, 8
- [46] Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. Iconqa: A new benchmark for abstract diagram understanding and visual language reasoning. *arXiv preprint arXiv:2110.13214*, 2021. 2
- [47] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022. 2
- [48] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023. 2
- [49] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019. 2
- [50] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *2019 international conference on document analysis and recognition (ICDAR)*, pages 947–952. IEEE, 2019. 2
- [51] Alfredo Nantes, Ross Brown, and Frederic Maire. A framework for the semi-automatic testing of video games. In *Pro-*

- ceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, pages 197–202, 2008. 1
- [52] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. 2
- [53] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. 2
- [54] Haotong Qin, Ge-Peng Ji, Salman Khan, Deng-Ping Fan, Fahad Shahbaz Khan, and Luc Van Gool. How good is google bard’s visual understanding? an empirical study on open challenges. *Machine Intelligence Research*, 20(5):605–613, 2023. 3
- [55] Farrukh Rahman. Weak supervision for label efficient visual bug detection. *arXiv preprint arXiv:2309.11077*, 2023. 1, 4
- [56] Geeta Rani, Upasana Pandey, Aniket Anil Wagde, and Vijaypal Singh Dhaka. A deep reinforcement learning technique for bug detection in video games. *International Journal of Information Technology*, 15(1):355–367, 2023. 1
- [57] Grand View Research. Video game market size, share and growth report, 2030, 2023. 1
- [58] Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*, 2018. 7
- [59] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 2
- [60] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *European Conference on Computer Vision*, pages 146–162. Springer, 2022. 8
- [61] Noah Shinn, Beck Labash, and Ashwin Gopinath. Reflexion: an autonomous agent with dynamic memory and self-reflection. *arXiv preprint arXiv:2303.11366*, 2023. 4
- [62] Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. Textcaps: a dataset for image captioning with reading comprehension. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 742–758. Springer, 2020. 2
- [63] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. 2
- [64] Mohammad Reza Taesiri, Giang Nguyen, Sarra Habchi, Cor-Paul Bezemer, and Anh Nguyen. Imagenet-hard: The hardest images remaining from a study of the power of zoom and spatial biases in image classification. 2
- [65] Mohammad Reza Taesiri, Moslem Habibi, and Mohammad Amin Fazli. A video game testing method utilizing deep learning. *Iran Journal of Computer Science*, 17(2), 2020. 1, 4
- [66] Mohammad Reza Taesiri, Finlay Macklon, and Cor-Paul Bezemer. Clip meets gamephysics: Towards bug identification in gameplay videos using zero-shot transfer learning. In *Proceedings of the 19th International Conference on Mining Software Repositories*, pages 270–281, 2022. 1, 3
- [67] Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. Winoground: Probing vision and language models for visiolinguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248, 2022. 2
- [68] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 5
- [69] Andrew Truelove, Eduardo Santana de Almeida, and Iftekhhar Ahmed. We’ll fix it in post: what do bug fixes in video game update notes tell us? In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 736–747. IEEE, 2021. 4
- [70] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. 5
- [71] Licheng Wen, Xuemeng Yang, Daocheng Fu, Xiaofeng Wang, Pinlong Cai, Xin Li, Tao Ma, Yingxuan Li, Linran Xu, Dengke Shang, Zheng Zhu, Shaoyan Sun, Yeqi Bai, Xinyu Cai, Min Dou, Shuanglu Hu, and Botian Shi. On the road with gpt-4v(ision): Early explorations of visual-language model on autonomous driving, 2023. 3
- [72] Benedict Wilkins and Kostas Stathis. Learning to identify perceptual bugs in 3d video games. *arXiv preprint arXiv:2202.12884*, 2022. 1
- [73] Yang Wu, Shilong Wang, Hao Yang, Tian Zheng, Hongbo Zhang, Yanyan Zhao, and Bing Qin. An early evaluation of gpt-4v(ision). *arXiv preprint arXiv:2310.16534*, 2023. 3
- [74] Zhiling Yan, Kai Zhang, Rong Zhou, Lifang He, Xiang Li, and Lichao Sun. Multimodal chatgpt for medical applications: an experimental study of gpt-4v, 2023. 3
- [75] Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*, 2023. 2
- [76] Zhuolin Yang, Wei Ping, Zihan Liu, Vijay Korthikanti, Weili Nie, De-An Huang, Linxi Fan, Zhiding Yu, Shiyi Lan, Bo Li, et al. Re-vilm: Retrieval-augmented visual language model for zero and few-shot image captioning. *arXiv preprint arXiv:2302.04858*, 2023. 2
- [77] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers

- large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023. 2
- [78] Zhenfei Yin, Jiong Wang, Jianjian Cao, Zhelun Shi, Dingning Liu, Mukai Li, Lu Sheng, Lei Bai, Xiaoshui Huang, Zhiyong Wang, et al. Lamm: Language-assisted multimodal instruction-tuning dataset, framework, and benchmark. *arXiv preprint arXiv:2306.06687*, 2023. 5
- [79] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023. 2
- [80] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6720–6731, 2019. 2
- [81] Ming Zhao, Peter Anderson, Vihan Jain, Su Wang, Alexander Ku, Jason Baldridge, and Eugene Ie. On the evaluation of vision-and-language navigation instructions. *arXiv preprint arXiv:2101.10504*, 2021. 7
- [82] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 5
- [83] Yan Zheng, Xiaofei Xie, Ting Su, Lei Ma, Jianye Hao, Zhaopeng Meng, Yang Liu, Ruimin Shen, Yingfeng Chen, and Changjie Fan. Wuji: Automatic online combat game testing using evolutionary deep reinforcement learning. In *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 772–784. IEEE, 2019. 1
- [84] Kankan Zhou, Eason Lai, Wei Bin Au Yeong, Kyriakos Mouratidis, and Jing Jiang. Rome: Evaluating pre-trained vision-language models on reasoning beyond visual common sense. *arXiv preprint arXiv:2310.19301*, 2023. 2
- [85] Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. Analyzing and mitigating object hallucination in large vision-language models. *arXiv preprint arXiv:2310.00754*, 2023. 7



# GlitchBench: Can large multimodal models detect video game glitches?

## Supplementary Material

### A1. Implementation Details

#### A1.1. Details about model inference

For each open source model, we used the provided sample code and demo from their respective repositories. Minor modifications were made to enable automatic processing of all images with designated prompts. The results were then stored in individual CSV files for each model. For OtterHD, which offers an API, we used the API to submit each image along with the appropriate prompt and recorded the responses. Our experiment was done prior to the official release of the GPT-4V API, and we used the ChatGPT web version for the benchmark, using a Chrome extension to assist in the process.

We kept the temperature and other parameters of each model unchanged. The only modification involved increasing the `max_token` limit, ensuring that the model’s response length was not restricted.

#### A1.2. Details about the judge

In our experiment, the Llama-2-70B model served as the judge. We utilized the API from [perplexity.ai](https://perplexity.ai), which is compatible with OpenAI’s Python package. Additionally, we employed a custom system message, as detailed below:

*Your task is to compare a model-generated text with a ground truth reference, assessing whether the key information and themes are similarly conveyed, even if worded differently. Focus on semantic content, thematic alignment, and intent, rather than exact phrasing or word usage. Recognize synonyms, paraphrases, and different stylistic expressions as valid, provided they faithfully represent the ground truth’s meaning. Offer feedback on the correlation between the texts and suggest improvements for alignment, while appreciating creative or varied linguistic expression that maintains the essence of the ground truth.*

*First analyze, then report the final answer in either of Yes or No*

## A2. Additional Results

### A2.1. Breakdown of Performance by Various Glitch Types

Table A1. Breakdown of Performance for Different LMMs by Various Glitch Types (%)

	Camera, User Interface, Lighting	Animation, Pose	Physics, Collision, Spawn	Rendering, Texture	Average Performance
GPT-4V	46.2	30.3	45.9	41.0	40.8
LLaVA-1.5-13B	36.5	27.0	35.7	43.3	35.6
LLaVA-1.5-7B	40.4	26.3	29.2	31.3	31.8
Long-SPHINX	30.8	20.4	29.7	23.9	26.2
SPHINX	25.0	17.8	26.1	26.9	23.9
OtterHD	15.4	19.1	25.9	20.9	20.3
InstructBLIP-7B	23.1	15.8	20.8	14.9	18.6
MiniGPT-v2	15.4	17.8	19.9	11.2	16.1
InstructBLIP-13B	15.4	13.2	17.0	8.2	13.4
Qwen-VL	13.5	10.5	17.2	9.0	12.5
Fuyu	3.8	5.3	10.0	4.5	5.9

### A2.2. Using multiple frames

As explained in Sec. 3, we went through our data to ensure that the glitches could be detected without using temporal information. To investigate whether current LMMs can effectively utilize temporal information between several input images, we conducted an experiment with a random sample of 150 videos. We extracted frames one second before and after the glitch from each video. We included the three frames and added ‘Given the sequence of images’ to the original GPT-4V prompt. The accuracy dropped to 36% (from 39%) for Q1 and to 28% (from 35%) for Q2. This suggests that current LMM performance cannot be improved by providing multiple frames.

### A3. Glossary of Model Failures

In this section, we offer a summary of our qualitative analysis, categorizing failures in multiple dimensions. These include instances where LMMs either fail to detect glitches or to note unusual aspects of images, as well as instances where the judge incorrectly labels the model’s responses as correct. Additionally, we address issues such as models producing hallucinations and other shortcomings.

- ➡ [Failures Related to Reasoning About the Content of the Image](#)
- ➡ [Failures Related to Facial Features](#)
- ➡ [Failures Related to Unnatural Body Positions](#)
- ➡ [Failures Related to Multimodal Hallucination](#)
- ➡ [Failures Related to the Judge Accepting the Wrong Answer](#)
- ➡ [Failures Related to Clipping Issues](#)
- ➡ [Failures Related to Unusual Circumstances](#)

### A3.1. Failures Related to Reasoning About the Content of the Image

InstructBLIP-7B fails to reason about the image and outputs *nothing*.



**Question:** What is wrong with this image?

**Ground truth:** A person is entirely black despite being in front of the flashlight.

**InstructBLIP-7B:** nothing

InstructBLIP-13B fails to reason about the image.



**Question:** What is wrong with this image?

**Ground truth:** A person is floating in the air in a horizontal position.

**InstructBLIP-13B:** There is nothing wrong with this image.

Figure A1. Samples for which models failed to reason about the content of the image and instead responded with *nothing* and similar phrases.



### A3.2. Failures Related to Facial Features

GPT-4V fails to recognize the absence of the facial skin



**Question:** What is wrong with this image?

**Ground truth:** The person is missing facial skin.

**GPT-4V:** This image appears to be from a character creation screen of a video game, where players can customize the appearance of their character. What's wrong or unusual is that the character's head has been replaced with that of a chimpanzee. This is likely a humorous or whimsical modification either due to a mod, a glitch, or a special feature within the game that allows such customization. Normally, in character customization screens, the options for head shapes and features are human, so this stands out as intentionally odd or comically absurd.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth states that a person is missing facial skin, which is a serious and potentially disturbing condition. In contrast, the generated description describes a humorous or whimsical modification to a video game character's appearance, which is a completely different theme and tone. The generated description does not convey the same level of seriousness or medical nature as the ground truth.

**Feedback:** The generated description could benefit from a more serious and realistic tone, as well as a focus on the medical aspects of the condition. It would be more accurate to describe the image as showing a person with a serious medical condition, rather than a humorous or whimsical modification to a video game character's appearance. Additionally, the description could benefit from more precise language, such as using terms like "skin absence" or "facial skin deficiency" to describe the condition, rather than the vague term "missing facial skin."

Figure A2. GPT-4V not only fails to recognize the absence of the facial skin, but it also hallucinates that the character's head is replaced with that of a chimpanzee.

GPT-4V fails to recognize the absence of the head.



**Question:** What is wrong with this image?

**Ground truth:** The person is missing their head.

**GPT-4V:** This image appears to be a screenshot from a video game, most likely “Cyberpunk 2077,” given the distinctive futuristic urban environment, user interface, and character designs that match the style of the game. There seems to be nothing inherently “wrong” with the image in terms of gameplay or the environment depicted, as it looks like a typical scene from within the game where the player is engaged in combat.

However, if you’re looking for discrepancies or errors:

1. The “Defeat all enemies” counter shows (6/7), suggesting one enemy is left, but we can’t see any enemies besides the character in the center. This might imply that the last enemy is not in the player’s view or that there is a glitch where the enemy isn’t appearing.
2. The mission titled “SPACE ODDITY” in the top right corner suggests a specific quest or objective, and the player has a sword named “BLACK UNICORN” equipped, as indicated by the weapon’s name at the bottom of the screen. Both are elements of the game and don’t indicate anything incorrect unless they are not supposed to appear together in the context of the game, which cannot be determined from this image alone.
3. The image shows a first-person perspective where the player is holding a sword. However, it could be considered unusual or suboptimal to use a melee weapon like a sword when the enemy appears to be wearing heavy armor with illuminated sections, possibly indicating a high level of threat where firearms or other strategies might be more effective.

Unless there’s a specific context or element not aligning with the game’s intended design or a technical issue you’re referring to, the image looks like a standard gameplay moment from “Cyberpunk 2077.”

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth.

The ground truth statement “The person is missing their head” is a straightforward description of a disturbing and abnormal situation. The generated description, on the other hand, is a detailed analysis of a screenshot from a video game, discussing various elements of the image such as the game’s environment, character designs, weaponry, and possible glitches or discrepancies. While it provides a thorough examination of the image, it completely misses the mark in terms of conveying the actual content and meaning of the ground truth. ...

Figure A3. GPT-4V provides a detailed description of the image, yet it fails to notice that the head is missing. The judge’s response truncated to save space.

### A3.3. Failures Related to Unnatural Body Positions

GPT-4V fails to detect unnaturally stretched hands.



**Question:** What is unusual about this image?

**Ground truth:** A person with unusually stretched hands.

**GPT-4V:** The image depicts a scene from what appears to be a video game, likely from the “Assassin’s Creed” series given the interface and character design. A character is interacting with a lock on a chest, while a downed enemy is nearby. Here are a few points that might be considered unusual, depending on context: The character is lockpicking the chest in broad daylight and seemingly in the open, which might be risky in a game focused on stealth. The nearby downed enemy suggests there was a recent confrontation, yet the character seems undisturbed while focusing on the lock. The interface shows various buttons, indicating the game is being played on a console. The button prompts suggest specific actions the player can take. The player’s status, such as “Elite Veteran” and username “Elvan8234”, is displayed on the screen.

Without further context, it’s hard to pinpoint exactly what might be unusual, as different viewers might interpret the scene in various ways based on their familiarity with the game or the specific situation presented.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth.

The ground truth describes a person with unusually stretched hands, which is not mentioned in the generated description. Instead, the generated description focuses on the details of a video game scene, such as the character’s interaction with a lock on a chest, the nearby downed enemy, and the game’s interface. While the generated description provides some context and details, it does not align with the main theme of the ground truth, which is the person’s unusual physical characteristic.

Figure A4. GPT-4V can detect some details from the image and the franchises of video games, yet it cannot recognize that the hands of the character shown in the image are unnaturally stretched to reach the box.



LLaVA-1.5 fails to notice the leg position of the character.



**Question:** What is unusual about this image?

**Ground truth:** An unnatural leg position of a football player.

**LLaVA-1.5-13B:** The unusual aspect of this image is that it is a computer-generated (CGI) image of a soccer player kicking a soccer ball. The player is wearing a red and white uniform, and the scene appears to be a soccer match. The CGI nature of the image suggests that it is not a real-life photograph but rather a digital representation of a soccer player in action. This type of image is typically used in video games, movies, or other digital media where realistic depictions of sports are required.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth.

The ground truth states that the image depicts an unnatural leg position of a football player, which is not explicitly mentioned in the generated description. The generated description focuses on the fact that the image is a CGI representation of a soccer player kicking a soccer ball, rather than the unusual leg position of the player.

While the generated description provides some relevant information about the image, such as the player's uniform and the possibility that it may be used in digital media, it does not accurately convey the essence of the ground truth. To improve the alignment, the generated description should focus more on the unusual leg position of the player and the fact that it is an unnatural position for a football player.

Figure A5. LLaVA-1.5 can describe the image content and details, such as the character's clothing, but it fails to notice that the character's leg is bent in an unnatural way.



Qwen-VL fails to notice the unnatural positions of the hands.



**Question:** What is unusual about this image?

**Ground truth:** A person with unnatural hand positions.

**Qwen-VL:** The unusual aspect of this image is that it is a screenshot taken from a video game, featuring a character pointing a gun at another character. This is not something one would typically encounter in real life.

**Llama-2 (Judge):** No.

The generated description does not accurately convey the key information and themes of the ground truth. The ground truth focuses on the unusual hand positions of a person, while the generated description describes a scene from a video game featuring a character with a gun. The themes of violence and gaming are present in the generated description, but they are not related to the theme of unusual hand positions.

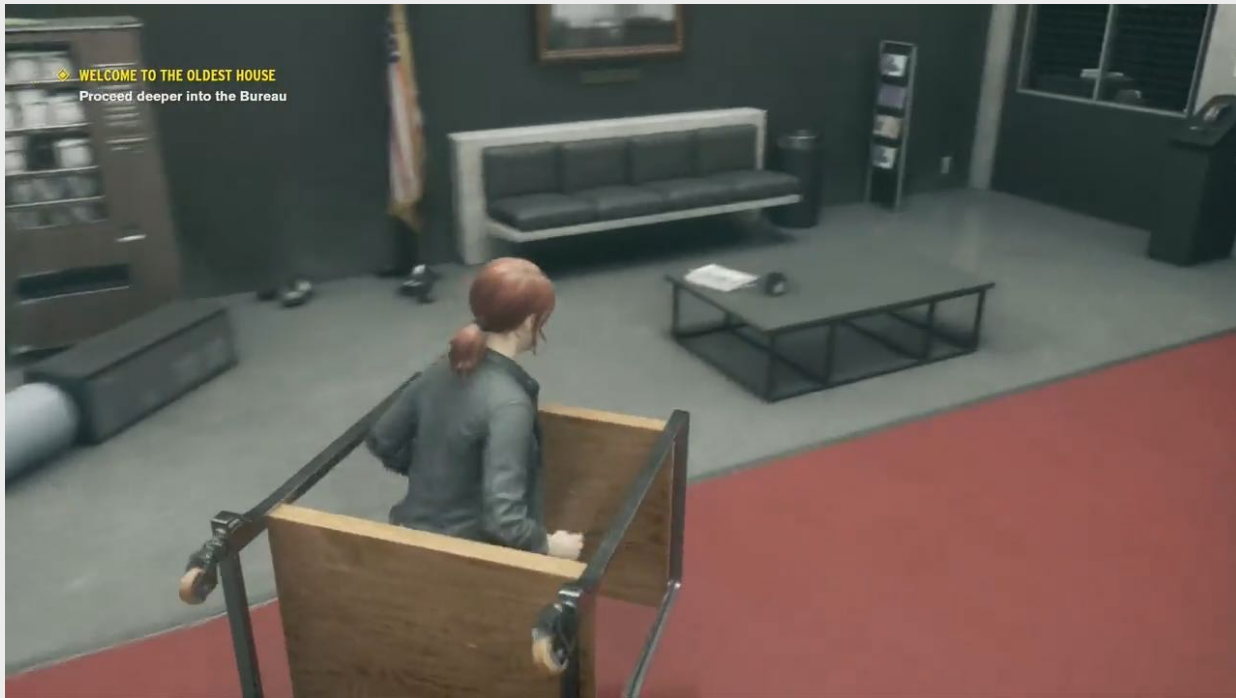
Even though the generated description uses different words and phrasing, it does not faithfully represent the meaning of the ground truth. The semantic content and thematic alignment are not present, and the intent of the generated description is different from the ground truth.

To improve the alignment, the generated description could focus on the unusual hand positions and exclude the elements of violence and gaming. For example, “A person with their hands in unnatural positions, such as fingers spread widely or palms facing upwards.” This would maintain the essence of the ground truth while allowing for creative or varied linguistic expression.

Figure A6. Qwen-VL can describe the image correctly, but it fails to understand the character’s body configuration and the unnatural hand positions.

### A3.4. Failures Related to Multimodal Hallucination

GPT-4V hallucinates about a person pushing a photocopier.



**Question:** What is unusual about this image?

**Ground truth:** A person stuck in a piece of furniture.

**GPT-4V:** The image appears to be from a video game, which can be inferred from the visual style and the overlaid text that reads "WELCOME TO THE OLDEST HOUSE - Proceed deeper into the Bureau". The unusual aspect is that it depicts a third-person view of a character pushing a photocopier inside an office environment. This is an uncommon action within the context of a video game, as photocopiers are typically static objects in such settings. The scenario suggests a gameplay element or a task within the game that involves moving office equipment, which is not a conventional video game objective.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth.

The ground truth describes a person stuck in a piece of furniture, which is not mentioned in the generated description. Instead, the description focuses on the video game aspect and the character's actions, which are not related to the ground truth. The themes of being stuck and struggling to get out are not present in the generated description, and the intent of the image is not accurately conveyed.

While the generated description does contain some creative and varied linguistic expressions, such as "WELCOME TO THE OLDEST HOUSE" and "Proceed deeper into the Bureau," they do not align with the meaning of the ground truth. To improve the correlation between the texts, the generated description could focus more on the person's predicament and their attempts to escape, rather than the video game context and actions.

Figure A7. GPT-4V starts by describing the image as a video game screenshot and then begins to read the text on the image. Regarding the actual content of the image, it fails to understand the content and hallucinates about a person pushing a photocopier.

GPT-4V hallucinates about a mirror.



**Question:** What is unusual about this image?

**Ground truth:** A towel is floating in the air.

**GPT-4V:** The image appears to be a screenshot from a video game, featuring a first-person perspective where the player's hand and a gun are visible, aiming towards a mirror. What is unusual is that the mirror does not reflect the player's hand or the gun; instead, it shows a room with furniture, which suggests a graphical error or a deliberate design choice within the game. Mirrors in video games are often challenging to render correctly due to the computational demand of real-time reflections.

**Llama-2 (Judge):** No.

The generated description does not accurately convey the key information and themes of the ground truth. The ground truth simply states that a towel is floating in the air, whereas the generated description describes a scene from a video game with a first-person perspective, a mirror, and furniture. The themes of a towel floating and a video game are unrelated, and the generated description does not capture the essence of the ground truth.

While the generated description recognizes synonyms and paraphrases, such as "a first-person perspective" instead of "a person holding a gun," it does not maintain the essence of the ground truth. The creative or varied linguistic expression in the generated description does not align with the ground truth, and the correlation between the two texts is weak.

To improve the alignment, the generated description could focus on the visual elements in the image, such as the floating towel, and avoid introducing unrelated elements like a video game or furniture. A revised description that accurately conveys the key information and themes of the ground truth could be: "A white towel is seen floating in mid-air, with no visible means of support." This description maintains the essence of the ground truth while using different words and phrasing.

Figure A8. While GPT-4V detects that the image is a screenshot of a first-person video game with a person holding a weapon, it fails to detect the floating towel; instead, it hallucinates about a mirror.



**Question:** What is wrong with this image?

**Ground truth:** Part of the house's external wall has a low-resolution texture.

**InstructBLIP-7B:** The image is unusual because it features a horse standing in front of a wooden house with a thatched roof. Horses are typically associated with outdoor environments, such as fields or stables, rather than urban settings like the one depicted in this image. Additionally, the presence of a clock and a dog in the scene adds to the unconventional nature of the image, as these elements are not typically found in close proximity to a horse and a wooden house with a thatched roof.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth statement simply describes a physical characteristic of a house, whereas the generated description describes a scene with a horse, clock, and dog, and focuses on the unusual nature of the image. The generated description does not mention the low-resolution texture of the house's external wall, which is the main focus of the ground truth.

To improve the alignment between the generated description and the ground truth, the model could be fine-tuned to focus more on the physical characteristics of the house, rather than the objects in the scene. Additionally, the model could be trained to recognize and use more specific vocabulary related to architecture and building materials, such as "low-resolution texture" instead of "unusual."

However, it's worth noting that the generated description does show some creativity and varied linguistic expression, and it could be considered a valid description of the image from a different perspective.

Figure A9. Sample hallucination made by InstructBLIP-7B. The model mentions the presence of a horse in front of a wooden horse.



LLaVA-1.5-7B hallucination about person holding a book.



**Question:** What is wrong with this image?

**Ground truth:** A person with their hand stretched up.

**LLaVA-1.5-7B:** The unusual aspect of this image is the presence of a person wearing a red coat and a hat, who is standing in a room with a typewriter and a person lying on the floor. The person in the red coat is holding a book, and there is a computer mouse nearby. The scene appears to be a mix of old-fashioned and modern elements, as the person is dressed in a red coat and hat, while the presence of a computer mouse and a typewriter suggests a more traditional setting. The combination of these elements creates an intriguing and unexpected visual experience.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth.

The ground truth image contains a person with their hand stretched up, which is not mentioned in the generated description. The description focuses on the person's clothing, the presence of a typewriter and computer mouse, and the mixing of old-fashioned and modern elements, which are not the main aspects of the ground truth image.

To improve the alignment between the generated description and the ground truth, the model should focus on the person's gesture of stretching their hand up and the overall simplicity of the image, rather than the details of the person's clothing and the surrounding objects. A possible revised description could be:

"A person with their hand stretched up, surrounded by a minimalist background, creating a sense of isolation and focus."

This revised description maintains the essence of the ground truth image while using different wording and emphasizing the key elements of the image.

Figure A10. While LLaVA-1.5-7B provides some details about the image, such as a person lying on the ground, it hallucinates the presence of a computer mouse and a person holding a book.

Fuyu-8B hallucinates all the details about the image.



**Question:** What is wrong with this image?

**Ground truth:** A car is flying in the air.

**Fuyu-8B:** In the image, there is a person standing on a street corner, wearing a mask and holding a weapon. The scene appears to be altered, possibly to make it appear as if the person is a zombie or a masked person. The weapon, a revolver, is visible, and the scene appears to be set in a surreal or futuristic setting. The altered nature of the image, combined with the altered perspective, creates an interesting and unusual visual effect.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth describes a car flying in the air, while the generated description describes a person standing on a street corner with a weapon, in a possibly surreal or futuristic setting. The themes of transportation and movement are not present in the generated description, and the focus on a person and a weapon suggests a different theme altogether.

While the generated description does use creative and varied linguistic expressions, such as “altered perspective” and “surreal or futuristic setting,” the overall meaning and intent of the description does not align with the ground truth. To improve the alignment, the generated description could focus more on the movement and transportation themes present in the ground truth, and strive to convey the idea of a car flying in the air in a creative and visually engaging way.

Figure A11. The screenshot displays a car floating in the air, but Fuyu-8B provides incorrect details and generates numerous hallucinations. All details are incorrect.

MiniGPT-v2 hallucinates about everything, including a red fire truck.



**Question:** What is wrong with this image?

**Ground truth:** A horse is in an unusual position with its front legs raised and resting on a ledge of a wooden feeding trough.

**MiniGPT-v2:** In the image, there are two men standing near an old-looking building with wood planks. One man is pointing a bow and arrow at the other man who is holding a large, long-handled spoon that looks like it could be used for a bowling ball. There's also a wooden table nearby, but the scene appears to be quite peculiar as there are a few horses in the vicinity, as well as two dogs, one of which is sitting on the ground.

There is also a large red fire truck that appears to be partially buried in the ground nearby. The combination of people and animals along with the unconventional weapons and objects in this setting may imply some sort of unusual or fictional activity, possibly a video game scenario.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth describes a horse in an unusual position, resting its front legs on a ledge, while the generated description focuses on the presence of men, horses, dogs, and a fire truck in a setting that appears to be a video game scenario. The generated description does not mention the horse's unusual position or the wooden feeding trough, and instead introduces unrelated elements such as the bow and arrow, the large spoon, and the partially buried fire truck.

The lacunarity in the generated description suggests that the model has not fully captured the essence of the ground truth, and may have been influenced by other factors such as the presence of people and animals in the image, or the suggestion of a fictional or video game-like scenario. To improve the alignment between the generated description and the ground truth, the model could benefit from additional training on images that depict unusual or unexpected scenarios, as well as reinforcement of the importance of focusing on the main subject and its unique features, rather than peripheral elements.

That being said, the generated description does display some creative and varied linguistic expression, using phrases such as "unusual or fictional activity" and "large, long-handled spoon that looks like it could be used for a bowling ball." These expressions suggest that the model has some ability to generate engaging and descriptive language, which could be leveraged to improve the accuracy and fidelity of its descriptions.

Figure A12. The screenshot shows a scene from The Witcher 3 game in which a horse has an unusual posture. MiniGPT-v2 provides incorrect details and generates numerous hallucinations. Almost all details are incorrect.



SPHINX hallucinates about “a dog is standing on its hind legs next to a car”.



**Question:** What is wrong with this image?

**Ground truth:** A person is attached to a car's door.

**SPHINX:** The unusual aspect of this image is that a dog is standing on its hind legs next to a car, which is not a typical behavior for a dog. This scene is likely a staged or edited photograph, as it is not common for a dog to be able to stand on its hind legs in such a manner. The dog's position and the car's location on the road create an interesting and unexpected visual effect.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth.

The ground truth states that a person is attached to a car's door, which suggests a potentially dangerous or emergency situation. In contrast, the generated description focuses on the unusual behavior of a dog standing on its hind legs next to a car, which is not the main focus of the ground truth. Additionally, the generated description implies that the scene is staged or edited, which is not mentioned in the ground truth.

To improve the alignment between the generated description and the ground truth, the model should focus more on the person's attachment to the car's door and the potential danger or urgency of the situation. The description could be rephrased to emphasize the unexpected nature of the person's attachment and the car's location, while still acknowledging the dog's unusual behavior. For example: “An unexpected scene unfolds as a person is seen attached to the door of a car, while a dog stands on its hind legs nearby. The person's attachment to the car creates a sense of urgency, and the dog's unusual behavior adds to the overall unexpected nature of the scene.”

Figure A13. Sample image in which SPHINX hallucinates about a dog. However, some details about the image, such as the presence of a car, are correct.



### A3.5. Failures Related to the Judge Accepting the Wrong Answer

Llama-2 aligns “desert-like environment” with the “pixelated character floating in the air”.



**Question:** What is wrong with this image?

**Ground truth:** A pixelated character is floating in the air.

**LLaVA-1.5-13B:** In the image, there is a man standing in front of a tent, wearing a long coat and holding a sword. The man appears to be a character from a video game, possibly a Viking or a warrior. The scene seems to be set in a desert-like environment.

There are a few minor issues in the image. The man's sword is positioned in a way that it appears to be floating in the air, rather than being held properly. Additionally, there is a small glitch in the image where the man's head appears to be cut off, making it difficult to see his facial features. These issues suggest that the image might be a screenshot from a video game, and the game's developers might need to address these bugs and glitches to improve the overall gaming experience.

**Llama-2 (Judge):** Yes.

The generated description successfully conveys the key information and themes of the ground truth, despite some differences in wording and phrasing. The reference to a “pixelated character” in the ground truth is interpreted by the model as a “man standing in front of a tent” who appears to be a video game character, possibly a Viking or a warrior. The mention of a “desert-like environment” aligns with the pixelated character floating in the air.

The model's description also creatively includes details not present in the ground truth, such as the character's long coat and the sword being held in a way that appears to be floating. These additional elements maintain the essence of the ground truth while offering a unique expression.

However, there are minor issues in the image that the model has identified, such as the glitch in the man's head and the sword appearing to be floating. These observations suggest that the image might be a screenshot from a video game, and the game's developers might need to address these bugs and glitches.

Overall, the generated description effectively conveys the key information and themes of the ground truth and offers some creative variations while maintaining the essence of the original meaning. The model's output demonstrates a good understanding of the semantic content and thematic alignment, and its feedback on the minor issues in the image is appreciated.

Figure A14. Llama-2 is confused and accepts the wrong response. The response generated by LLaVA-1.5-13B hallucinates a floating sword, which is incorrect. However, Llama-2 matches “desert-like environment” with the ground truth: “A pixelated character is floating in the air.”

Llama-2 aligns “unique and exciting match-up” with “clipping into each other”.



**Question:** What is wrong with this image?

**Ground truth:** Two characters are clipping into each other.

**MiniGPT-v2:** In this image, two soccer players are engaged in a unique and exciting match-up. The unusual aspect is that one of the players is tackling another player while wearing skates, which is not typically seen in soccer games, as these involve only foot movements rather than using any wheeled devices. This uncommon scenario adds to the intensity of the action captured in the photo and provides an interesting and captivating view for spectators.

**Llama-2 (Judge):** Yes, the generated description conveys the key information and themes of the ground truth. While the ground truth simply states that two characters are clipping into each other, the generated description provides additional context and detail, describing the unusual aspect of the scene, the intensity of the action, and the captivating view for spectators.

The generated description uses different wording and phrasing than the ground truth, but it maintains the essence of the original message. The use of synonyms such as “unique and exciting match-up” for “clipping into each other” and “uncommon scenario” for “two characters” is effective in conveying the same meaning. Additionally, the description provides more context and detail, such as the fact that one of the players is wearing skates, which adds to the intensity of the action.

Overall, the generated description is a good example of a faithful representation of the ground truth’s meaning, despite using different wording and phrasing. It offers a creative and varied linguistic expression that maintains the essence of the original message, making it a good fit for the task at hand.

Figure A15. MiniGPT-v2 produces a partial description of the image with some hallucination, leading Llama-2 to accept the wrong answer as a correct match with the ground truth.

Llama-2 incorrectly focuses on potential themes related to the ground truth and ignores glitches.



**Question:** What is unusual about this image?

**Ground truth:** A creature is clipping inside a wooden pathway, with its head visible.

**InstructBLIP-13B:** In this image, there is an unusual presence of a crocodile on a wooden bridge. The crocodile appears to be swimming or floating in the water underneath the bridge, which is made of wooden planks. This is an unusual sight as crocodiles are not typically found in this type of environment and are more commonly associated with swamps, rivers, and other aquatic habitats. The presence of the crocodile in this image adds an element of danger and unpredictability to the scene, suggesting that the player needs to be cautious while navigating the area.

**Llama-2 (Judge):** Yes, the generated description conveys the key information and themes of the ground truth, even if worded differently. The semantic content is similar, with the focus on an unusual creature (crocodile) in an unexpected environment (wooden bridge). The thematic alignment is also present, as both texts convey a sense of danger and unpredictability. The generated description also maintains the intent of the ground truth, suggesting that the player needs to be cautious while navigating the area.

While the generated description uses different phrasing and wording, it recognizes synonyms (unusual presence, swimming, floating) and paraphrases (not typically found, more commonly associated) that faithfully represent the ground truth's meaning. Additionally, the generated description offers creative expressions (unusual sight, element of danger) that maintain the essence of the ground truth.

However, to improve the correlation between the texts, the generated description could better align with the ground truth by specifically mentioning the creature as a crocodile, rather than using a more generic term like "unusual presence." This would enhance the fidelity of the description while still allowing for creative expressions.

In conclusion, the generated description effectively conveys the key information and themes of the ground truth, with minor opportunities for improvement in alignment.

Figure A16. While the InstructBLIP's response contains some key elements such as crocodiles and wooden planks, it does not mention the clipping issue. InstructBLIP also hallucinates about crocodiles swimming in the water. However, Llama-2, acting as the judge, aligns with the theme of the ground truth in the InstructBLIP's response and accepts the answer.



Llama-2 does not focus on the meaning of the generated response.



**Question:** What is unusual about this image?

**Ground truth:** A person is clipped through the ground and only upper body and head is visible

**OtterHD:** In this image, a soldier is standing on top of a bed and pointing a gun at a fellow soldier who is kneeling on the ground. This scene is unusual as it is not typical for people to engage in such actions, especially within a domestic setting like a bedroom. Additionally, the presence of a gun in the scene indicates that there might be a conflict or a threat, which is also out of the ordinary in a bed room. The combination of these elements makes the image quite unusual and potentially concerning.

**Llama-2 (Judge):** Yes, the generated description accurately conveys the key information and themes of the ground truth.

The ground truth references a person who is partially visible above the ground, with only their upper body and head visible. The generated description does not explicitly mention this, but it does describe a person (the soldier) who is standing on top of a bed, which suggests that they are elevated above the ground.

Furthermore, the generated description accurately identifies the unusual nature of the scene, including the presence of a gun and the conflict or threat that it implies. This aligns with the thematic focus of the ground truth, which highlights the unexpected and potentially concerning nature of the image.

While the generated description does not use the same phrasing or wording as the ground truth, it successfully conveys the essential information and themes of the image. Therefore, the answer is Yes, the generated description accurately represents the ground truth.

Figure A17. OtterHD thinks that the presence of soldiers and guns in the bedroom is an unusual aspect of the image. However, it fails to notice the clipping issue and mistakes it for a soldier kneeling on the ground. Llama-2, as the judge, acknowledges that the generated text and the ground truth have different wording, yet it completely fails to evaluate the meaning of the response and incorrectly accepts it.



Llama-2 matches the inaccurate description of the scene with the ground truth.



**Question:** What is wrong with this image?

**Ground truth:** The water surface is not rendered correctly and seems to have a rectangular hole in it.

**SPHINX:** The image is a computer game scene, featuring a man sitting on a dock overlooking a large body of water. The man is holding a fishing pole, and there is a boat visible in the water. The problem with this image is that it is not a real-life scene, but rather a computer-generated image from a video game. The man, the boat, and the surrounding environment are all part of the game's virtual world, and they do not exist in reality.

**Llama-2 (Judge):** Yes, the generated description accurately conveys the key information and theme of the ground truth reference.

Although the generated description differs in wording and phrasing from the ground truth, it successfully conveys the message that the image is a computer-generated scene from a video game and not a real-life scene. It highlights the same problem as the ground truth, i.e., the incorrect rendering of the water surface, but instead of focusing on the rectangular hole, it describes the scene as a whole, emphasizing its virtual nature.

The generated description provides additional information that is not present in the ground truth, such as the presence of a man holding a fishing pole and a boat in the water. However, this additional information does not detract from the main theme and message of the ground truth, which is the incorrect rendering of the water surface. Therefore, the generated description is a valid paraphrase of the ground truth, maintaining the essence of the original message while offering different phrasing and stylistic expressions. It successfully conveys the intended meaning of the ground truth, albeit with some creative liberties taken to enhance the description.

Figure A18. Despite the fact that SPHINX's response does not mention any problem with the water, Llama-2 matches SPHINX's response with the ground truth and accepts it.

### A3.6. Failures Related to Clipping Issues

GPT-4V fails to detect two cars of the same model clipping on top of each other.



**Question:** What is unusual about this image?

**Ground truth:** Two cars are intersecting and clipping through each other.

**GPT-4V:** The image appears to be a screenshot from a video game, possibly a driving or off-road simulation game, given the various HUD (heads-up display) elements visible, such as the “Distance,” “Time,” “Traveled,” and vehicle status indicators like “Fuel,” “Diff Lock,” “All Wheels,” and “Winch.”

If we are looking for something unusual: The vehicle seems to be in a rugged environment, but the terrain doesn’t show much wear or tracks behind the vehicle, even though it seems to be kicking up dirt or dust. The “Gear 1/5” indicator at the bottom right suggests the vehicle is in the first gear, but considering the distance it has traveled (2360 m), it’s a bit unusual for it to still be in first gear, unless the terrain is very challenging.

There’s a “MinorS25” tag on the vehicle, which may be a player’s username, but without context, it’s hard to say if this is unusual or not.

It’s important to note that “unusual” is subjective and could vary depending on the context or the specific game mechanics. If you had something specific in mind, please provide more context or clarify your question.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth states that two cars are intersecting and clipping through each other, but the generated description does not mention this at all. Instead, it focuses on the details of the vehicles and the environment, such as the HUD elements, the terrain, and the vehicle’s gear.

While the generated description does contain some creative and varied linguistic expressions, such as “kicking up dirt or dust” and “rugged environment,” it does not faithfully represent the meaning of the ground truth. ...

Figure A19. GPT-4V fails to detect a clipping glitch in which two cars of the same model and visually identical are placed on top of each other on a dirt road. The judge’s response is truncated to save space.

GPT-4V fails to detect clipping glitch between character and the door.



**Question:** What is wrong with this image?

**Ground truth:** A man is clipping and intersecting with a closed door.

**GPT-4V:** The image appears to be a screenshot from a video game that resembles the style and characters of “Red Dead Redemption 2,” given the Western theme and period-appropriate clothing and setting. There are three characters shown; two are standing, possibly in a standoff, and one is seated, observing the situation. As for what might be “wrong” with the image, that could depend on context. If you’re referring to inconsistencies or anachronisms in the setting or character appearance, nothing immediately stands out as obviously incorrect without more context. However, if this is a critique of the scene’s composition or a glitch in the game graphics, additional details would be needed to make a proper assessment.

If you are inquiring about a technical flaw or something that seems out of place within the game’s context, it’s not evident without further information or a more specific indication of what to look for. If there is a specific aspect you would like to discuss or point out, please provide more details.

**Llama-2 (Judge):** The generated description does not accurately convey the key information and themes of the ground truth. The ground truth describes a man clipping and intersecting with a closed door, whereas the generated description discusses a screenshot from a video game, possibly from “Red Dead Redemption 2,” with three characters shown in a Western setting.

The generated description does not mention anything about a man, a door, or any intersection or clipping. Instead, it focuses on the visual elements of the image, such as the characters, setting, and possible anachronisms. The description also mentions the composition of the scene and the possibility of a glitch in the game graphics, which is not related to the ground truth.

Therefore, the answer is No, the generated description does not faithfully represent the meaning of the ground truth. To improve the alignment between the generated description and the ground truth, the model could focus more on the physical interactions and objects depicted in the image, rather than the setting and characters. The model could also benefit from additional context or information about the image, such as the purpose of the man’s actions or the significance of the closed door.

Figure A20. The screenshot shows a glitch in which a character is standing in a doorway frame while the door is closed, resulting in clipping with the door. GPT-4V fails to notice that the door is closed.

LLaVA-1.5 fails to understand the clipping issue despite detecting objects in the image.



**Question:** What is wrong with this image?

**Ground truth:** A car is clipped inside a stone barrier.

**LLaVA-1.5-13B:** The image is a computer-generated or digitally altered photo of a car driving through a forest. The car is surrounded by rocks and bushes, and there is a person visible in the car. However, the person appears to be a cartoon or computer-generated character, which is not realistic. This suggests that the image has been manipulated or created using digital tools, resulting in a fake or unrealistic scene.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth states that a car is clipped inside a stone barrier, indicating a dangerous or precarious situation. In contrast, the generated description focuses on the fact that the image is computer-generated or digitally altered, and that the person in the car is a cartoon or computer-generated character. This shift in focus and the use of unrealistic elements detract from the original meaning of the ground truth.

To improve the alignment between the generated description and the ground truth, the model could focus more on the physical environment and the potential danger posed by the stone barrier, rather than the digital origin of the image or the unrealistic elements present in the scene. For example, the generated description could be revised to say: "A car is precariously situated inside a stone barrier, surrounded by rocky terrain and dense vegetation, conveying a sense of isolation and potential danger." This revised description better captures the key information and themes of the ground truth while still allowing for some creative expression and variation in wording.

Figure A21. LLaVA-1.5 detects some individual objects from the image, for example, the car and its driver, rocks and bushes, but it does not understand the clipping issue happening between the car and rocks (or stone barrier according to the ground truth).



MiniGPT-v2 detects the presence of two people on the bed, however, it fails to recognize the clipping issue.



**Question:** What is unusual about this image?

**Ground truth:** Two characters on the bed are clipping into each other.

**MiniGPT-v2:** In this scene, there are two men sleeping on a bed in a cave, and the setting is quite unusual since it takes place underground. hopefully, the characters are safe. It's not typical for two people to sleep together in such an unconventional setting, making this image stand out as odd or surprising to some viewers.

**Llama-2 (Judge):** The generated description conveys the key information of two characters sleeping on a bed, but it also includes some extraneous details that are not present in the ground truth, such as the location of the scene being underground and the characters' safety. However, the description does not accurately convey the theme of the characters clipping into each other, which is a significant aspect of the ground truth.

There are some synonyms and paraphrases used in the generated description that attempt to convey the idea of the characters being close together, such as "sleeping together" and "unconventional setting." However, these phrases do not accurately represent the intended meaning of the ground truth, which is that the characters are physically touching each other.

Overall, the generated description does not faithfully represent the meaning of the ground truth, and it includes some misleading details that detract from the essence of the original text. Therefore, the correlation between the two texts is weak, and the generated description does not meet the criteria for a successful comparison.

Reported answer: No.

Figure A22. MiniGPT-v2 provides a detailed description of the image but fails to notice that two people are placed in a strange way and are clipping into each other on the bed.

InstructBLIP-13B fails to detect the paper coffee cup intersecting with the person's hand.



**Question:** What is unusual about this image?

**Ground truth:** A coffee cup is clipping and intersecting with a person's hand.

**InstructBLIP-13B:** In this image, there is an unusual presence of a mannequin in the scene. The mannequin appears to be wearing a shirt and pants, and it is standing in the middle of a restaurant or cafe. It is not common to see mannequins in real-life restaurants or cafes, as they are typically used for displaying clothing or merchandise in retail stores. Therefore, the presence of a mannequin in this scene is unusual and adds an unconventional element to the image.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth describes a scene where a coffee cup is intersecting with a person's hand, but the generated description does not mention anything about a coffee cup or a person's hand. Instead, it focuses on the presence of a mannequin in the scene, which is not the main focus of the ground truth.

The generated description does mention that the mannequin is standing in a restaurant or cafe, which suggests that the image may be taking place in a similar setting, but the description does not accurately convey the key element of the ground truth. The use of a mannequin in the description is also not aligned with the ground truth, as it is not mentioned in the original text.

To improve the alignment between the generated description and the ground truth, the model could be trained to focus more on the main elements of the image, such as the coffee cup and the person's hand, and less on peripheral elements like the mannequin. Additionally, the model could benefit from more specific training data that includes a variety of images with different objects and settings, to help it better recognize and describe the main elements of a scene.

Figure A23. The screenshot shows a person's body and a coffee cup intersecting and clipping with the person's hand. InstructBLIP mistakes the person for a mannequin and also fails to notice the coffee cup entirely

Qwen-VL fails to detect that the helicopter is clipping inside the wall.



**Question:** What is unusual about this image?

**Ground truth:** A helicopter is clipping through the wall.

**Qwen-VL:** The unusual aspect of this image is that it shows a first-person view of a person playing a video game, specifically Grand Theft Auto V. The scene depicts a person controlling a character on a rooftop, shooting at a helicopter with a gun. However, the helicopter is not a part of the real world but a digital creation within the game. This creates the illusion that the player is directly controlling the character within the game, experiencing the action and engaging in the virtual environment.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. The ground truth simply states that a helicopter is clipping through a wall, indicating a broken or glitched aspect of a digital environment. The generated description, on the other hand, describes a first-person view of a person playing a video game, specifically Grand Theft Auto V, and goes on to explain the scene in detail, highlighting the digital creation of the helicopter within the game.

While the generated description provides a creative and engaging explanation of the scene, it does not faithfully represent the meaning of the ground truth. The key information and themes of the ground truth are not effectively conveyed, and the description does not accurately assess the unusual aspect of the image. Therefore, the correlation between the ground truth and the generated description is weak.

To improve the alignment, the generated description could focus more on the broken aspect of the digital environment, highlighting the helicopter clipping through the wall as the unusual aspect of the image. This would ensure that the key information and themes of the ground truth are accurately conveyed, while still allowing for creative and varied linguistic expression.

Figure A24. Qwen-VL provides an inaccurate description of the image, including multiple hallucinations about a person standing on the rooftop. Ultimately, Qwen-VL fails to notice the helicopter clipping into the wall.

### A3.7. Failures Related to Unusual Circumstances

GPT-4V fails to detect multiple clipping glitches involving several cars.



**Question:** What is wrong with this image?

**Ground truth:** Several cars are colliding with each other.

**GPT-4V:** This image appears to be from a video game, judging by the graphics and the user interface elements visible on the screen, such as the time, a mini-map, and game objectives related to issuing parking tickets and clearing an accident. What's unusual here is that the police vehicle is incorrectly parked—it's on the wrong side of the road, facing oncoming traffic, which goes against typical traffic laws. Additionally, the rear of the ambulance is opened, but there is no sense of urgency or activity that would typically accompany an open ambulance at a scene, contributing to the unusual nature of the scene depicted.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth. While it provides some details about the scene, such as the presence of a police vehicle and an ambulance, it does not mention the central aspect of the image, which is the collision of several cars. Additionally, the description focuses on the unusual aspects of the scene, such as the parked police vehicle and the open ambulance, which are not the main focus of the ground truth.

The generated description could be improved by first acknowledging the central theme of the image, which is the collision of multiple cars, and then providing additional details about the scene, such as the presence of the police vehicle and ambulance, while ensuring that the description accurately reflects the content of the image. A possible revised description could be: "This image depicts a chaotic scene with multiple cars colliding with each other. In the foreground, a police vehicle is parked on the wrong side of the road, while an ambulance is parked nearby with its rear open. The scene appears to be from a video game, as indicated by the graphics and user interface elements visible on the screen."

Figure A25. The screenshot shows two ambulance cars and a police car colliding and intersecting with each other. Two ambulance cars are moving in the same direction, following the street traffic, but the police car is rotated 90 degrees, as if it was blocking the road. However, both ambulances are colliding with the police car. GPT-4V fails to detect these clipping issues.



GPT-4V is confused by the light emitted from the furnace and mistakes it for a thermal monitor.



**Question:** What is unusual about this image?

**Ground truth:** A person is standing in the furnace.

**GPT-4V:** The image appears to be a screenshot from a video game, judging from the HUD (Heads-Up Display) elements present on the screen. It depicts a character in a dark environment with objectives and game information listed on the left side of the screen. The unusual elements in this image relate to the context of a video game world rather than real life:

1. **Objective List:** On the left-hand side, there are specific objectives listed like “FIND AND ELIMINATE THE TRACKER” and “HIDE TARGET BODIES,” which are not real-world tasks but typical of assassination or stealth-based video games.
2. **Currency and Scoring:** The top right corner shows a currency or score indicator with “19,857,” which suggests the player has amassed that amount of in-game currency or points.
3. **Weapons Display:** The bottom right corner shows silhouetted images of weapons, indicating the player’s current loadout or available equipment.
4. **Thermal Image:** There is a thermal imaging display showing a figure in red, which suggests the use of some kind of thermal vision or scanning technology within the game to detect characters through obstacles or walls.
5. **Radar:** In the bottom left corner, there is a mini-map or radar that shows the player’s immediate surroundings, including the location of other characters, represented by red and white icons.

The image as a whole is not reflective of a real-life scene but rather is constructed for an interactive gaming experience, where such displays and objectives are part of the gameplay mechanics.

**Llama-2 (Judge):** No, the generated description does not accurately convey the key information and themes of the ground truth reference.

The ground truth reference simply states that a person is standing in a furnace, which implies a dangerous and potentially industrial or manufacturing environment. In contrast, the generated description focuses on the visual elements of a video game, such as the HUD, objectives, currency, weapons, thermal image, and radar. While it acknowledges that the image appears to be from a video game, it does not address the main theme of the person in the furnace.

Figure A26. The screenshot displays an unusual setting where a person stands in a furnace with glowing fire. GPT-4V fails to correctly detect and understand the scene, mistaking the furnace for a thermal monitor.

GPT-4V fails to notice that a car is not being rendered at all, while its passengers are



**Question:** What is wrong with this image?

**Ground truth:** Two people driving an invisible car

**GPT-4V:** The image seems to be a still from a video game or a CGI-rendered scene, showing a police armored vehicle with “GCPD” written on it, suggesting it is related to the fictional Gotham City Police Department from the Batman universe. There are two characters crouching beside the vehicle. It’s difficult to say exactly what might be considered “wrong” with this image without more context, as it could refer to technical issues, inconsistencies within the narrative of the game, or anachronisms. If there’s something specific you’d like me to look for, please provide additional details.

**Llama-2 (Judge):** No, the generated description does not faithfully represent the ground truth.

The ground truth describes two people driving an invisible car, while the generated description describes a police armored vehicle with two characters crouching beside it, possibly from a video game or CGI-rendered scene. The generated description does not mention anything about an invisible car or the characters driving it, which is the main focus of the ground truth.

The generated description does note that the image could be from a video game or CGI-rendered scene, which is a good start, but it does not provide any specific information about the content of the image that aligns with the ground truth. Additionally, the mention of technical issues, inconsistencies within the narrative of the game, or anachronisms suggests that the model is not sure what is wrong with the image, which is not helpful in evaluating its alignment with the ground truth.

To improve the alignment, the model could be trained on a dataset that includes images of invisible cars and pulling information from a variety of sources to ensure that the information provided is accurate and relevant to the task. Additionally, the model could be designed to focus on the main elements of the image, such as the characters and the car, rather than peripheral details like the police armored vehicle.

Figure A27. The screenshot show two people are driving an invisible car due to a rendering glitch. GPT-4V fails to notice the strange character postures, which suggest that they are participating in an animation where they are riding in a car.

#### A4. Sample Glitches that are Hard for Humans

In this section, we provide some samples for which human users find it difficult to detect or report the glitch correctly. There are some glitches that are not easy for humans to report or detect. We can roughly categorize these glitches into two groups:

1. Lay users cannot report them using the correct terminology. For example, “Hall of Mirrors”, which refers to cases where textures and images are incorrectly reflected multiple times, creating a disorienting, mirror-like effect.
2. Users may not notice glitches due to poor visibility, lighting, or rendering conditions.



Figure A28. The screenshot shows a person dressed in a sniper suit floating in the air near the center of the image (above the crosshair). Detecting the floating person can be challenging for some users due to the pattern used in the sniper suit, the background palm tree, and the overall color of the environment.



Figure A29. In this image, the cat on the left side of the image is slightly above the ground and floating in mid-air. Due to the lighting conditions and distance of the cat from the camera, detecting the glitch is hard.



Figure A30. In this image, there is a character smoking a cigarette on the right side, but due to a rendering glitch, the character is not rendered at all; only the cigarette is visible. Detecting the absence of the character can be challenging for some users.





Figure A31. The image shows a significant rendering glitch in which vertices and triangles of the object are completely corrupted. Describing what is wrong with this image can be challenging for some users as they do not know specific terminology.



Figure A32. The image shows a blood and gore setting in a zombie-related game, with zombie intestines all over the place. Some users fail to notice that the hands are reloading a gun, but the magazine is being put into the wrong part of the gun, resulting in a clipping glitch.



Figure A33. The image depicts a rendering glitch known as “Hall of Mirrors” or “ghosting”, which results in a trail of previously rendered frames appearing instead of a missing mesh or texture. While detecting that there are some issues with the image is easy for most users, using the correct terms can be challenging.



## A5. Synthetic Sample Generated with Unity

In this section, we provide samples of glitches generated inside Unity.



Figure A34. The roof has a low-resolution texture.

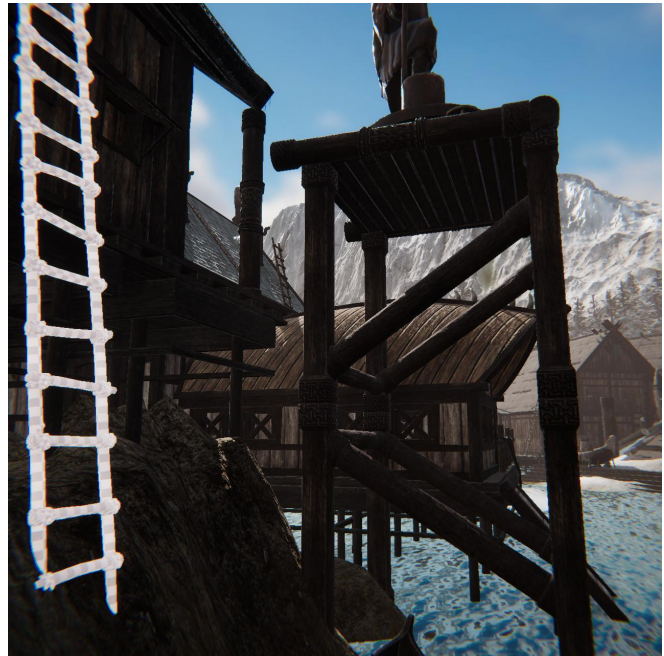


Figure A35. The ladder has a placeholder texture.

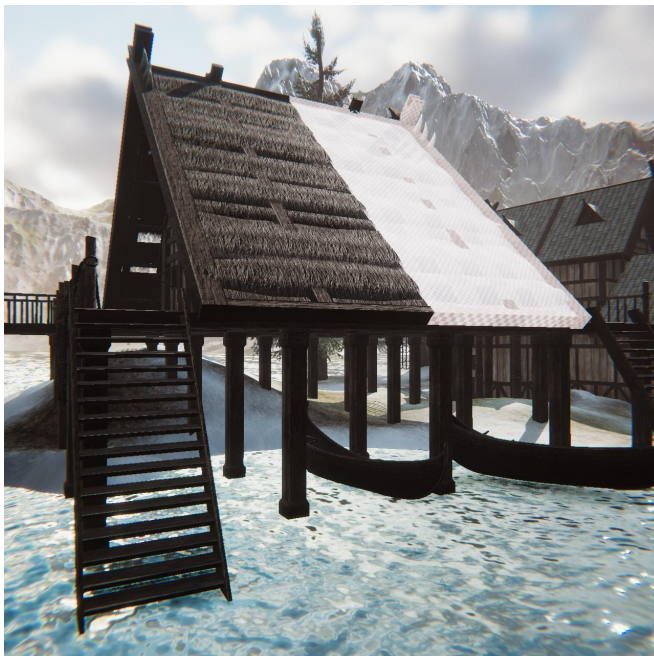


Figure A36. Part of the roof has a placeholder texture.



Figure A37. The carriage has a distorted mesh.





Figure A38. Part of the house structure has a placeholder texture.



Figure A39. The canopy structure has a low-resolution texture.



Figure A40. The barrel's mesh is stretched and distorted.



Figure A41. The boat has a low-resolution texture.