LONGHALQA: LONG-CONTEXT HALLUCINATION EVALUATION FOR MULTIMODAL LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

Abstract

Hallucination, a phenomenon where multimodal large language models (MLLMs) tend to generate textual responses that are plausible but unaligned with the image, has become one major hurdle in various MLLM-related applications. Several benchmarks have been created to gauge the hallucination levels of MLLMs, by either raising discriminative questions about the existence of objects or introducing LLM evaluators to score the generated text from MLLMs. However, the discriminative data largely involve simple questions that are not aligned with real-world text, while the generative data involve LLM evaluators that are computationally intensive and unstable due to their inherent randomness. We propose LongHalQA, an LLM-free hallucination benchmark that comprises 6K long and complex hallucination text. LongHalQA is featured by GPT4V-generated hallucinatory data that are well aligned with real-world scenarios, including object/image descriptions and multi-round conversations with 14/130 words and 189 words, respectively, on average. It introduces two new tasks, hallucination discrimination and hallucination completion, unifying both discriminative and generative evaluations in a single multiple-choice-question form and leading to more reliable and efficient evaluations without the need for LLM evaluators. Further, we propose an advanced pipeline that greatly facilitates the construction of future hallucination benchmarks with long and complex questions and descriptions. Extensive experiments over multiple recent MLLMs reveal various new challenges when they are handling hallucinations with long and complex textual data.

006

008 009 010

011

013

014

015

016

017

018

019

021

024

025

026

027

028

029

1 INTRODUCTION

Multi-modal Large Language Models (MLLMs) (Dai et al., 2024; Hu et al., 2024; ?; Bai et al., 2023; Liu et al., 2023b; 2024a; Zhu et al., 2023) have achieved great progress in understanding 037 multi-modal contents, by generating detailed descriptions of images, conducting sophisticated, consecutive conversations with humans, etc. Despite the remarkable advancements, MLLMs often experience severe hallucination problems (Yin et al., 2023; Leng et al., 2023; Huang et al., 2023; Zhu 040 et al., 2024; Yue et al., 2024; Bai et al., 2024) by generating textual responses that are not aligned 041 with the corresponding image contents. While hallucinations significantly compromise MLLMs' 042 reliability and applicability in various vision-language tasks and applications, effective and efficient 043 measurement of the hallucination level of MLLMs has become a prerequisite for diagnosis and 044 mitigation of hallucination in MLLMs.

Several related benchmarks (Yifan et al., 2023; Qiu et al., 2024; Jiang et al., 2024; Liu et al., 2024b; Lovenia et al., 2023; Wang et al., 2023; 2024a) have been proposed to gauge the hallucination level of MLLMs in two representative approaches. The first approach conducts discriminative evaluations, where MLLMs are queried with simple questions about whether some objects exist in the image, as illustrated in the upper part of Fig. 1. The second approach conducts generative evaluations, which first apply MLLMs to describe the image and then adopt LLM evaluators to examine whether MLLMs generate hallucinatory content. However, most existing benchmarks share several constraints: 1)
Most discriminative benchmarks merely require a yes-or-no answer, which is often too simple to tell much on the cause of hallucinations. 2) Discriminative benchmarks usually come with very short questions like "Is there an {object} in the image?" which are oversimplified and insufficient for

087

880

089

091



Figure 1: LongHalQA is featured with two novel tasks, namely, *Hallucination Discrimination* and *Hallucination Completion*, which unify both discriminative and generative evaluations into the same multiple-choice-question form without requiring costly LLM evaluations. It comprises three types of long-context data, including *Object-level Description*, *Image-level Description*, and *Multi-round Conversation*. Compared with short and simple questions in existing benchmarks like "Is there an {object} in the image?", the three types of data are more open-ended, richer in contextual information, and closer to real-world data. White circle in image emphasizes the hallucination of passengers.

examining hallucination in sophisticated real-world scenarios. 3) Both discriminative and generative
benchmarks (Yifan et al., 2023; Qiu et al., 2024; Kaul et al., 2024) usually leverage off-the-shelf object
annotations to construct questions or detect hallucinatory objects, leading to limited variability (e.g.,
a fixed set of 80 object categories for COCO) and biased evaluations toward a small set of objects. 4)
Generative benchmarks (Kaul et al., 2024; Jiang et al., 2024; Qiu et al., 2024; Sun et al., 2023; Liu
et al., 2023a) generally employ LLMs for hallucination evaluations, but LLMs are computationally
intensive and often unstable due to their inherent randomness.

098 We design LongHalQA, an LLM-free hallucination benchmark that comprises 6K long and contextrich hallucination questions. LongHalQA is built from GPT4V-generated hallucinatory data that is 100 well aligned with various real-world scenarios. It features two multiple-choice-question (MCQ) tasks, 101 namely, hallucination discrimination and hallucination completion as illustrated at the bottom of 102 Fig. 1. Specifically, hallucination discrimination requires MLLMs to determine whether the given 103 text contains hallucinations and pick the *right causes* of the hallucinations. Hallucination completion 104 instead transforms generative evaluations into a discriminative task, asking MLLMs to continue 105 the text and pick the right option that does not contain hallucinations. LongHalQA thus unifies discriminative and generative evaluations into the same MCQ form, assessing MLLMs' understanding 106 of hallucinations and their tendency to generate hallucinations concurrently. LongHalQA queries 107 come in three data formats as illustrated in the middle of Fig. 1, including object-level descriptions,

image-level descriptions, and multi-round conversation, which are much longer and complex, covering
a wide range of 12 types of hallucinations. Such long and complex questions allow LongHalQA to
gauge the hallucination levels of MLLMs in more practical applications and scenarios. Compared
to existing generative benchmarks, we demonstrate that our MCQ hallucination completion task
exhibits similar trends to free-form generative evaluation. Additionally, LonghalQA achieves much
higher speed in evaluating MLLMs generative hallucination, especially for extremely large models,
facilitating the expansion of evaluation samples and the need for fast testing and evolution of MLLMs.

Additionally, we propose LongHallGen, an automated pipeline for Long-context Hallucination Data Generation. LongHallGen is featured with a set of prompt templates for GPT4v that allow generating hallucination data and converting the generated data into multiple-choice-questions automatically. By modifying the prompt sets, LongHallGen can adjust the type of generated hallucinations, content topics, and data formats. We believe that LongHallGen will serve as a strong basis while creating new or expanding existing hallucination datasets for evaluating and training MLLMs in future research.

121 Based on LongHalQA, we evaluate ten mainstream MLLMs on long-context hallucinations and 122 provide a comprehensive analysis. The evaluations reveal constraints of MLLMs in discerning 123 and explaining hallucination in long texts, as well as in generating hallucinatory content when completing long texts. Additionally, we observe that the Chain-Of-Thought (COT), a simple but 124 125 effective hallucination mitigation method (Jiang et al., 2024; Qian et al., 2024), is effective for short queries and generative hallucinations but degrades the performance of most MLLMs on long-context 126 hallucinations discrimination in LongHalQA, especially for those with small sizes, This suggests that 127 COT may be limited by MLLMs' capability on long context processing. We believe that LongHalQA 128 will serve as a basis for mitigating long-context hallucinations in various real-world MLLM tasks. 129

131 2 RELATED WORKS

Hallucination Benchmarks for MLLMs. Various benchmarks have been proposed to measure 133 the hallucination level of MLLMs, including both discriminative and generative benchmarks. For 134 discriminative benchmarks, POPE (Yifan et al., 2023), CIEM (Hu et al., 2023), AMBER (Wang 135 et al., 2023), NOPE (Lovenia et al., 2023), and MME (Fu et al., 2023) query MLLMs with simple 136 questions about the existence or attributes of specific objects in images. PhD (Liu et al., 2024b) 137 and Hal-Eval (Jiang et al., 2024) introduce more types of intrinsic hallucinations into evaluations, 138 such as multi-modal conflicting or event hallucinations. Most of these benchmarks feature simple 139 and short questions and seek solely simple binary "yes" or "no" answers. Generative evaluations 140 introduce LLM evaluators to analyze hallucinations in MLLM-generated text. Various generative 141 benchmarks (Jiang et al., 2024; Kaul et al., 2024; Qiu et al., 2024; Wang et al., 2023) are proposed to 142 improve the scope of evaluated hallucinations and the efficiency and accuracy of LLM evaluators. 143 Our LonghalQA takes a different perspective on long-context hallucinations and encompasses both 144 discriminative and generative evaluations in a unified and efficient MCQ format.

145

130

132

146 147 148

149

150

151

3 LONGHALQA: LONG-CONTEXT HALLUCINATION BENCHMARK

LongHalQA comprises 6485 multiple-choice questions (MCQ) that cover two tasks with long-context text: hallucination discrimination and hallucination completion. This section presents the task format in Sec. 3.1, the data format and distributions in Sec. 3.2, and the evaluation metrics in Sec. 3.3.

152 3.1 TASK FORMAT

In order to comprehensively evaluate the hallucination level of MLLMs, we introduce two tasks,
namely, hallucination discrimination and hallucination completion, which conduct discriminative and
generative MLLM evaluations, respectively.

Hallucination Discrimination. For discriminative evaluations, we propose a set of multiple-choice
questions to query MLLMs whether object/image descriptions or text responses in a conversation
match the contents of images as illustrated in Fig. 1. Each hallucination question is equipped with
multiple answer choices and corresponding explanations. One of the choices starts with "yes,"
suggesting that the text matches the image contents, while the other three start with "no," followed
by explanations. MLLMs are required not only to identify the presence of hallucination but also to

Table 1: Statistics of 12 types of hallucinations in LongHalQA. "Object", "Description", and 163 "Conversation" denote the data formats of object-level descriptions, image-level descriptions, and 164 multi-round conversations, respectively. We use both hallucinatory and non-hallucinatory data to 165 construct the discrimination task, while only hallucinatory data for the completion task. 166

| Halluc | ination Types | Object | Description | Conversation | Total |
|--------|---------------------------------|-------------|-------------|--------------|--------------|
| H1 | (Non) Existent Objects | 234 | 261 | 323 | 818 |
| H2 | Object Attributes | 89 | 130 | 175 | 394 |
| H3 | Object Color | 122 | 90 | 86 | 298 |
| H4 | Object States | 50 | 67 | 84 | 201 |
| H5 | Number of Objects | 80 | 92 | 134 | 306 |
| H6 | Object Locations | 45 | 76 | 86 | 207 |
| H7 | Object Relationships | 49 | 54 | 49 | 152 |
| H8 | Text / Sign Meaning | 27 | 61 | 91 | 179 |
| H9 | Environment Description | 10 | 13 | 31 | 54 |
| H10 | Background Description | 13 | 14 | 21 | 48 |
| H11 | Time | 2 | 7 | 5 | 14 |
| H12 | Weather | 2 | 5 | 13 | 20 |
| Halluc | inatory Data | 723 (52.7%) | 869 (63.3%) | 1098 (68.5%) | 2690 (61.9%) |
| Non-ha | allucinatory Data | 647 (47.2%) | 503 (36.7%) | 506 (31.5%) | 1656 (38.1%) |
| Data f | or Hallucination Discrimination | 1370 | 1372 | 1604 | 4346 |
| Data f | or Hallucination Completion | - | 869 | 1270 | 2139 |
| Total | | 1370 | 2241 | 2874 | 6485 |

- 18
- 185 186

187

188

understand why hallucination happens and choose the correct explanation. LongHalQA comes with 4346 such image-question samples for the hallucination discrimination task.

Hallucination Completion. Previous generative benchmarks require MLLMs to generate descrip-189 tions for the image and employ LLM evaluators to score descriptions. To obviate the slow generation 190 process and costly LLM evaluations, we transform generative evaluations into MCQ format, as shown 191 in Fig. 1. Specifically, we provide an image and a related incomplete description or conversation and 192 ask MLLMs to continue the text. Four answer choices of possible completing sentences are provided, 193 with one correct choice and three hallucinatory choices. Compared to generative benchmarks based 194 on LLM evaluators, the format of generative MCQ significantly reduces evaluation costs and allows 195 for more detailed annotation and analysis of hallucination data. LongHalQA comes with 2149 196 samples for the hallucination completion task.

197 The MCQ hallucination completion task simulates the sampling process in MLLM inference, where 198 MLLMs first generate several potential outputs and then select the most appropriate one free from 199 hallucinations. Furthermore, the hallucination completion task can be adapted into a free-form 200 continuation task, where MLLMs are prompted to freely continue long-context data. Our experiments 201 demonstrate that the MCQ format of the hallucination completion task and the free-form generation 202 format yield similar trends in evaluating generative hallucinations of MLLMs.

203 204 205

206

3.2 DATA FORMAT AND DISTRIBUTION

This section presents the format of long-context hallucination data from two aspects, namely, data 207 formats and types of hallucinations, more details to be elaborated in the ensuing subsections. 208

209 Data Formats. As shown in Tab. 1, LongHalQA consists of three formats of image-text hallucinatory 210 data, including 1370 Object-level Description, 1372 Image-level Description, and 1604 Multi-round 211 Conversation. Specifically, Object-level Description describes a specific object in the image, such as 212 its attributes, states, or relations with other objects. Image-level Description covers the main contents 213 and more details of an image, such as objects, background, weather, etc, in one paragraph. For the *Multi-round Conversation*, we simulate a human user who communicates with an assistant, querying 214 the image content. The three types of data formats are highly compatible with the actual application 215 scenarios of MLLMs and thus can better simulate real hallucination situations.

216 Continue the following conversation: 217 User: Can you describe what's on the table in this image? 218 Assistant: The table is set for a meal with blue and yellow patterned plates. 219 A. There are four such plates in total, and each has a different arrangement of food. 220 B. There are four such plates in total, and they are all empty. C. There are five such plates for salads placed around the table. 222 D. There are five plates in total, and they all have various types of food. 224 Does the following description match the image content? 225 "The image displays an organized wardrobe with various compartments containing 226 clothes and accessories. On the left section, there are folded items on shelves 227 including white and crimson towels. The central part shows two levels of 228 hanging space: the upper rack holds beige shirts, while the lower section displays similar items. The right section contains a series of lighter clothes, predominantly 229 white shirts, arranged in two rows, hanging neatly above each other." 230 231 232

233 234 Figure 2: LongHalQA involves complex hallucination annotations involving logic and textual consistency, which are closer to hallucinations in real-world MLLM application scenarios.

Types of Hallucination. As shown in Tab. 1, We analyze the description texts of objects/images
 and categorize the wrong descriptions into twelve types of hallucinations. The detailed definitions
 and visualization of each type of hallucination are provided in Appendix A. Compared to existing
 benchmarks that focus on the existence and attributes of objects, LongHalQA contains a much broader
 collection of hallucinations with detailed annotations.

In addition, LongHalQA includes hallucination samples involving logic and contextual consistency,
which are untouched in most existing benchmarks. Such complex hallucinations are often observed
with contextually inconsistent descriptions such as 'four such plates' or 'five plates' as illustrated in
the upper part of Fig. 2, or incorrectly mixed descriptions such as 'two rows of shirts in the central
part' vs 'one row of shirts in the right part' as illustrated in the lower part of Fig. 2. We foresee that
LongHalQA will inspire more in-depth studies of MLLMs regarding such complex hallucinations.

246 Complexity and Length of Text in LongHalQA. We derive certain statistics of LongHalQA data 247 to verify their quality, including the length and the number of object nouns appearing in the text 248 data. The study shows LongHalQA has an average of 14/130 and 189 words, respectively, for 249 object/image-level descriptions and multi-round conversations, clearly longer and more informative 250 than descriptions of around 80 words in existing generative benchmarks (Jiang et al., 2024; Kaul et al., 2024). In addition, LongHalQA contains approximately 4K object names, indicating more 251 informative data compared with existing benchmarks (Yifan et al., 2023; Hu et al., 2023) with fixed 252 annotations (e.g., 80 object names within COCO dataset). 253

254 255

256

257

258

259

260

3.3 EVALUATION METHODS AND METRICS

We adopt both binary and multiple-choice settings for the Hallucination Discriminiation task, and multiple-choice setting for Hallucination Completion task. For binary answers, we use Accuracy, Precision, and "Yes" ratios as metrics following previous practices (Yifan et al., 2023; Jiang et al., 2024). For multiple-choice setting, we adopt (mc-)accuracy (Liu et al., 2023d) as the evaluation metric, which requires MLLMs to generate the letter (e.g., A, B, C, or D) of the answer option. We randomly shuffle the order of the four options for each MCQ to reduce the impact of option order.

4 LONGHALLGEN: AUTOMATED LONG-CONTEXT HALLUCINATION DATA GENERATION

265 266

Given the lack of long-context image-text hallucination data in the broad area of vision language
 understanding, we dig deep into the proposed LongHallQA and distill LongHallGen, a generic
 pipeline that aims to facilitate the construction of long-context hallucination benchmarks or datasets
 in various multimodal tasks, more detailed processes to be elaborated in the ensuing subsections.

| | Hallucinati | on Discrimination | Hallucination Completion | Auorogo |
|---------------|-------------|-------------------|--------------------------|---------|
| Model | bi-Acc. | mc-ACC. | mc-ACC. | Average |
| MiniCPM-V2-2B | 44.15 | 40.80 | 46.25 | 43.73 |
| Qwen2-VL-2B | 48.21 | 38.76 | 50.36 | 45.78 |
| Fuyu-8B | 43.31 | 23.86 | 23.67 | 30.28 |
| LLaVA-1.5-7B | 38.52 | 35.04 | 36.08 | 36.55 |
| LLaVA-1.5-13B | 41.83 | 43.60 | 37.58 | 41.00 |
| LLaVA-1.6-7B | 44.13 | 45.56 | 43.40 | 44.36 |
| Qwen-VL-Chat | 43.21 | 37.03 | 36.57 | 38.94 |
| LLaVA-1.6-34B | 46.99 | 57.40 | 56.03 | 53.47 |
| Qwen2-VL-72B | 50.36 | 54.78 | 61.50 | 55.55 |
| GPT4o | 52.80 | 47.63 | 56.15 | 52.19 |
| | | | | |

Table 2: Evaluations of MLLMs on LongHalOA with Hallucination Discrimination and Hallucination 271 ').

287 Image Collection and Filtering. To generate informative hallucinatory data, the first step involves 288 selecting images that contain rich content and discarding those with overly simple scenes or rare 289 objects. This can be simply achieved by leveraging off-the-shelf image understanding techniques. We 290 adopt images from the validation set of VisualGenome (Krishna et al., 2017) and Objects365 (Shao et al., 2019) to avoid images being used for training MLLMs, and then analyze and filter them based 291 on dataset annotations and GroundingDINO (Liu et al., 2023c). This straightly leads to 1200 images 292 with complex contents for further hallucinatory data generation. 293

Positive Data Generation. With selected images, long-context data can be generated by leveraging 295 state-of-the-art MLLMs. This process involves designing certain prompts according to specific 296 tasks and domains. We adopt GPT4V to generate long-context texts. Unified prompt templates are designed to allow adjusting the format and scope of generated texts, as illustrated in Appendix B. 297 Note current MLLMs, even GPT4V, suffer from severe hallucinations while generating long-context 298 data. A hallucination check process is required to ensure the quality of the generated data. 299

300 Hallucination Check. The MLLM-generated data are then examined comprehensively for detecting 301 any inherent hallucinations. The examination can be achieved in three steps. First, the generated 302 data undergo a per-sentence self-check by GPT4V, where each piece of data is checked twice to reduce randomness. Second, names of objects present in the data are extracted, and certain image 303 understanding tools such as GroundingDINO (Liu et al., 2023c) are then conducted with detection 304 results feeding to GPT4V for further checking. Finally, the summarized analysis are revised manually, 305 and the revised analysis can serve for the generation of hallucination-explanation data pairs. 306

307 Though MLLMs such as GPT-4V can generate lengthy descriptions, they tend to generate numerous 308 hallucinations as well. Take LongHallQA as an example. Among its generated descriptions for the 309 500 images on image-level descriptions, 394 descriptions (78.8%) contain at least one hallucination. The ratio goes up to 82.4% for the generated conversations. One major cause of the high hallucination 310 rate is due to the increased length of the generated descriptions. In addition, most selected images 311 in LongHalQA have complex scenes, further boosting the possibility of descriptive hallucinations. 312 Nevertheless, such data realistically simulate the hallucinations in actual applications of MLLMs. 313

314 Hallucination-Explanation Pair Generation. Two different prompts are formulated to construct hallucination-explanation (HE) data pairs with the checking outcome. For the data without hallucina-315 tions detected, GPT4V is prompted to modify the data to produce a misleading error within the range 316 of hallucination types suggested in the prompt. For the data containing hallucinations, GPT-4V is 317 prompted to modify the data to contain only one error to form HE pairs. The generated HE pairs are 318 then adopted to construct MCQs for MLLM evaluations. 319

320 Question and Answer Generation. With the generated HE pairs, MLLMs such as GPT-4V can be 321 employed to generate questions for tasks such as hallucination discrimination and hallucination completion in LongHalQA. For discriminative tasks, questions like 'Does the following {Hallucination 322 Data} match the image?' can be formulated to prompt MLLMs to generate four candidate options 323 with explanations. For completion tasks with questions like 'Complete the following {Hallucination

| | Object- | Object-level Description Image- | | | | cription | Multi-round Conversation | | |
|---------------|---------|---------------------------------|-------|-------|-------|----------|--------------------------|-------|------|
| Model | Acc. | Pre. | YR | Acc. | Pre. | YR | Acc. | Pre. | Y |
| MiniCPM-V2-2B | 59.71 | 54.67 | 74.23 | 36.66 | 36.64 | 99.85 | 36.10 | 31.91 | 89. |
| Qwen2-VL-2B | 64.31 | 58.01 | 71.97 | 36.88 | 36.74 | 99.78 | 43.45 | 32.05 | 69.4 |
| Fuyu-8B | 50.29 | 48.52 | 83.65 | 43.29 | 37.28 | 78.79 | 36.35 | 30.72 | 82. |
| LLaVA-1.5-7B | 45.18 | 45.99 | 94.60 | 36.59 | 36.62 | 99.93 | 33.79 | 47.53 | 94.: |
| LLaVA-1.5-13B | 52.70 | 49.96 | 84.16 | 36.95 | 36.95 | 99.71 | 35.85 | 31.93 | 90.0 |
| LLaVA-1.6-7B | 60.51 | 55.31 | 72.85 | 37.10 | 36.82 | 99.56 | 34.79 | 31.90 | 92.8 |
| Qwen-VL-Chat | 58.69 | 53.71 | 79.64 | 36.66 | 36.66 | 100.0 | 34.29 | 31.78 | 93. |
| LLaVA-1.6-34B | 68.61 | 61.65 | 67.96 | 38.26 | 37.22 | 98.10 | 34.10 | 32.16 | 96. |
| Qwen2-VL-72B | 71.60 | 64.46 | 65.11 | 40.08 | 37.87 | 95.84 | 39.40 | 33.59 | 88. |
| GPT40 | 73.94 | 68.31 | 57.81 | 37.92 | 37.13 | 97.03 | 46.32 | 35.71 | 77. |

Table 3: Experiments on LongHalOA for Hallucination Discrimination task with binary an-

Data ..., MLLMs are employed to construct a completion task by providing prefix text from HE pairs 341 and options of candidate sentences for completion. 342

343 LongHallGen exploits MLLMs for most processes in generating long-context hallucination data, 344 except the hallucination checking that involves optional human verification. We expect LongHallGen to serve as a basis for constructing more long-context hallucination data for training and evaluating 345 MLLMs, thereby enhancing their capability and reliability in complex application scenarios. 346

348 Table 4: Experiments on Hallucination Discrimination under multi-choice settings. 349 "Desc" indicates description. 350

Table 5: Experiments on Hallucination Completion under multi-choice settings."Desc" indicates discription.

| Accuracy | Image Desc. | Conversation | Accuracy | Image Desc. | Conversation |
|-------------------------------|-----------------------|-----------------------|-------------------------------|-----------------------|-----------------------|
| MiniCPM-V2-2B Owen2-VL-2B | 39.65 41.55 | 41.96 35.97 | MiniCPM-V2-2B Owen2-VL-2B | 44.07 47.18 | 48.43 53.54 |
| Fuyu-8b | 23.47 | 24.25 | Fuyu-8b | 23.25 | 24.09 |
| LLaVA-1.5-7B LLaVA-1.5-13B | 37.17 45.99 | 32.92 41.21 | LLaVA-1.5-7B LLaVA-1.5-13B | 32.80 31.53 | 39.37 43.62 |
| LLaVA-1.6-7B Qwen-VL-Chat | 49.42 37.97 | 41.71 36.10 | LLaVA-1.6-7B Qwen-VL-Chat | 39.47 33.14 | 47.32 40.00 |
| LLaVA 1.6-34B Qwen2-VL-72B | 60.93 53.57 | 53.86 55.98 | LLaVA-1.6-34B Qwen2-VL-72B | 53.16 59.38 | 58.90 63.62 |
| GPT-40 | 46.57 | 48.69 | GPT-40 | 50.97 | 61.33 |
| | | | | | |

5 EXPERIMENTS

5.1 OVERALL EXPERIMENTS

We adopt LMMs-Eval (Li* et al., 2024) to employ LongHalQA to gauge the hallucination level of 368 MLLMs. The evaluations are performed over nine widely adopted open-source MLLMs, including 369 MiniCPM-V2 (Hu et al., 2024), Qwen series (Bai et al., 2023; Wang et al., 2024b), Fuyu (Bavishi et al., 370 2023), LLaVA series (Liu et al., 2023b; 2024a), and the closed-source GPT-40, covering MLLMs' 371 sizes from 2B to 72B and larger. We present the overall experiments in Tab. 2. We use GPTQ-Int8 372 quantization for Qwen2-VL-72B due to memory constraints. Notably, Qwen2-VL-72B achieves the 373 best accuracy on average for hallucination completion tasks, demonstrating its superior ability to 374 identify hallucinated information and produce reliable content. Such performance advantage may be 375 attributed to the superior capabilities of LLM and their proposed naive dynamic resolution mechanism. Following Qwen2-VL-72B are LLaVA-v1.6-34B and GPT-40. This performance advantage suggests 376 GPT's potential capability of self-correction for hallucinations, given that the LongHalQA is primarily 377 based on hallucination data from GPT. Next, smaller models like MiniCPM-V2, Qwen2-VL-2B, and

340

347

363

364

366 367

378 and LLaVA 1.6-7B also achieved excellent results, surpassing many larger models. It is worth noting 379 that both MiniCPM-V2 and Qwen2-VL-2B adopts reinforcement learning to mitigate hallucinations, 380 indicating that this is an effective method for improving the reliability of MLLMs. One common 381 feature of these leading MLLMs is that they all support high-resolution images, suggesting that 382 resolution plays a significant role in alleviating hallucinations.

384

386

391

392

399 400

401

5.2 EXPERIMENTS ON HALLUCINATION DISCRIMINATION

Binary-Answer Setting. Tab. 3 shows experiments on the hallucination discrimination task under the 387 388 binary answer setting. GPT-40 performs the best among all MLLMs, particularly for the multi-round conversations, with an accuracy gain of 9.5% over other MLLMs. Fuyu-8b shows superior capabilities 389 in identifying hallucinations in long text and achieves the best accuracy among all open-source 390 MLLMs, scoring 45.0% for image-level descriptions and 36.8% for multi-round conversations. We observe that most MLLMs produce a high yes ratio of over 70%, even 99% for image-level description, largely deviating from the ratio of non-hallucinatory data with answers 'yes' in LongHalQA (38.1%). 393 Moreover, Qwen-VL-Chat, MiniCPM-V2-2B, and LLaVA series exhibit unbalanced capabilities in 394 handling text of varying lengths. For object-level descriptions, LLaVA1.6-7B achieves an accuracy 395 of 60.6%, but this drops to 37.1% and 34.7% for image-level descriptions and conversations that are 396 about ten times longer. The lower accuracy, coupled with a significantly high Yes ratio, demonstrates 397 the constraints of existing MLLMs in recognizing hallucinations in long contexts. 398

Table 6: Experiments on different types of hallucinations on discrimination task with multiple-choice setting. The indexes of hallucination types are consistent with Tab. 1.

| | H1 | H2 | H3 | H4 | H5 | H6 | H7 | H8 | H9 | H10 | H11 | H12 |
|---------------|------|------|------|------|------|------|------|------|------|------|------|------|
| MiniCPM-V2-2B | 20.1 | 19.5 | 16.5 | 18.3 | 22.7 | 20.7 | 21.8 | 16.7 | 16.7 | 16.7 | 33.3 | 17.6 |
| Qwen-VL-Chat | 12.4 | 10.1 | 10.0 | 9.8 | 10.0 | 12.9 | 11.9 | 8.3 | 16.7 | 6.7 | 33.3 | 11.7 |
| Fuyu-8B | 27.9 | 26.9 | 30.1 | 25.3 | 26.3 | 22.0 | 20.8 | 27.8 | 26.2 | 30.0 | 40.0 | 29.4 |
| LLaVA-1.5-7B | 9.0 | 9.1 | 8.8 | 4.2 | 10.0 | 8.8 | 12.9 | 6.9 | 19.0 | 10.0 | 20.0 | 5.9 |
| LLaVA-1.5-13B | 22.0 | 22.9 | 26.6 | 23.2 | 26.4 | 30.8 | 26.7 | 26.4 | 35.7 | 26.7 | 26.7 | 35.3 |
| LLaVA-1.6-7B | 32.8 | 40.4 | 40.2 | 30.9 | 30.8 | 43.4 | 43.5 | 38.8 | 47.6 | 20.0 | 20.0 | 47.0 |

408 409 410

411 **Multiple-choice Setting.** Tab. 4 shows experiments on the hallucination discrimination task under the 412 MCQ setting. LLaVA-1.6-34B and Qwen2-VL-72B achieve the highest accuracy in description and conversation data formats, respectively, followed closely by GPT-40. Notably, most MLLMs achieve 413 much higher accuracy in both image description and conversation formats compared to the accuracy 414 under binary settings. This is likely because answer choices include detailed explanations for the 415 involved hallucinations, giving the model a clearer understanding and aiding in selecting the correct 416 option. However, the much lower ranking-based accuracy suggests that MLLMs struggle to correctly 417 discern hallucinations and provide accurate reasons when they cannot directly access all options for 418 reference, consistent with the low accuracy observed in binary settings where explanations are also 419 inaccessible. Interestingly, Fuyu-8B achieves the highest ranking-based accuracy, even surpassing its 420 generation-based accuracy, which may be attributed to its unique decoder-only structure.

421 422 423

424

5.3 **EXPERIMENTS ON HALLUCINATION COMPLETION**

425 Tab. 5 shows experiments on the hallucination completion task. Among open-sourced MLLMs, 426 Qwen2-VL-72B achieves the best performance for both image description and multi-round conversa-427 tion. Two small models, MiniCPM-V2 and Qwen2-VL-2B, also achieved excellent results in this 428 task, ranking just behind three much larger models with the least size of 34B. This further reflects the 429 significance of high-resolution representation and multi-modal RLHF (Yu et al., 2023), which aligns vision and language for trustworthy behavior against object hallucinations in training. Other MLLMs 430 achieve similar ranking performance as those in the hallucination discrimination task, with LLaVA 431 1.6-7B prevailing, followed by Qwen-VL-Chat and LLaVA 1.5-13B.

| | H | all. Discrimination | L | Hall. Completion | A |
|---------------|--------------------------|---------------------|----------------|------------------|---------------|
| Model | (Object) bi-Acc. | (Long) bi-Acc. | mc-ACC. | mc-ACC | Average |
| MiniCPM-V2 | 63.87 (+4.14) | 36.24 (-0.14) | 38.12 (-2.68) | 48.13(+1.88) | 43.90(+0.17) |
| Qwen2-VL-2B | 61.61 (-2.70) | 38.06 (-2.10) | 34.95 (-3.81) | 50.75(+0.39) | 43.87(-1.91) |
| Fuyu-8B | 47.23 (-3.06) | 34.16 (-5.66) | 23.74 (-0.12) | 23.35 (-0.32) | 28.55 (-1.73) |
| LLaVA-1.5-7B | 49.64 (+4.46) | 34.84 (-0.35) | 34.67 (-0.37) | 38.81 (+2.73) | 37.75 (+1.2) |
| LLaVA-1.5-13B | 54.38(+1.68) | 36.29 (-0.11) | 44.12(+0.52) | 40.90(+3.32) | 42.45(+1.45) |
| LLaVA-1.6-7B | 60.36 (-0.15) | 34.68 (-1.26) | 40.68 (-4.88) | 44.48(+1.08) | 42.80 (-1.56) |
| Qwen-VL-Chat | 59.85(+1.16) | 34.76 (-0.71) | 35.38 (-1.65) | 39.06(+2.49) | 39.19(+0.25) |
| LLaVA-1.6-34B | 66.64 (+1.97) | 36.03 (-0.15) | 45.56 (-11.84) | 57.49(+1.46) | 49.76 (-3.71) |
| Qwen2-VL-72B | 73.58(+1.98) | 46.08(+6.34) | 53.11 (-1.67) | 64.82(+3.32) | 57.73(+2.18) |
| GPT-40 | 74.84(+0.90) | 45.50(+3.38) | 50.12(+2.49) | 58.97(+2.82) | 54.79(+2.6) |

Table 7: Experiments on LongHalQA with the modified prompt. 'bi-ACC' and 'mc-ACC' denote binary answer and multiple-choice accuracies. "Object" and "Long" indicate data formats of objectlevel descriptions and long context data of image-level descriptions and conversation

5.4 ANALYSIS OF HALLUCINATION TYPES.

We conduct a detailed analysis of different types of hallucinations in the discrimination task as shown in Tab. 6. We find that most MLLMs exhibit relatively higher accuracy in hallucinations of object existence (H1), attributes (H2), and colors (H3), which are relatively simple to discern because they can be directly observed from images and rely less on detailed comprehension of image content. From the perspective of MLLMs, MiniCPM-V2-2B and Qwen-VL-Chat show balanced strength across different types of hallucinations. Fuyu-8B is competitive across multiple types, but struggles with object location (H6) and relationships (H7). LLaVA1.6-7B outperforms other MLLMs on most types of hallucination, especially for object location (H6) and Text/Sign semantic meanings (H8).

460 5.5 PROMPT ANALYSIS 461

462 We additionally examine the Chain-Of-Thought (COT) on LongHalQA, which has been verified in 463 previous benchmarks (Qian et al., 2024; Jiang et al., 2024) for mitigating hallucinations. We modify 464 the prompt to guide MLLMs to think step by step, and provide more instructions, such as possible 465 types of hallucinations and suggestions for per-sentence verification (Refer to Appendix C for details). As shown in Tab. 7, apart from GPT-40, most MLLMs experience a drop in performance across 466 different tasks when using Chain of Thought (COT)-notably, the larger the language model, the 467 smaller the performance drop. Qwen2-VL-72B only experienced a drop in the multiple-choice task 468 for hallucination Discrimination, with an average increase of 2.18 accuracy. We conjecture that this is 469 largely due to the limited capability of MLLMs to interpret long-context data. Besides, the modified 470 prompts improve short query discrimination and hallucination completion the most but have little 471 effect on discriminate hallucinations in long texts.

472 473 474

475

432

433

434

448 449 450

451

452

453

454

455

456

457

458

459

COMPARISON WITH FREE-GENERATION EVALUATION 6

476 We further compare LongHalQA with evaluations in free-generation scenarios to examine whether 477 the multiple-choice (MCQ) format of the hallucination completion task in LongHalOA accurately captures the true generative capabilities of MLLMs. As described in Sec. 3.1, the multiple-choice 478 hallucination completion task could be transformed into a similar counterpart in a free-form generation 479 setting. Specifically, we randomly selected 200 image-text data from image description and multi-480 round conversation data for the hallucination completion task, respectively, and apply MLLMs to 481 complete the text freely. As shown in Tab. 8, the ranking of hallucination levels, using GPT-4 482 evaluation, is largely consistent with the results from LongHalQA, demonstrating that the proposed 483 MCQ task is able to capture the generative capabilities of MLLMs. 484

Additionally, we compare the efficiency of different approaches in evaluating MLLM's generative 485 hallucination. As shown in Fig fig. 3, for some extremely large models, the time required for

Table 8: Comparison of Multi-Choice (mc-ACC) and Free-Generation (gen-ACC) settings for Hallucination Completion. Under the free-generation setting, MLLMs are provided with preceding contexts and images and are prompted to freely continue the pretext to assess their generative hallucinations.



Figure 3: Comparison of evaluation times under different settings and MLLMs. Our proposed multiple-choice hallucination completion task is significantly faster than other(existing) setups, especially for large models. We measure the time taken to evaluate 1,000 image-text pairs under three different evaluation settings. Only the time that MLLMs take to generate text is measured without considering the time for evaluations by other LLM evaluators. All MLLMs are tested on one A100 except LLaVA 1.6-34B and Qwen2-VL-72B on an H100.

generating descriptions from scratch may be excessively long and impractical. The multiple-choice
 format of Longhalqa, while preserving evaluation effectiveness, significantly improves evaluation
 efficiency and facilitates future expansion of evaluation data in both scale and diversity.

- 7 CONCLUSION
- ⁵²⁶ 527 Т

This paper presents LongHalQA, a novel benchmark with long-context data for evaluating MLLMs' level of hallucinations in more practical scenarios. LongHalQA consists of 6.4k question-answer pairs with long-context data covering 12 types of hallucination. It features two multiple-choice tasks: hallucination recognition and hallucination completion, implementing both discriminative and generative hallucination evaluation in one unified format. It also offers additional assessments of the causes of hallucinations without involving LLM evaluators as in existing benchmarks. We also propose an automated pipeline for generating long-context hallucination data. Extensive tests reveal the constraints of existing MLLMs in handling long-context hallucinations, showing the necessity for more research on robust MLLMs with respect to long-context hallucinations.

537 REFERENCES

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou,
 and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization,
 text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023.

570

582

| 540 | Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. |
|-----|---|
| 541 | Hallucination of multimodal large language models: A survey, arXiv preprint arXiv:2404.18930. |
| 542 | 2024. |
| 543 | |

- Rohan Bavishi, Erich Elsen, Curtis Hawthorne, Maxwell Nye, Augustus Odena, Arushi Somani, and
 Sağnak Taşırlar. Introducing our multimodal models, 2023. URL https://www.adept.ai/
 blog/fuyu-8b.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang,
 Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: Towards general-purpose visionlanguage models with instruction tuning. *Advances in Neural Information Processing Systems*, 36,
 2024.
- ⁵⁵¹ Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023.
- Hongyu Hu, Jiyuan Zhang, Minyi Zhao, and Zhenbang Sun. Ciem: Contrastive instruction evaluation
 method for better instruction tuning. *arXiv preprint arXiv:2309.02301*, 2023.
- Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang,
 Yuxiang Huang, Weilin Zhao, et al. Minicpm: Unveiling the potential of small language models
 with scalable training strategies. *arXiv preprint arXiv:2404.06395*, 2024.
- Qidong Huang, Xiaoyi Dong, Pan Zhang, Bin Wang, Conghui He, Jiaqi Wang, Dahua Lin, Weiming
 Zhang, and Nenghai Yu. Opera: Alleviating hallucination in multi-modal large language models
 via over-trust penalty and retrospection-allocation. *arXiv preprint arXiv:2311.17911*, 2023.
- Chaoya Jiang, Wei Ye, Mengfan Dong, Hongrui Jia, Haiyang Xu, Ming Yan, Ji Zhang, and Shikun
 Zhang. Hal-eval: A universal and fine-grained hallucination evaluation framework for large vision
 language models. *arXiv preprint arXiv:2402.15721*, 2024.
- Prannay Kaul, Zhizhong Li, Hao Yang, Yonatan Dukler, Ashwin Swaminathan, CJ Taylor, and
 Stefano Soatto. Throne: An object-based hallucination benchmark for the free-form generations of
 large vision-language models. *arXiv preprint arXiv:2405.05256*, 2024.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie
 Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language
 and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017.
- Sicong Leng, Hang Zhang, Guanzheng Chen, Xin Li, Shijian Lu, Chunyan Miao, and Lidong
 Bing. Mitigating object hallucinations in large vision-language models through visual contrastive
 decoding. *arXiv preprint arXiv:2311.16922*, 2023.
- Bo Li*, Peiyuan Zhang*, Kaichen Zhang*, Fanyi Pu*, Xinrun Du, Yuhao Dong, Haotian Liu, Yuanhan Zhang, Ge Zhang, Chunyuan Li, and Ziwei Liu. Lmms-eval: Accelerating the development of large multimoal models, March 2024. URL https://github.com/EvolvingLMMs-Lab/
 lmms-eval.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Mitigating
 hallucination in large multi-modal models via robust instruction tuning. In *The Twelfth International Conference on Learning Representations*, 2023a.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 tuning. *arXiv preprint arXiv:2310.03744*, 2023b.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee.
 Llava-next: Improved reasoning, ocr, and world knowledge, January 2024a. URL https: //llava-vl.github.io/blog/2024-01-30-llava-next/.
- Jiazhen Liu, Yuhan Fu, Ruobing Xie, Runquan Xie, Xingwu Sun, Fengzong Lian, Zhanhui Kang, and Xirong Li. Phd: A prompted visual hallucination evaluation dataset. *arXiv preprint arXiv:2403.11116*, 2024b.

| 594 595 596 597 | Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. <i>arXiv preprint arXiv:2303.05499</i> , 2023c. |
|---------------------------------|--|
| 598 599 600 | Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. Mmbench: Is your multi-modal model an all-around player? <i>arXiv:2307.06281</i> , 2023d. |
| 601 602 603 | Holy Lovenia, Wenliang Dai, Samuel Cahyawijaya, Ziwei Ji, and Pascale Fung. Negative object presence evaluation (nope) to measure object hallucination in vision-language models. <i>arXiv</i> preprint arXiv:2310.05338, 2023. |
| 605 606 | Yusu Qian, Haotian Zhang, Yinfei Yang, and Zhe Gan. How easy is it to fool your multimodal llms? an empirical analysis on deceptive prompts. <i>arXiv preprint arXiv:2402.13220</i> , 2024. |
| 607 608 609 | Haoyi Qiu, Wenbo Hu, Zi-Yi Dou, and Nanyun Peng. Valor-eval: Holistic coverage and faithfulness evaluation of large vision-language models. <i>arXiv preprint arXiv:2404.13874</i> , 2024. |
| 610 611 612 | Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In <i>Proceedings of the</i> <i>IEEE/CVF international conference on computer vision</i> , pp. 8430–8439, 2019. |
| 613 614 615 616 | Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. <i>arXiv preprint arXiv:2309.14525</i> , 2023. |
| 617 618 619 | Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Ming Yan, Ji Zhang, and Jitao Sang. Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation. <i>arXiv preprint arXiv:2311.07397</i> , 2023. |
| 620 621 622 623 | Lei Wang, Jiabang He, Shenshen Li, Ning Liu, and Ee-Peng Lim. Mitigating fine-grained hallucina- tion by fine-tuning large vision-language models with caption rewrites. In <i>International Conference</i> <i>on Multimedia Modeling</i> , pp. 32–45. Springer, 2024a. |
| 624 625 626 | Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. <i>arXiv preprint arXiv:2409.12191</i> , 2024b. |
| 627 628 629 630 631 | Li Yifan, Du Yifan, Zhou Kun, Wang Jinpeng, Zhao Wayne Xin, and Wen Ji-Rong. Evaluating object hallucination in large vision-language models. In <i>The 2023 Conference on Empirical Methods in Natural Language Processing</i> , 2023. URL https://openreview.net/forum?id=xozJw0kZXF. |
| 632 633 634 | Shukang Yin, Chaoyou Fu, Sirui Zhao, Tong Xu, Hao Wang, Dianbo Sui, Yunhang Shen, Ke Li, Xing Sun, and Enhong Chen. Woodpecker: Hallucination correction for multimodal large language models. <i>arXiv preprint arXiv:2310.16045</i> , 2023. |
| 635 636 637 638 | Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, and Tat-Seng Chua. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. <i>arxiv</i> , 2023. |
| 639 640 641 | Zihao Yue, Liang Zhang, and Qin Jin. Less is more: Mitigating multimodal hallucination from an eos decision perspective. <i>arXiv preprint arXiv:2402.14545</i> , 2024. |
| 642 643 644 | Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. <i>arXiv preprint arXiv:2304.10592</i> , 2023. |
| 645 646 647 | Lanyun Zhu, Deyi Ji, Tianrun Chen, Peng Xu, Jieping Ye, and Jun Liu. Ibd: Alleviating hallucinations in large vision-language models via image-biased decoding. <i>arXiv preprint arXiv:2402.18476</i> , 2024. |

| 648 649 | А | Appendix |
|-------------------|------------|---|
| 650 651 | В | HALLUCINATION TYPES IN LONGHALQA |
| 652 653 654 | We in F | present the definitions of different types of hallucinations as follows. Some examples are show g. 4 and Fig. 5. |
| 655 | | H1 (Non) Existent Objects. The described objects do not exist in the given image. |
| 656 657 | | H2 Object Attributes. The appearances (shape, pattern, etc.), types, or other attributes objects are incorrectly described. |
| 659 660 661 | | H3 Object Color. The colors of objects are incorrectly described. (Due to the high numb of hallucinations of object color descriptions, we list "object color" separately from object attributes.) |
| 662 663 | | H4 Object States. The states of objects, such as movement, orientation, the actions of t person, etc, are incorrectly described. |
| 664 | | H5 Number of Objects. The number of objects is incorrectly stated. |
| 665 | | H6 Object Locations. The locations of objects in the image are incorrectly described. |
| 666 | | H7 Object Relationships . The relations or relative positions between two or multiple object |
| 667 | | are incorrectly described. |
| 668 | | H8 Text/Sign Meaning The text in the image is wrongly discerned or the meanings of sig |
| 669 | | such as street signs, advertisements, price tags, etc, are incorrectly described. |
| 670 | | H9 Environment Description. Wrong descriptions or adjectives of the environment or location |
| 672 | | for example, indoors, outdoors, rural, urban, bookstore, food market, etc. |
| 673 | | 110 Background Description. The descriptions of activities, scenes, objects, etc., in t |
| 674 675 | | background of the image are hallucinatory. Examples of descriptions include the presence mountains, buildings, skyscrapers people in the background, or the description of no visil |
| 676 | | people in the background. |
| 677 678 | | judgment of the time without clear evidence. |
| 679 | | H12 Weather. Incorrect descriptions of weather or sky conditions. |
| 680 | | |
| 681 | C | DETAILS OF LONGHALL GEN |
| 683 | C | DETAILS OF LONGHALLOEN |
| 684 | We | present the prompts for different steps in LongHallGen in Fig. 6 Fig. 7 Fig. 8 and Fig. 9 |
| 685 | | |
| 686 | Л | CHAIN OF THOUGHT BROMPT FOR EVALUATION |
| 687 | D | CHAIN-OF-THOUGHT PROMPT FOR EVALUATION |
| 688 | Wa | avaming the impact of Chain of Thought on LongHelOA. We append additional prompts |
| 689 | sho | vn in Fig. 10 before the questions for the hallucination discrimination task and the hallucinati |
| 690 | con | pletion task. |
| 691 | | • |
| 692 | | |
| 693 | | |
| 694 | | |
| 695 | | |
| 696 | | |
| 697 | | |

709

710 711

712

713

714

715

716

717

718

719 720

721

722

723

724

725

726

727

728 729

730

731

732

733

734

735

736

H2

H3

H4

H5

H6

702



The image showcases a wooden desk in a room scattered with various electronic devices and objects. A laptop with an open page is centrally placed on the desk. To the right of the laptop, there is a second closed laptop and a printer. On the left, there is a desk phone connected to a cord. Multiple cards, notepads, and a pen are also visible on the desk. There is a lamp with a golden base and a cream-colored shade to the right of the desk. The floor has a geometric pattern, and a portion of a bed can be seen on the far left.

The image shows four musicians performing onstage with a red curtain as the background. From left to right: a man seen from the back wearing a dark suit is playing a grand piano, a man with dreadlock hair is playing the double bass, another man in a dark jacket and brown pants is playing the saxophone, and a man in glasses wearing a dark suit is playing a drum set labeled 'YAMAHA'. The ambience suggests an intimate live jazz performance with dim, warm lighting.





The image displays a merchandise booth with a variety of t-shirts for sale. Hanging on a makeshift clothesline above the booth are twelve t-shirts in different colors, including maroon, black, grey, blue, orange, red, and yellow, each with different graphic designs. Below, on the booth table, more t-shirts and folded clothing items are neatly arranged in rows, with price signs indicating "\$15 TEES" and "\$10 TEES." Two male individuals appear to be attending the booth,

In the image, a man stands to the left side in what appears to be a kitchen or a kitchenette within an office environment. He is dressed in a blue buttoned-down shirt, grey trousers, and wears a belt. The man's expression is neutral, and his hands are in his pockets. In front of him is a table laden with various food items, likely set up for a meal or a gathering. The table presents a variety of dishes including what appears to be pies, bowls filled with what might be berries or salads, plates with pastries or rolls, and at least one large silver mixing bowl. The table also holds napkins





746 747

748

749

750

751

752

753

754

755



The image depicts an office environment. There are four people in the image, with one standing man and three seated individuals, two men, and a women. The standing man, wearing glasses and a green striped shirt, is handing a piece of paper to the seated man with a beard, who is wearing a dark shirt and is reaching out to take the paper. The second seated man, wearing a pink shirt, is looking at the paper, while a seated woman, in a white jacket and red top.....

The image showcases an urban city square with a mixture of classic European architecture. There are multi-story buildings with different facades, colors, and architectural details. To the left, there is a market area with various tents and a concentration of people, suggesting a bustling atmosphere. The foreground features construction with visible orange cones and barriers, indicating ongoing development or maintenance. There are several trees scattered



Figure 4: Visualizations of hallucination types from H1 to H6.



Figure 5: Visualizations of hallucination types from H7 to H12.

| 811 | |
|---|---|
| 812 | Prompt for Positive Data Generation |
| 013 | System Prompt: You are a nowerful visual expert a nowerful image captioner and an Al visual assistant who can |
| 014 | understand and analyze image content and construct relevant data. |
| 010 | { Prompt for Specific Data Format }. Content that might be mentioned include { Possible Content }, etc. Only |
| 010 | describing the content one can determine confidently from the image. Minimize aesthetic descriptions as much as |
| 017 | possible. The following is example templete for your response: |
| 010 810 | { Response Template } |
| 820 | Prompt for Object-level Description |
| 821 | |
| 822 | { System Prompt } You are a powerrul visual expert, a powerrul image captioner, and an AI visual assistant who can |
| 823 | specific object or specific objects in the image. Each description should be in one sentence. Content that might be |
| 824 | mentioned for describing objects include object types, colors, states, actions, number of objects, precise object |
| 825 | locations, texts or OCR results, relationships or relative positions between objects, etc. Only describing the content |
| 826 | one can determine confidently from the image. Minimize aesthetic descriptions as much as possible. Do not describe what you can not discover from the image. |
| 827 | The following is an example template for your response: |
| 828 | Object Descriptions:\n1.\n2.\n3.\n4.\n5.\n6.\n7.\n8.\n9.\n10. |
| 829 | Prompt for Image-level Description |
| 830 | System Prompt Create a detailed description to describe the content of the given image. Compress the description |
| 831 | in one paragraph. Content that might be mentioned include object types, colors, states, actions, number of objects, |
| 832 | precise object locations, texts or OCR results, relationships and relative positions between objects, environment, |
| 833 | background, time, weather, etc. Only describing the content one can determine confidently from the image. |
| 834 | The following is an example template for your response: |
| 835 | Image Description: |
| 836 | Prompt for Multi-round Conversation |
| 837 | |
| 501 | |
| 838 | { System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to |
| 838 839 | {System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about |
| 838 839 840 | {System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. |
| 838 839 840 841 | {System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. |
| 838 839 840 841 842 | System Prompt Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects objects and the conversations are previous and avoid asking or relations between objects. |
| 838 839 840 841 842 843 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure |
| 838 839 840 841 842 843 844 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize |
| 838 839 840 841 842 843 844 845 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. |
| 838 839 840 841 842 843 844 845 846 | System Prompt Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's nawer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content leaving most of the content unchanged. You can modify abject states. |
| 838 839 840 841 842 843 844 845 846 847 | System Prompt Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time. |
| 838 839 840 841 842 843 844 845 846 847 848 | System Prompt Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 | System Prompt Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 851 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 851 852 853 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. The following is an example template for generated data. Original Conversation: User:[Start the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 851 852 853 | System Prompt Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversation include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. The following is an example template for generated data. Original Conversation: User:[Start the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by 1)seeking clarification on information in the previous answer; 2)asking |
| 838 839 840 841 842 843 844 845 844 845 846 847 848 849 850 851 852 853 854 | System Prompt } Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. The following is an example template for generated data. Original Conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by 1]seeking clarification on information in the previous answer; 2 |
| 838 839 840 841 842 843 844 845 844 845 846 847 848 849 850 851 852 853 854 855 | (System Prompt) Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. The following is an example template for generated data. Original Conversation: User:[Start the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by alseed on the image content.] User:[Continue the conversation until the image is julk discussed 1 |
| 838 839 840 841 842 843 844 845 844 845 846 847 848 849 850 851 852 853 854 855 856 | (System Prompt) Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation: User:[Start the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by 1]seeking clarification on information in the previous answer; 2]asking additional questions based on previous question; 3] asking new questions.] Assistant:[Answer the user's question. Maintain a fluid and coherent dialogue flow.] [continue the |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 851 852 853 854 855 855 856 857 | (System Prompt) Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. The following is an example template for generated data. Original Conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by alsing a question about the image.] Assistant:[Answer the user's question. Maintain a fluid and coherent dialogue flow.] [continue the conversation. [Continue the conversation until the |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 851 852 853 854 855 855 856 857 858 | (System Prompt) Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. The following is an example template for generated data. Original Conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by asking a question shoul the image.] Assistant:[Answer the user's question. Maintain a fluid and coherent dialogue flow.] [continue the conversation until the image is fully discussed.]< |
| 838 839 840 841 842 843 844 845 846 845 846 847 848 849 850 851 852 853 854 855 855 855 856 857 858 859 860 | (System Prompt) Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize aesthetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation: User:[Continue the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by asking a question about the image.] Assistant:[Answer the user's question. Maintain a fluid and coherent dialogue flow.] [continue the conversation until the image is fully discussed.] Modified Conversation:[Edit the original conver |
| 838 839 840 841 842 843 844 845 844 845 846 847 848 849 850 851 852 853 854 855 855 856 855 856 857 858 859 860 861 | (System Prompt) Generate conversation data for a situation where a user and an assistant discuss the content and details of the image. The user's role is to ask various questions about the image. The role of the assistant is to respond to the user's questions. During the conversation, the user can ask multiple progressive questions about certain content or objects in the assistant's previous answers to inquire about their details from general to specific. Except when starting a new topic, users should try to ask specific questions and avoid asking overly broad questions. Contents that can be discussed in the conversations include types, colors, states, and actions of objects, the number of objects, object locations, text or OCR results, doublechecked relative positions or relationships between objects, environment, background, time, weather, etc. The conversation should be natural, continuous, and logical. Ensure that the conversation remains focused on the content one can determine confidently from the image. Minimize asethetic descriptions and keep the interaction informative and engaging. Then, edit the assistant's answer in the generated conversation data and deliberately introduce one misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify object states, actions, number of objects, relationships or relative positions between objects, environment, background, time, etc., or introduce plausible objects that may exist according to the scene of the image. Explain the error in the modified conversation. User:[Start the conversation by asking a question about the image.] Assistant:[Answer the user's question based on the image content.] User:[Continue the conversation by 1]seeking clarification on information in the previous answer; 2]asking additional questions based on previous questions; 3] asking new questions.] Assistant:[Answer the user's question. Maintain a fluid and coherent dialogue flow.] [continu |

| | Q Drownt for Hollycinstion Check |
|---|--|
| _ | |
| | { System Prompt: You are a powerful visual expert, a powerful image captioner, and an Al visual assistant who can understand and analyze image content, construct and analyze relevant data.} { Prompt for Specific Data Format } Pay special attention to check { Possible Content }. |
| | Following is the { Data Format } : { Generated Data } |
| | Following is an example template for your response: { Response Template } |
| | Prompt for Object-level Description |
| | <pre>{ System Prompt } You will be given 10 sentences describing the objects in the image. Check each sentence of description carefully to see if any content does not match the image content. If the description exactly matches the contents in the image, output [Match]. If anything in the description does not fit the image, output [Do not Match] and explain. Pay attention to the content of the description including objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationships and relative positions between objects, environment, background, time, weather, etc. The following are the descriptions of objects: Following is an example answer template: 1.Origin Sentence:[The 1st description]. Result:[Match] / [Do not match][Explanations if it does not match.] 2.Origin Sentence:[The 2st description]. Result:[Match] / [Do not match][Explanations if it does not match.]</pre> |
| | Prompt for Image-level Description |
| | { System Prompt } Check each sentence in the following description carefully to see if any content does not match |
| | the image content. If the sentence exactly matches the image content, output [Match]. If there is anything in the sentence that does not fit the image, output [Do not Match] and give an explanation. Pay attention to the content of the description including objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationships and relative positions between objects, environment, background, time, weather, etc. The following is the description: Following is an example template for your response: 1.Origin Sentence:[The 1st sentence]. Result:[Match] / [Do not match][Explanations if it does not match.] 2.Origin Sentence:[The 2st sentence]. Result:[Match] / [Do not match][Explanations if it does not match.] |
| | Prompt for Multi-round Conversation |
| | { System Prompt } You will be given a conversation between the user and the assistant. Carefully check whether there is any content in each sentence of each assistant's answer that does not match the image content. If the sentence exactly matches the image content, output [Match]. If there is anything in the sentence that does not fit the image, output [Do not Match] and give an explanation. First, analyze each sentence and find its mentioned objects and corresponding content, such as objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationships and relative positions between objects, environment, background, time, weather, etc. Then, check whether each mentioned object exists in the image. If it exists, then check the correctness of each piece of content corresponding to that object. The following is the conversation: |
| | 1.1.Origin Sentence:[1st sentence in 1st answer]. Result:[Match] / [Do not match][Explanations if it does not match.] |
| | 1.2.Origin Sentence:[2nd sentence in 1st answer]. Result:[Match]/[Do not match][Explanations if it does not match.] |
| | 3.1.Origin Sentence:[The 1nd sentence in the 3st answer] |
| C | |
| | Figure 7: Prompt for Hallucination Check in LongHallGen. |
| | |
| | |
| | |
| | |

| 918 | |
|------|---|
| 919 | |
| 920 | Q Drownt for Hellusinstian Evaluation Dair Consertion |
| 921 | Prompt for Hallucination-Explanation Pair Generation |
| 922 | { System Prompt: You are a powerful visual expert who can understand and analyze the content of images and |
| 923 | construct relevant data.} |
| 924 | Prompt for Hallucinatory Data: |
| 925 | You will be given { Data Format } of the image, and an analysis about incorrect sentences in the { Data Format } that |
| 926 | do not match the image content and corresponding reasons. Your task is to 1) generate the correct { Data Format } has ed on the above content. 2) Select at most three errors and then modify the { Data Format } for each selected |
| 927 | error (correct other errors in the { Data Format } so that it only contains this one error). Output the correct { Data |
| 028 | Format }, the { Data Format } corresponding to each error, and the explanation of the error. Pay attention to |
| 020 | maintaining the logic and coherence of { Data Format } when making modifications. |
| 020 | Prompt for Non-Hallucinatory Data: |
| 930 | You will be given { Data Format } of the image. Your task is to edit the { Data Format } and deliberately introduce one |
| 931 | misleading error that doesn't match the image's content, leaving most of the content unchanged. You can modify { |
| 932 | Ivpes of Hallucinations }. Pay attention to maintaining the logic and coherence of description when making |
| 933 | Thouncations. Output the original description, the modified description and the explanation of the introduced error. |
| 934 | Following is the description (Concerned Date) |
| 935 | Following is the applysis: { Generated Data } |
| 936 | Following is the analysis, Tranuchation Analysis <i>f</i> |
| 937 | Tonowing is an example template for the response template j |
| 938 | Prompt for Image-level Description |
| 939 | { System Prompt } |
| 940 | Prompt for Hallucinatory Data: |
| 941 | Following is an example template for your response: |
| 0/12 | Explanation [Briefly summarize the content of the correct description in one sentence to suggest the description |
| 042 | matches the image content.] |
| 943 | Incorrect Description:[Description with one error preserved.] |
| 944 | Wrong Sentence in Description : [the incorrect sentence in the above modified description.] |
| 945 | Explanation:[Explain the error in the description.] |
| 946 | [Repeat the process to generate incorrect descriptions for each error in the description.] |
| 947 | Prompt for Non-Hallucinatory Data: |
| 948 | Original Description: [Description without errors.] |
| 949 | matches the image content.] |
| 950 | Incorrect Description:[Description with one error introduced.] |
| 951 | Wrong Sentence in Description :[the incorrect sentence in the above modified description] |
| 952 | Explanation:[Explain the error in the description.] |
| 953 | Prompt for Multi-round Conversation |
| 954 | {System Prompt} |
| 955 | Prompt for Hallucinatory Data: |
| 956 | Following is an example template for your response: |
| 957 | Correct conversation:[Complete conversation without errors.] |
| 958 | Explanation: [Briefly summarize the content of the correct conversation in one sentence to suggest the conversation |
| 050 | matches the image content.] |
| 939 | Incorrect Conversation 1:[The modified complete conversation with one error preserved.] |
| 061 | Explanation: [Explain the error in this incorrect conversation.] |
| 301 | Incorrect Conversation 2: |
| 962 | Prompt for Non-Hallucinatory Data: |
| 963 | Original Conversation:[Conversation without errors.] |
| 964 | Explanation: [Briefly summarize the content of the conversation in one sentence to suggest the conversation |
| 965 | matches the image content.] |
| 966 | Incorrect Conversation: [Conversation with one error introduced.] |
| 967 | Final Sentence in conversation. [the incorrect sentence in the above modified conversation] |
| 968 | |
| 969 | |
| 970 | Figure 8: Prompt for Hallucination-Explanation Pair Generation in LongHallGen. |

| | { System Prompt: You are a powerful visual expert who can understand and analyze the content of images and |
|---|--|
| | construct relevant data.} |
| | You will be given { Data Format }, and an answer of whether the { Data Format } match the image content and |
| | The given answer should be paraphrased as one of the options and the correct option, then add \"(correct)\" as a |
| | tag after this correct option. Options are ordered by $\mathbb{Z} \times \mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}$ and $\mathbb{Z} \times \mathbb{Z}$ |
| | of a \"Yes\" or \"No\" and an explanation. One of the options should begin with yes, and give an statement |
| | indicating that { Data Format } match the image content. The other three options should start with \"No\" and poin |
| | out a possible error in { Data Format }. Options other than the correct one should contain errors. Errors can be of |
| | two types: 1). the incorrect \"yes\" or \"no\" compared to the correct option; 2) the given explanation is incorrect. |
| | Do not construct explanations by indicating that the { Data Format } does not mention something in the image. For |
| | all answers, the explanations given must be logical, appropriate to the question, relevant to the image content, and |
| | be compressed in one sentence. |
| | Following is the answer() all |
| | Following is an example template for your output |
| | Ouestion: { Ouestion } |
| | Answer Options:\nA.\nB.\nC.\nD.\n |
| | |
| _ | Prompt for Constructing the Hallucination Completion task |
| | { System Prompt: You are a powerful visual expert who can understand and analyze the content of images and |
| | construct relevant data.} |
| | You will receive { Data Format }, and an analysis of why and which sentence of { Data Format } do not match the |
| | content of the image. Your task is to create a continuation task. First, output the text of the { Data Format } before the incorrect centence. Then, provide four continuation options; one is the original incorrect centence, another is |
| | the corrected sentence hased on the analysis, the third ontion introduces a notential error by modifying the correct |
| | sentence, and the last option is a sentence containing an error that you write based on the preceding text and the |
| | image content. Add \"(correct)\" as a tag after the option of the correct sentence. Options are ordered by \"A.\", |
| | \"B.\", \"C.\", and \"D.\". |
| | Following is the { Data Format } : |
| | Following is the wrong sentence in the { Data Format }: |
| | Following is the analysis: |
| | Following is an example template for your output. |
| | Answer Ontions: $\nA \nB \nC \nD \n$ |
| ł | Figure 9: Prompt for <i>Question and Answer Generation</i> for the hallucination discrimination task the hallucination completion task, respectively. |
| | Modified Prompt for Hallucination Discrimination |
| | given a text about the image content. Your task is to check whether the text data match the image content. Please think step by step. First, analyze the image content including objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationships and relative positions between objects, |
| - | any ironmont background time weather ate. Then carefully check each contence in the text data based on the |
| - | above analysis of image content. |
| | above analysis of image content. |
| | above analysis of image content. |
| | A modified Prompt for Hallucination Completion You are a powerful Al assistant who can understand the image content. You will be given a text about the image content. Your task is to continue the text. When continue writing. Pay attention to the image content including. |
| | A modified Prompt for Hallucination Completion You are a powerful AI assistant who can understand the image content. You will be given a text about the image content. Your task is to continue the text. When continue writing, Pay attention to the image content including objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationshin |
| | A modified Prompt for Hallucination Completion You are a powerful AI assistant who can understand the image content. You will be given a text about the image content. Your task is to continue the text. When continue writing, Pay attention to the image content including objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationship and relative positions between objects, environment, background, time, weather, etc. Be careful to avoid from |
| | A modified Prompt for Hallucination Completion You are a powerful AI assistant who can understand the image content. You will be given a text about the image content. Your task is to continue the text. When continue writing, Pay attention to the image content including objects' types, colors, states, actions, number of objects, precise object locations, texts or OCR results, relationship and relative positions between objects, environment, background, time, weather, etc. Be careful to avoid from responding text that contains hallucinations or error. |