

MiCEval: Unveiling Multimodal Chain of Thought’s Quality via Image Description and Reasoning Steps

Xiongtao Zhou^{1*} Jie He^{2*} Lanyu Chen¹ Jingyu Li⁴ Haojing Chen⁴
V́ctor Gutírrez-Basulto^{3†} Jeff Z. Pan^{2†} and Hanjie Chen^{5†}

¹ Waseda University, Japan ² School of Informatics, University of Edinburgh, UK

³ School of Computer Science and Informatics, Cardiff University, UK

⁴ Independent Researcher ⁵ Rice University, US

alenai.tao@ruri.waseda.jp, j.he@ed.ac.uk

j.z.pan@ed.ac.uk, hanjie@rice.edu

Abstract

Multimodal Chain of Thought (MCoT) is a popular prompting strategy for improving the performance of multimodal large language models (MLLMs) across a range of complex reasoning tasks. Despite its popularity, there is a notable absence of automated methods for evaluating the quality of reasoning steps in MCoT. To address this gap, we propose **Multimodal Chain-of-Thought Evaluation (MiCEval)**, a framework designed to assess the correctness of reasoning chains by evaluating the quality of both the *description* and each *reasoning step*. The evaluation of *description* component focuses on the accuracy of the image descriptions, while the *reasoning step* evaluates the quality of each step as it is conditionally generated based on the preceding steps. MiCEval is built upon a fine-grained dataset with annotations that rate each step according to *correctness*, *relevance*, and *informativeness*. Extensive experiments on four state-of-the-art MLLMs show that step-wise evaluations using MiCEval align more closely with human judgments compared to existing methods based on cosine similarity or fine-tuning approaches. MiCEval datasets and code can be found in <https://github.com/alennai97/MiCEval>.

1 Introduction

Multimodal chain-of-thought (MCoT) enhances the performance of multimodal large language models (MLLMs) on complex reasoning tasks by generating explicit multi-step reasoning chains to achieve the final goal (Zhang et al., 2024; Gao et al., 2023; Chen et al., 2024d; Shao et al., 2024). It not only improves the interpretability of MLLM outputs but also offers a valuable framework for assessing their reasoning capabilities. Most existing evaluations (Lu et al., 2022; Zhang et al., 2024; Chen et al., 2024d) focus exclusively on the

* Equal Contribution.

† Corresponding author




Figure 1: We exemplify how CLIP and ReCEval did not choose the correct MCoT answer from the two model-generated MCoT answers, but MiCEval succeeded.

correctness of the final prediction in MCoT, while overlooking the quality of the reasoning process within the chains, failing to capture errors and unreliable steps that may occur throughout the chain (Lyu et al., 2023; Turpin et al., 2023). To establish a principled framework to more accurately reflect the quality of MCoT, it is essential to evaluate each step in the reasoning chain individually.

A common approach for evaluating (unimodal) CoT is to compare model-generated reasoning chains with human-written reference chains (Cliniciu et al., 2021; Welleck et al., 2022; Saparov and He, 2023). However, this method incurs substantial costs. As a result, recent research has shifted towards reference-free evaluation methods, fine-tuning models on human-annotated reasoning chain datasets (Golovneva et al., 2022; Prasad et al., 2023). Yet, these methods often fall short in delivering robust evaluation metrics due to the inherent limitations of fine-tuning. The use of large language models (LLMs) as evaluators has emerged as a promising alternative (He et al., 2024;

Xia et al., 2024) but they cannot be directly applied to MCoT, as they are unable to effectively evaluate the image descriptions in MCoT; e.g., in Figure 1, the statement “*In Pair 1, both magnets’ North poles face each other directly*” is an image description. Methods that focus solely on the textual modality cannot properly evaluate the correctness of image descriptions. In the multimodal domain, several evaluation metrics have been proposed to assess the alignment between images and descriptions by calculating their cosine similarity (Hessel et al., 2021; Li et al., 2023a). However, their effectiveness in evaluating MCoT remains unexplored.

The answers generated by MCoT involve both image descriptions (Chen et al., 2023b,c) and their corresponding reasoning steps. Therefore, for an MCoT to be correct, all steps related to both image descriptions and reasoning must be accurate, relevant, and contribute effectively to reaching the final answer. This highlights the need for a method that evaluates both the visual and textual components, providing a comprehensive assessment of the overall quality of the MCoT output. As noted in the context of CoT by Jacovi et al. (2024), the creation of a multi-step evaluation dataset is essential for accurately assessing the abilities of MLLMs as MCoT verifiers. Such a comprehensive dataset will not only enable detailed evaluation of an MCoT verifier’s effectiveness but will also offer insights into the overall quality of MCoT outputs.

In this paper, we propose MiCEval, a framework designed to evaluate MCoT by breaking down MCoT chains into *description* and *reasoning* steps, and assessing the quality of each. MiCEval decomposes MCoT verification into five core tasks across both visual and textual modalities: description correctness, image relevance, logical correctness, logical relevance, and informativeness. Additionally, MiCEval introduces a dataset comprising 903 MLLM-generated MCoT answers and 2,889 human-annotated MCoT steps, which serve as the foundation for our comprehensive step-wise evaluations of both MCoT and its verifiers.

Each MCoT step in the MiCEval dataset is initially categorized as either a “Description Step”, a “Reasoning Step”, or a “Both Step” (a combination of description and reasoning). For the description steps, we provide fine-grained annotations for correctness and relevance, while for the reasoning steps, we collect labels for (1) correctness: each step is generated based on previously valid information; (2) relevance: every step is related to either

the image or answering the question; and (3) informativeness: steps should provide new, relevant information. Furthermore, we annotate the specific error types for steps identified as description incorrect or logically incorrect. At the MCoT level, we annotate whether the overall MCoT answer is correct.

Leveraging the diverse annotation tasks in the MiCEval dataset, the MiCEval framework provides a comprehensive range of evaluation dimensions. From the MCoT perspective, MiCEval not only facilitates step-level reasoning evaluation but also assesses the overall quality of MCoT. Additionally, MiCEval enables a detailed, fine-grained evaluation of the reasoning capabilities of MLLMs. In this work, we focus primarily on evaluating: (1) the effectiveness of MLLMs as verifiers, and (2) the correlation between MiCEval correctness metrics and human preferences.

Our **main contributions** are as follows:

1. We propose a comprehensive step-level annotation protocol to construct a high-quality, fine-grained, human-annotated MCoT dataset.
2. We assess MLLM verifiers’ capabilities on each fine-grained task, identifying specific weaknesses, such as poor performance on complex reasoning tasks.
3. We introduce the MiCEval multidimensional correctness metrics for evaluating MCoT answers. Moreover, we demonstrate that MiCEval enables MLLM evaluators to better align with human judgment.

2 Problem Definition and Formalization

In this section, we define MCoT and formalize fine-grained tasks for validating and evaluating MCoT depending on the type of steps. An MCoT consists of an image \mathcal{I} , a question \mathcal{Q} , and a reasoning chain $\mathcal{RC} = r_1, \dots, r_n$:

$$MCoT = \text{Prompt}(\mathcal{I}, \mathcal{Q}, \mathcal{RC})$$

$\text{Prompt}()$ represents the transformation of multimodal inputs into task-specific instruction formats. The main difference between of CoT and MCoT is the inclusion of a visual input in the latter. In our study, we focus on *MCoT answer*, namely the chain \mathcal{RC} generated from \mathcal{I} and \mathcal{Q} . Each step r_i , $i \in [1, n]$, is a complete sentence, and r_n is the final step that contains the prediction corresponding to the question.

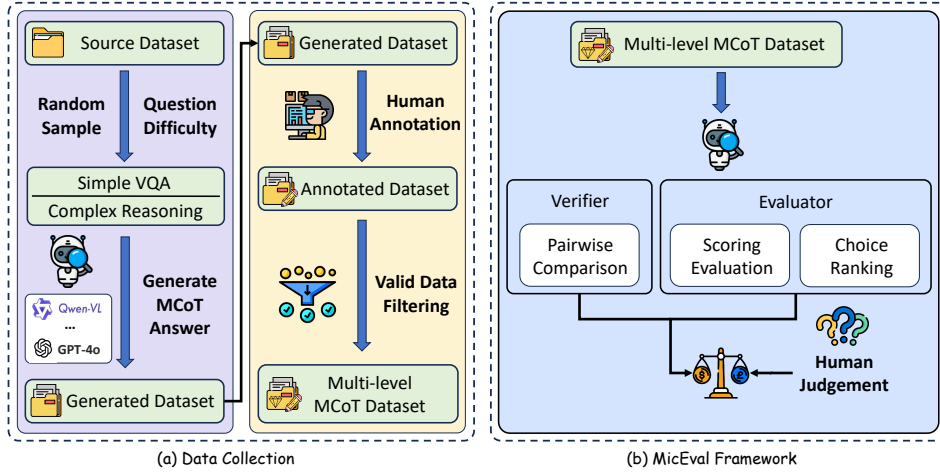


Figure 2: Our work consists of two main parts: (a) sampling questions from the source datasets, generating MCoT answers using four MLLMs, followed by high-quality human annotation and filtering to create the MiCEval dataset; (b) a detailed illustration of our MiCEval framework.

2.1 Verification and Evaluation Tasks

In MiCEval we define two evaluation tasks: Verification (Jacovi et al., 2024) and Evaluation (Chen et al., 2024a). Figure 2 illustrates the MiCEval framework.

Verification. In this task, we aim to assess the ability of MLLMs as verifiers. Given an input MCoT, the MLLM outputs either 0 (“Incorrect”) or 1 (“Correct”).

Evaluation. The goal of this task is to evaluate the correlation between MiCEval correctness metrics and human performance, using MLLMs as evaluators. Additionally, we aim to determine whether MiCEval enables MLLMs to achieve closer alignment with human judgments. For a given input MCoT, the MLLM generates a correctness score within the range of $[0, 1]$.

2.2 MCoT-level and Step-level Evaluation

There are two distinct evaluation settings in MiCEval: MCoT-level and step-level evaluation.

MCoT-level. Given the input $\mathcal{I}, \mathcal{Q}, \mathcal{RC}$, where the \mathcal{RC} comprising $\{r_1, r_2, \dots, r_n\}$ is evaluated as a whole. MLLM generates a score between $[0, 1]$ that reflects the correctness of the entire sequence \mathcal{RC} .

Step-level. For the given input \mathcal{I}, \mathcal{Q} , the current step r_i , and the previous steps $\mathcal{RC}_{i-1} = \{r_1, r_2, \dots, r_{i-1}\}$, only the current step r_i is evaluated, the MLLM outputs a score between $[0, 1]$ that reflects the correctness of r_i .

2.3 Step Type

The definition of correctness varies depending on the type of step. For a *description step*, the focus is on description correctness and relevance. For a *reasoning step*, the emphasis is on logical correctness, relevance, and informativeness.

In an MCoT sequence, the steps are generated from information extracted from both visual and textual modalities. We define steps derived solely from visual information (i.e., describing visual content) as *description steps*. Steps that involve information inferred beyond the visual content, incorporating the question and previous steps, are referred to as *reasoning steps*.

2.3.1 Description Step

For description steps, we assess their correctness and relevance.

Description Correctness. A step is labeled as “Fully Correct” if it contains no incorrect information about the image. If all the information in the step is incorrect, it is marked as “Unsupported.” If a step contains both correct and incorrect information, it is labeled as “Partially Correct.” Additionally, we further categorize error types for “Partially Correct” or “Unsupported” steps based on Huang et al. (2023)—Entity False: Incorrectly identifying the presence or type of entities; Attribute False: Incorrect description of entity attributes (e.g., color, shape, texture, count, state, or text recognition); Spatial Relationship False: Incorrect description of spatial relationships between entities (e.g., left, right); Non-Spatial Relationship False: Incorrect description of non-spatial relationships (e.g., speak-

ing to).

Description Relevance. A step is labeled “Image Relevant” if its description pertains to the content of the image. If the description is relevant to answering the question, it is labeled “Logic Relevant”. If it is relevant to both the image and the question, the step is marked as “Both”. If it is relevant to neither, the step is labeled “None”.

2.3.2 Reasoning Step

For reasoning steps, we focus on Logical Correctness, Logical Relevance, and Informativeness.

Logical Correctness. A reasoning step is considered “Correct” if it is logically deduced from the facts and previous steps. A step is labeled “Incorrect” if it contains logical errors or contradicts the facts. Additionally, following Prasad et al. (2023), we further categorize error types for logically incorrect steps—Inter-step Incorrect: When the step cannot be logically inferred or contradicts previous steps; Intra-step Incorrect: When there are internal contradictions within the step; Both: When both inter-step and intra-step errors are present.

Logical Relevance. A step is labeled as “Relevant” if it contributes to answering the question; otherwise, it is marked as “Irrelevant”.

Informativeness. A step is considered “Informative” if it introduces new information and does not repeat or provide redundant details from previous steps (Chen et al., 2023a; Prasad et al., 2023). Otherwise, it is labeled as “Uninformative”.

2.3.3 Full MCoT Sequences

When evaluating an MCoT, we also assess its overall correctness. An MCoT is considered high-quality (Good MCoT) if all description steps are correct and relevant, and all reasoning steps are correct, relevant, and informative (Li et al., 2023b; Jacovi et al., 2024). If these conditions are not met, it is classified as low-quality (Bad MCoT).

3 Dataset Annotation Process

In this section, we introduce the annotation process for the MiCEval dataset, which encompasses: *type of step*, *correctness of step*, and *correctness of MCoT*. The complete data annotation process is illustrated in Figure 3. Three annotation tasks are performed sequentially, with each task building on the results of the preceding one.

Stage 1: Type of Step. In this stage, annotators will label the MCoT answers step-by-step according to the definitions of step types provided in

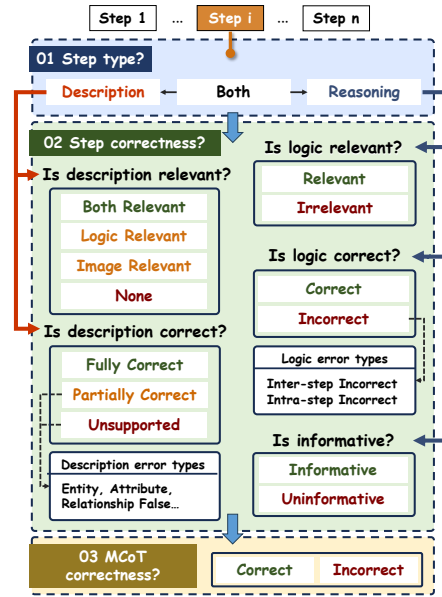


Figure 3: A complete flowchart of the MCoT annotation process. We first determine the type of each step and then annotate its correctness based on the type of step. Once all steps in an MCoT answer are annotated, we evaluate the correctness of the entire MCoT.

Sec. 2. Following this, the annotation proceeds to the corresponding labeling tasks based on the type of each step.

Stage 2: Correctness of Step. We annotate each step from different perspectives based on the results of Stage 1. For description steps, we label correctness and relevance. For reasoning steps, we label logical correctness, relevance, and informativeness. For steps classified as both description and reasoning, we annotate all aspects of each.

When the previous steps contain incorrect information, the annotators need to make additional judgments about the current step. If the current step logically follows or references the incorrect information from the previous steps, it should also be labeled as “Incorrect”. For more details, refer to the annotation illustrated in Figure 9.

Stage 3: Correctness of MCoT. For MCoT, we evaluate the correctness of the MCoT answer. An MCoT answer is considered high-quality (Correct MCoT answer) when every step in the reasoning chain is a Correct Step. This means that each step, according to its respective type, is annotated as Correct, Relevant, or Informative, with no attribute marked as erroneous based on the criteria defined earlier. If any step contains an attribute that is annotated with an error label, the MCoT answer is deemed low-quality (Incorrect MCoT answer).

4 Data Collection

In this section, we present a comprehensive explanation of the construction process for the MiCEval dataset. The full data collection process is presented in Figure 2, which consists of three stages: *question collection for the dataset*, *MCoT answer collection*, and *human annotation*.

4.1 Sampling and Generation

We randomly sampled 700 questions from 8 multi-modal datasets. Two authors manually conducted coarse-grained annotations based on the difficulty of the questions and source datasets to get HARD and NORMAL splits. We used four representative MLLMs to generate MCoT answers. Ultimately, we obtained 1,000 MCoT answers generated by the MLLMs. Detailed descriptions are provided in the Appendix A.2.

4.2 High-Quality Annotation Protocol

It is important to note that the requirement for fine-grained and multi-task annotations necessitates that annotators have relevant expertise. Therefore, all annotators possess the necessary background knowledge related to the annotation tasks.

Preliminary Training. To ensure high-quality data, four of the authors annotated the entire MCoTs (each CoT has three annotations). Previous to the generation of the official annotation of the MCoT answers in the dataset, all annotators underwent preliminary training. We randomly selected 30 questions from the remaining 800 questions not included in the dataset in Sec. A.2 and used GPT-4o and InstructBLIP to generate a total of 60 MCoT answers. The annotators then completed the annotation tasks on these 60 MCoT answers, and based on the annotation results, the annotators engaged in discussions with two additional authors, to reach an internal consensus on some of the more challenging cases.

Human Annotation. Each MCoT answer underwent annotation based on the tasks outlined in Sec. 3. We established a rigorous process to filter for valid data. If any annotation for a task on a given step could not be determined through majority voting (i.e., a 1:1:1 or 1:1 tie), the CoT was considered invalid. In the end, we obtained 903 valid MCoTs. More details can be found in the App. B.2. We computed Bennett, Alpert, and Goldstein’s S inter-rater agreement for the MiCEval-HARD and MiCEval-NORMAL, with scores of 0.888 and

| Dataset(Eval) | Hard | Normal |
|--------------------------------|-------|--------|
| Question | 323 | 320 |
| MCoT Answer | 457 | 446 |
| MCoT Step | 1,745 | 1,144 |
| Avg. step per MCoT | 3.8 | 2.6 |
| Description Step | 899 | 852 |
| Reasoning Step | 769 | 288 |
| Description Fully Correct Step | 679 | 745 |
| Logic Fully Correct Step | 515 | 201 |
| Fully Correct MCoT | 185 | 307 |

Table 1: Statistics on MiCEval-NORMAL and HARD.

0.795, respectively. The S agreement for the description step and reasoning step annotations was 0.877 and 0.859, respectively.

4.3 Data Statistics

We provide the statistics of MiCEval-NORMAL and MiCEval-HARD in Table 1. MiCEval-NORMAL contains 1,144 MCoT steps and MiCEval-HARD contains 1,745 MCoT steps for evaluation. We observe significant differences between the two splits in the total number of reasoning steps, the average of MCoT steps per MCoT answer, and the number of fully correct MCoT answer. These differences reflect that our splits based on question difficulty are meaningful. Further analysis of the dataset can be found in the Appendix A.3.

5 THE MICEVAL FRAMEWORK

The MiCEval framework also introduces two evaluation metrics: one at the step-level and the other at the MCoT-level. The dataset and metrics are designed based on the correctness of the various MCoT step types.

5.1 Evaluation of Step Correctness

We propose two methods for measuring step correctness in MCoT answers.

Description Step Correctness. Our goal is to instruct the MLLM to generate a correctness score for a description step r_i using a prompt-based input $(\mathcal{I}, \mathcal{Q}, \mathcal{RC}_i)$, where $\mathcal{RC}_i = \{r_1, r_2, \dots, r_i\}$. Based on the definitions provided in Sec. 3, the following metric is derived:

$$S_{d_correct}, S_{d_relevant} = M_{prompt}(\mathcal{I}, \mathcal{Q}, \mathcal{RC}_i)$$

$$Correctness_D^{(i)} = S_{d_correct} \odot S_{d_relevant}$$

where M_{prompt} represents the prompt-based MLLM, $S_{d_correct}$ refers to the description correctness score, and $S_{d_relevant}$ represents the description

relevance score. \odot indicates the geometric mean operation.

Reasoning Step Correctness. Similarly, using a prompt-based input $(\mathcal{I}, \mathcal{Q}, \mathcal{RC}_i)$, we guide the MLLM for generating the correctness score for the reasoning step r_i . The corresponding metric is defined as:

$$S_{l_correct}, S_{l_relevant}, S_{info} = M_{prompt}(\mathcal{I}, \mathcal{Q}, \mathcal{RC}_i)$$

$$Correctness_R^{(i)} = S_{l_correct} \odot S_{l_relevant} \odot S_{info}$$

where M_{prompt} represents the prompt-based MLLM, $S_{l_correct}$ denotes the logical correctness score, $S_{l_relevant}$ refers to the logical relevance score, and S_{info} captures the informativeness score.

5.2 Evaluation of MCoT Correctness

MCoT Correctness. We compute the overall score for an MCoT answer by generating a score for each step using the prompt-based MLLM. First, we need to obtain the type of each step based on MLLMs:

$$Type^{(i)} = M_{prompt}(\mathcal{I}, \mathcal{Q}, \mathcal{RC}_i)$$

Then, based on the type of each step, we calculate its correctness score and the score of the entire MCoT as follows:

$$Correctness^{(i)} = \begin{cases} Correctness_D^{(i)}, & \text{description step} \\ Correctness_R^{(i)}, & \text{reasoning step} \end{cases}$$

$$Correctness_{type} = \sqrt[n]{\prod_{i=1}^n Correctness^{(i)}}$$

where $Correctness_D^{(i)}$ represents the correctness score for the description step r_i and $Correctness_R^{(i)}$ represents the correctness score for the reasoning step r_i . The overall MCoT score is computed as the geometric mean of all step correctness scores.

Given the limitations of current research on MCoT step type classification, we also propose another method for calculating correctness that is not based on step type. This method involves calculating the description step correctness and reasoning step correctness for each step, and then deriving the overall score for the entire MCoT:

$$Correctness^{(i)} = Correctness_D^{(i)} \odot Correctness_R^{(i)}$$

$$Correctness_{all} = \sqrt[n]{\prod_{i=1}^n Correctness^{(i)}}$$




Figure 4: The relationship between the average accuracy of three MLLMs across all Pairwise Comparison tasks and the number of shots on two splits.

where $Correctness_D^{(i)}$ represents the correctness score for the step r_i and $Correctness_R^{(i)}$ represents the correctness score for the step r_i , \odot denotes the geometric mean operation. The overall MCoT score is computed as the geometric mean of all step correctness scores.

6 Experiments

In the proposed MiCEval framework, we conducted MCoT verification and MCoT evaluation experiments on existing MLLMs at both the step-level and MCoT-level.

6.1 MLLM-as-a-Verifier

Pairwise comparison. In this experiment, we integrate MLLM-as-a-verifier into the MiCEval framework to validate fine-grained tasks. The labels generated by the MLLMs serve as predictions, which are subsequently compared to human annotations to assess accuracy. Our evaluation focuses on two key aspects: the correlation between MLLM outputs and human annotations, and the effectiveness of MLLMs as MCoT verifiers.

6.1.1 Experimental Setting

We assess the performance of seven mainstream MLLMs as MCoT verifiers in both zero-shot and few-shot settings. Details of the experimental setting and metrics can be found in Appendix C.1.

6.1.2 Results

The results are presented in Tables 2 and 9. The results of the few-shot preliminary experiments are shown in Figure 4, while the results of the exploration experiments for few-shot types and more results can be found in Appendix C.3.

▷ *Few-shot does not always outperform zero-shot.* On MiCEval-NORMAL, only LLaVA-1.6-Mistral-7B demonstrated improvement, with a few-shot

| Evaluation Setting | Model | Step Type | Description | | | Logic | | | | MCoT | Avg. |
|--------------------|-------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|-----------------|--------------|--------------|--------------|
| | | | Relevance | Correctness | Error Types | Relevance | Correctness | Informativeness | Error Types | | |
| Zero-Shot | LLaVA-1.6-Mistral-7B | 0.653 | 0.456 | 0.203 | 0.175 | 0.901 | 0.487 | 0.922 | 0.354 | 0.379 | 0.503 |
| | MiniCPM-V-2.6 | 0.121 | 0.693 | 0.702 | 0.248 | 0.918 | 0.664 | 0.837 | 0.836 | 0.564 | 0.620 |
| | Llama-3.2-11B-Vision-Instruct | 0.137 | 0.806 | 0.028 | 0.309 | 0.851 | 0.541 | 0.375 | 0.534 | 0.703 | 0.493 |
| | LLaVA-1.6-Yi-34B | 0.348 | 0.575 | 0.704 | 0.443 | 0.949 | 0.693 | 0.917 | 0.690 | 0.508 | 0.647 |
| | GPT-4o | 0.672 | 0.768 | 0.631 | 0.545 | 0.903 | 0.649 | 0.732 | 0.765 | 0.726 | 0.710 |
| | Qwen-VL-Max | 0.433 | 0.853 | 0.630 | 0.455 | 0.931 | 0.691 | 0.910 | 0.227 | 0.580 | 0.634 |
| | Gemini-1.5-Pro | 0.868 | 0.838 | 0.674 | 0.398 | 0.926 | 0.584 | 0.330 | 0.827 | 0.670 | 0.679 |
| Few-Shot | LLaVA-1.6-Mistral-7B | 0.452 | 0.310 | 0.255 | 0.371 | 0.934 | 0.688 | 0.812 | 0.863 | 0.584 | 0.585 |
| | MiniCPM-V-2.6 | 0.104 | 0.413 | 0.719 | 0.401 | 0.873 | 0.679 | 0.764 | 0.811 | 0.547 | 0.590 |
| | Llama-3.2-11B-Vision-Instruct | 0.044 | 0.257 | 0.601 | 0.004 | 0.638 | 0.670 | 0.069 | 0.220 | 0.376 | 0.320 |
| | LLaVA-1.6-Yi-34B | 0.703 | 0.206 | 0.638 | 0.488 | 0.939 | 0.670 | 0.927 | 0.751 | 0.716 | 0.671 |
| | GPT-4o | 0.871 | 0.455 | 0.642 | 0.687 | 0.891 | 0.671 | 0.784 | 0.773 | 0.742 | 0.724 |
| | Qwen-VL-Max | 0.359 | 0.262 | 0.603 | 0.370 | 0.924 | 0.776 | 0.749 | 0.650 | 0.689 | 0.598 |
| | Gemini-1.5-Pro | - | - | - | - | - | - | - | - | - | - |

Table 2: The overall performance on Pairwise Comparison of different MLLMs in MiCEval-HARD. Due to limited funding, Gemini-1.5-Pro was not evaluated in the few-shot setting.

average accuracy of 0.657 compared to 0.575 in the zero-shot setting. The other MLLMs, however, showed performance declines, with Llama-3.2-11B-Vision-Instruct experiencing the most significant drop. A detailed analysis of Llama-3.2’s performance is provided in Appendix C.4. On MiCEval-HARD, few-shot performance improved for LLaVA-1.6-7B, LLaVA-1.6-Yi-34B, and GPT-4o, while MiniCPM-V-2.6 and Gemini-1.5-Pro experienced declines.

▷ *MLLMs still exhibit shortcomings in handling complex reasoning tasks.* In comparison to the logical relevance and informativeness tasks on MiCEval-HARD, the average accuracy of most MLLMs across other tasks, in both zero-shot and few-shot settings, remains below 0.7, with a few notable exceptions: MiniCPM achieved 0.702 (zero-shot) and 0.719 (few-shot) in description correctness, while GPT-4o scored 0.726 (zero-shot) and 0.742 (few-shot) on the MCoT task. Most MLLMs demonstrated a low correlation with human judgments in description correctness and error type, indicating substantial room for improvement in the visual modality of current MLLMs.

▷ *The number of shots in few-shot settings does not always correlate with better model performance.* As the number of shots increased, the performance of LLaVA-1.6-7B declined from 0.500 to 0.393 on MiCEval-HARD and from 0.394 to 0.303 on MiCEval-NORMAL. The other two MLLMs exhibited minimal changes in performance. Consequently, we selected 1-shot for all subsequent multimodal few-shot experiments.

6.2 MLLM-as-a-Evaluator

In this experiment, we have two distinct evaluation tasks: *scoring evaluation* and *choice ranking* (Zheng et al., 2023; Chen et al., 2024a).

Scoring Evaluation. In this experiment, we utilize

MLLMs as MiCEval evaluators to score the steps of MCoT answers on a scale from 0 to 10. Human annotations are mapped to 0 (“Incorrect”), 1 (“Correct”), or 0.5 (only for the “Partially Correct” label in the description correctness task). We then compute the correlation between human annotations and the scores generated by the MiCEval evaluators to assess the effectiveness of MiCEval metrics in detecting the quality of MCoT responses. The prompt templates used in this experiment can be found in App. D.2.

Choice Ranking. Using the common split adopted by four MLLMs in Sec. A.2, we constructed a high-quality MCoT selection dataset consisting of 70 questions and 257 MCoT answers. Details about Choice Ranking are presented in the Appendix.

6.2.1 Experimental Setting

We conduct correlation evaluations based on two splits of MiCEval: MiCEval-HARD, MiCEval-NORMAL. Details of the experimental setting and metrics can be found in Appendix D.1.

6.2.2 Results

Table 3 presents the performance of various methods on the *scoring evaluation*. Results for *choice ranking* across different methods are shown in Table 4. For additional analysis, please refer to Appendix D.3.

▷ *Existing metrics are not well-suited for evaluating MCoT answers.* MLLM-based methods show a stronger correlation with human judgment in MCoT answers compared to existing approaches. Among current metrics, LLM-Score achieves the highest performance on the entire MiCEval dataset, with a Somer’s-D score of 0.162. In general, the MLLM-based methods outperform LLM-Score, with the exception of LLaVA-1.6-Mistral-7B.

▷ *MiCEval brings the MLLM evaluator closer to human preferences.* The evaluation metric

| Metric | Normal | Hard | MiCEval |
|-----------------------------------|--------------|--------------|--------------|
| CLIP | 0.079 | 0.019 | 0.060 |
| BLIP2-ITM | 0.088 | -0.065 | 0.031 |
| BLIP2-ITC | 0.072 | -0.006 | 0.075 |
| ReCEval | -0.006 | 0.015 | 0.040 |
| LLM-Score | 0.078 | 0.210 | 0.162 |
| MiniCPM-V-2.6 (1) | 0.154 | 0.123 | 0.130 |
| LLaVA-1.6-Mistral-7B (1) | 0.109 | -0.024 | 0.080 |
| Llama-3.2-11B-Vision-Instruct (1) | 0.090 | 0.218 | 0.176 |
| GPT-4o (1) | 0.154 | 0.256 | 0.208 |
| MiniCPM-V-2.6 (2) | 0.112 | 0.123 | 0.169 |
| LLaVA-1.6-Mistral-7B (2) | 0.083 | 0.178 | 0.140 |
| Llama-3.2-11B-Vision-Instruct (2) | 0.031 | 0.090 | 0.075 |
| GPT-4o (2) | 0.178 | 0.282 | 0.257 |
| MiniCPM-V-2.6 (3) | 0.264 | 0.188 | 0.273 |
| LLaVA-1.6-Mistral-7B (3) | 0.081 | 0.133 | 0.121 |
| Llama-3.2-11B-Vision-Instruct (3) | 0.157 | 0.159 | 0.186 |
| GPT-4o (3) | 0.229 | 0.255 | 0.265 |
| MiniCPM-V-2.6 (4) | 0.270 | 0.191 | 0.277 |
| LLaVA-1.6-Mistral-7B (4) | 0.083 | 0.141 | 0.129 |
| Llama-3.2-11B-Vision-Instruct (4) | 0.156 | 0.172 | 0.194 |
| GPT-4o (4) | 0.245 | 0.272 | 0.284 |

Table 3: Somer’s-D scores of different evaluation metrics on two splits of MiCEval. (1) A holistic evaluation of the entire MCoT without the MiCEval; (2) “Step-by-Step” evaluation without MiCEval; (3) “Step-by-Step” evaluation with $Correctness_{type}$ -based MiCEval; (4) “Step-by-Step” evaluation with $Correctness_{all}$ -based MiCEval.

based on the description and reasoning correctness of MiCEval shows better performance and more closely aligns with human judgments compared to evaluations without MiCEval. We observe that MiCEval-based methods, specifically $Correctness_{type}$ and $Correctness_{all}^{(i)}$ using MiniCPM-V-2.6, Llama-3.2-11B-Vision-Instruct, and GPT-4o, exhibit a stronger correlation in MCoT answer evaluation than the other two MLLM-based approaches. Notably, GPT-4o with $Correctness_{all}^{(i)}$ achieved the highest Somer’s-D score of 0.284 across the entire MiCEval dataset.

▷ *MiCEval helps filtering out high-quality MCoTs.* In the choice ranking task, CLIP achieved the best result with an accuracy of 0.700. The performance of the three MLLM-based methods varied: MiniCPM (CoT, w/o MiCEval) reached an accuracy of only 0.686 and all LLaVA-based evaluation metrics were below 0.700. However, the remaining MLLM-based methods outperformed CLIP. Compared to the other two approaches (without MiCEval), MiniCPM, Llama-3.2, and GPT-4o performed better using the MiCEval metrics evaluation, with GPT-4o and MiniCPM achieving the highest result.

6.3 Analysis

Instruction-following ability of MLLMs. Despite our efforts to guide the model’s predictions within the defined label categories by experimenting with

| Model | Acc. |
|-----------------------------------|--------------|
| CLIP | 0.700 |
| BLIP2-ITM | 0.657 |
| BLIP2-ITC | 0.657 |
| ReCEval | 0.600 |
| LLM-Score | 0.614 |
| MiniCPM-V-2.6 (1) | 0.686 |
| LLaVA-1.6-Mistral-7B (1) | 0.686 |
| Llama-3.2-11B-Vision-Instruct (1) | 0.729 |
| GPT-4o (1) | 0.757 |
| MiniCPM-V-2.6 (2) | 0.700 |
| LLaVA-1.6-Mistral-7B (2) | 0.571 |
| Llama-3.2-11B-Vision-Instruct (2) | 0.757 |
| GPT-4o (2) | 0.771 |
| MiniCPM-V-2.6 (3) | 0.800 |
| LLaVA-1.6-Mistral-7B (3) | 0.600 |
| Llama-3.2-11B-Vision-Instruct (3) | 0.743 |
| GPT-4o (3) | 0.786 |
| MiniCPM-V-2.6 (4) | 0.814 |
| LLaVA-1.6-Mistral-7B (4) | 0.600 |
| Llama-3.2-11B-Vision-Instruct (4) | 0.757 |
| GPT-4o (4) | 0.814 |

Table 4: Accuracy of different evaluation metrics on Choice Ranking. Settings are provided in Table 3.

| Model | Zero Shot | | Few Shot | |
|-------------------------------|---------------|---------------|--------------|--------------|
| | NORMAL | HARD | NORMAL | HARD |
| LLaVA-1.6-Mistral-7B | 1.29% | 3.94% | 2.46% | 8.59% |
| MiniCPM-V-2.6 | 0.00% | 1.08% | 0.53% | 3.49% |
| LLaVA-1.6-Yi-34B | 0.18% | 0.32% | 0.00% | 0.87% |
| GPT-4o | 0.74% | 1.40% | 0.64% | 0.91% |
| Qwen-VL-Max | 14.27% | 18.61% | 0.96% | 5.40% |
| Llama-3.2-11B-Vision-Instruct | 7.48% | 7.12% | 3.08% | 2.87% |
| Gemini-1.5-Pro | 0.25% | 0.19% | - | - |

Table 5: Invalid outputs proportions of each MLLMs in zero shot and few shot evaluations under all tasks of NORMAL and HARD splits.

various instructions, prompt templates, and multiple iterations, a small percentage of invalid outputs persisted due to the current limitations of MLLMs in following instructions. We conducted a statistical analysis of the invalid predictions across all models in the **Pairwise Comparison** experiments. The overall proportion of invalid outputs was 2.93%, with the zero-shot evaluation yielding 3.48% invalid outputs, compared to 2.34% in the few-shot evaluation. This suggests that few-shot evaluation improves the model’s ability to follow instructions. The distribution of invalid output proportions for each model across different settings and datasets is shown in Table 5. For additional analysis, please refer to Appendix C.4 and D.4.

7 Related Work

Reasoning Chain Construction and Evaluations. Chain of Thought (CoT) is an innovative method for solving reasoning tasks using large language models (LLMs) (Wei et al., 2022). Kojima et al. (2022) showed that adding a simple prompt, such as “Let’s think step by step”, significantly improves

LLMs performance in zero-shot settings. Jung et al. (2022) demonstrated that even erroneous reasoning chains can lead to correct final predictions, which has increasingly shifted the focus towards verifying the correctness of the entire reasoning chain (Welleck et al., 2022; Chen et al., 2024b). To address the challenge of automatically generating good quality CoTs, one idea is to connect explicit knowledge, such as knowledge graphs (KG) (Pan et al., 2017b,a), and parametric knowledge from LLMs (Pan et al., 2023). Wu et al. (2024) proposed to generate CoTs from relevant knowledge graphs (Huang et al., 2024). Wang et al. (2024c) proposed to use knowledge graph patterns, such as those widely used for query answering over KGs, to generate effective CoT, without having to rely on the content of KGs. Golovneva et al. (2022) introduced a comprehensive set of CoT evaluation metrics, validating CoT’s logical correctness, grammar, and informativeness. Prasad et al. (2023) evaluates the informativeness and logical correctness of CoTs, while Jacovi et al. (2024) focuses on verifying both attribution and logical correctness of CoT answers.

Broadly speaking, reasoning is also related to text entailments (Pekar, 2006; Saadat-Yazdi et al., 2023), entailment graphs (Hosseini et al., 2018; Li et al., 2024; Zhou et al., 2024), NL2SQL (Vougiouklis et al., 2023; Shen et al., 2024; Zheng et al., 2024), as well as planning (Wang et al., 2024b; Vyas et al., 2025) and API (Wang et al., 2024a).

Evaluation based on MLLMs. Recent works leverage LLM’s instruction-following capabilities as an evaluator (Fu et al., 2024; Zheng et al., 2023; Liu et al., 2023b). Huang et al. (2023) used MLLMs as evaluators to assess the degree of alignment between images and text. Chen et al. (2024a) explored the correlation between MLLM judgments as evaluators and human assessments.

8 Conclusion

In this paper, we introduce MiCEval, a novel automated evaluation framework designed to assess the correctness of MCoT and evaluate the capabilities of MLLMs in judging the quality of different reasoning steps. We create a multilevel human-annotated MCoT dataset and conduct three distinct experiments to evaluate the alignment between MLLMs’ judgments and human agreement: Pairwise Comparison, Scoring Evaluation, and Choice Ranking. The experimental results demonstrate

that MiCEval metrics align more closely with human preferences. Additionally, current MLLMs show notable weaknesses in both visual and language modalities, especially in handling complex reasoning tasks, highlighting an area in need of further improvement.

Acknowledgement

The computations detailed in this research were facilitated by the Edinburgh International Data Facility (EIDF) and the Data-Driven Innovation Programme at the University of Edinburgh.

Limitation

Gemini-1.5-Pro under few-shot evaluation. Compared to zero-shot evaluation, the token consumption in few-shot evaluation is approximately four times higher. In the zero-shot evaluation, the cost of using Gemini-1.5-Pro is twenty times that of GPT-4o. Due to budget constraints, we were unable to conduct a few-shot evaluation for Gemini-1.5-Pro.

Baseline in Pairwise Comparison. Due to the limitations of the Pairwise Comparison outputs, it is challenging to map the existing method’s outputs to 0 and 1 without setting a threshold. However, introducing a threshold could significantly affect the accuracy of the evaluation. Additionally, given the wide range of tasks in Pairwise Comparison, the current method is unable to effectively handle all tasks. Therefore, we did not test other baselines on Pairwise Comparison.

Step type classification. In the current zero-shot evaluation, the two best performing MLLMs in the step type task are GPT-4o and Gemini-1.5-Pro. Therefore, in our experiments, we used Gemini-1.5-Pro as the classifier for step type. From the experimental results, it is evident that the MLLM’s classification of step type affects the MiCEval performance based on the $Correctness_{type}$. This is an issue we need to address in future work.

The instruction-following abilities of current MLLMs. The instruction-following abilities of current MLLMs still have limitations and are not yet able to handle complex instruction tasks effectively. Despite our best efforts to improve the performance of each model in the evaluation, Qwen-VL-Max and Llama-3.2-11B-Vision-Instruct produced more invalid outputs compared to other models.

References

- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.
- Dongping Chen, Ruoxi Chen, Shilin Zhang, Yaochen Wang, Yinuo Liu, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. 2024a. **MLLM-as-a-judge: Assessing multimodal LLM-as-a-judge with vision-language benchmark**. In *Forty-first International Conference on Machine Learning*.
- Hanjie Chen, Faeze Brahman, Xiang Ren, Yangfeng Ji, Yejin Choi, and Swabha Swayamdipta. 2023a. **REV: Information-theoretic evaluation of free-text rationales**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2007–2030, Toronto, Canada. Association for Computational Linguistics.
- Justin Chen, Swarnadeep Saha, and Mohit Bansal. 2024b. **ReConcile: Round-table conference improves reasoning via consensus among diverse LLMs**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7066–7085, Bangkok, Thailand. Association for Computational Linguistics.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. 2024c. Are we on the right way for evaluating large vision-language models? *arXiv preprint arXiv:2403.20330*.
- Qiguang Chen, Libo Qin, Jin Zhang, Zhi Chen, Xiao Xu, and Wanxiang Che. 2024d. **M³cot: A novel benchmark for multi-domain multi-step multi-modal chain-of-thought**. In *Proc. of ACL*.
- Yiming Chen, Chen Zhang, Danqing Luo, Luis Fernando D’Haro, Robby Tan, and Haizhou Li. 2024e. **Unveiling the achilles’ heel of NLG evaluators: A unified adversarial framework driven by large language models**. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 1359–1375, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Zhuo Chen, Jiaoyan Chen, Wen Zhang, Lingbing Guo, Yin Fang, Yufeng Huang, Yichi Zhang, Yuxia Geng, Jeff Z. Pan, Wenting Song, and Huajun Chen. 2023b. **MEAformer: Multi-modal Entity Alignment Transformer for Meta Modality Hybrid**. In *In Proc. of the 31st ACM International Conference on Multimedia (MM 2023)*.
- Zhuo Chen, Lingbing Guo, Yin Fang, Yichi Zhang, Jiaoyan Chen, Jeff Z. Pan, Yangning Li, Huajun Chen, and Wen Zhang. 2023c. **Rethinking Uncertainly Missing and Ambiguous Visual Modality in Multi-Modal Entity Alignment**. In *In the 22nd International Semantic Web Conference (ISWC 2023)*, pages 121–139.
- Miruna-Adriana Clinciu, Arash Eshghi, and Helen Hastie. 2021. **A study of automatic metrics for the evaluation of natural language explanations**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2376–2387, Online. Association for Computational Linguistics.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. **Instructblip: Towards general-purpose vision-language models with instruction tuning**. *Preprint*, arXiv:2305.06500.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, , and et al. 2024. **The llama 3 herd of models**. *Preprint*, arXiv:2407.21783.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2024. **GPTScore: Evaluate as you desire**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6556–6576, Mexico City, Mexico. Association for Computational Linguistics.
- Jingying Gao, Qi Wu, Alan Blair, and Maurice Pagnucco. 2023. **Lora: A logical reasoning augmented dataset for visual question answering**. In *Advances in Neural Information Processing Systems*, volume 36, pages 30579–30591. Curran Associates, Inc.
- Olga Golovneva, Moya Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. 2022. **ROSCOE: A suite of metrics for scoring step-by-step reasoning**. In *In The Eleventh International Conference on Learning Representations*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. **Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering**. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. **Vizwiz grand challenge: Answering visual questions from blind people**. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hangfeng He, Hongming Zhang, and Dan Roth. 2024. **SocREval: Large language models with the socratic method for reference-free reasoning evaluation**. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2736–2764, Mexico City, Mexico. Association for Computational Linguistics.
- Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. **CLIPScore: A**

- reference-free evaluation metric for image captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mohammad Javad Hosseini, Nathanael Chambers, Siva Reddy, Xavier R. Holt, Shay B. Cohen, Mark Johnson, and Mark Steedman. 2018. [Learning typed entailment graphs with global soft constraints](#). *Trans. Assoc. Comput. Linguistics*, 6:703–717.
- Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. 2023. [T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 78723–78747. Curran Associates, Inc.
- Wenyu Huang, Guancheng Zhou, Hongru Wang, Pavlos Vougiouklis, Mirella Lapata, and Jeff Z. Pan. 2024. Less is More: Making Smaller Language Models Competent Subgraph Retrievers for Multi-hop KGQA. In *In the Findings of the Association for Computational Linguistics: EMNLP 2024 (EMNLP 2024 Findings)*, pages 15787–15803.
- Alon Jacovi, Yonatan Bitton, Bernd Bohnet, Jonathan Herzig, Or Honovich, Michael Tseng, Michael Collins, Roei Aharoni, and Mor Geva. 2024. [A chain-of-thought is as strong as its weakest link: A benchmark for verifiers of reasoning chains](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4615–4634, Bangkok, Thailand. Association for Computational Linguistics.
- Jaehun Jung, Lianhui Qin, Sean Welleck, Faeze Brahman, Chandra Bhagavatula, Ronan Le Bras, and Yejin Choi. 2022. [Maieutic prompting: Logically consistent reasoning with recursive explanations](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1266–1279, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.
- Juncai Li, Ru Li, Xiaoli Li, Qinghua Chai, and Jeff Z. Pan. 2024. Inference Helps PLMs Conceptual Understanding: Improving the Abstract Inference Ability with Hierarchical Conceptual Entailment Graphs. In *In Proc. of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, pages 22088–22104.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*.
- Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. 2023b. [Making language models better reasoners with step-aware verifier](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5315–5333, Toronto, Canada. Association for Computational Linguistics.
- Fangyu Liu, Guy Edward Toh Emerson, and Nigel Collier. 2023a. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruochen Xu, and Chenguang Zhu. 2023b. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*.
- Qing Lyu, Shreya Havaladar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. [Faithful chain-of-thought reasoning](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 305–329, Nusa Dua, Bali. Association for Computational Linguistics.
- OpenAI, : Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, and et al. 2023. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Jeff Z. Pan, Simon Razniewski, Jan-Christoph Kalo, Sneha Singhanian, Jiaoyan Chen, Stefan Dietze, Hajira Jabeen, Janna Omelinyanenko, Wen Zhang, Matteo Lissandrini, Russa Biswas, Gerard de Melo, Angela Bonifati, Edlira Vakaj, Mauro Dragoni, and Damien Graux. 2023. Large Language Models and Knowledge Graphs: Opportunities and Challenges. pages 1–38.
- J.Z. Pan, D. Calvanese, T. Eiter, I. Horrocks, M. Kifer, F. Lin, and Y. Zhao, editors. 2017a. *Reasoning Web: Logical Foundation of Knowledge Graph Construction and Querying Answering*. Springer.
- J.Z. Pan, G. Vetere, J.M. Gomez-Perez, and H. Wu, editors. 2017b. *Exploiting Linked Data and Knowledge Graphs for Large Organisations*. Springer.

- Viktor Pekar. 2006. [Acquisition of verb entailment from text](#). In *HLT-NAACL*. The Association for Computational Linguistics.
- Archiki Prasad, Swarnadeep Saha, Xiang Zhou, and Mohit Bansal. 2023. [ReCEval: Evaluating reasoning chains via correctness and informativeness](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10066–10086, Singapore. Association for Computational Linguistics.
- Ameer Saadat-Yazdi, Jeff Z. Pan, and Nadin Kokciyan. 2023. [Uncovering Implicit Inferences for Improved Relational Argument Mining](#). In *In the 17th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2023)*, pages 2484–2495.
- Abulhair Saparov and He He. 2023. [Language models are greedy reasoners: A systematic formal analysis of chain-of-thought](#). In *International Conference on Learning Representations*.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024. [Visual cot: Unleashing chain-of-thought reasoning in multi-modal language models](#). *Preprint*, arXiv:2403.16999.
- Zhili Shen, Pavlos Vougiouklis, Chenxin Diao, Kaustubh Vyas, Yuanyi Ji, and Jeff Z. Pan. 2024. [Improving Retrieval-augmented Text-to-SQL with AST-based Ranking and Schema Pruning](#). In *In Proc. of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, pages 7865–7879.
- Robert H. Somers. 1962. [A new asymmetric measure of association for ordinal variables](#). *American Sociological Review*, 27:799.
- C. Spearman. 1904. [The proof and measurement of association between two things](#). *The American Journal of Psychology*, 15(1):72–101.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, Soroosh Mariooryad, Yifan Ding, Xinyang Geng, Fred Alcober, and et al. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. [Eyes wide shut? exploring the visual shortcomings of multimodal llms](#). *Preprint*, arXiv:2401.06209.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. 2023. [Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 74952–74965. Curran Associates, Inc.
- Pavlos Vougiouklis, Nikos Papsarantopoulos, Danna Zheng, David Tuckey, Chenxin Diao, Zhili Shen, and Jeff Z. Pan. 2023. [FastRAT: Fast and Efficient Cross-lingual Text-to-SQL Semantic Parsing](#). In *In Proc. of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (JCNLP-AACL 2023)*.
- Kaustubh Vyas, Damien Graux, Yijun Yang, Sébastien Montella, Chenxin Diao, Wendi Zhou, Pavlos Vougiouklis, Ruofei Lai, Yang Ren, Keshuang Li, and Jeff Z. Pan. 2025. [From An LLM Swarm To A PDDL-Empowered HIVE: Planning Self-Executed Instructions In A Multi-Modal Jungle](#). In *In Proc. of the International Conference on Learning Representations (ICLR 2025)*.
- Hongru Wang, Yujia Qin, Yankai Lin, Jeff Z. Pan, and Kam-Fai Wong. 2024a. [Empowering Large Language Models: Tool Learning for Real-world Interaction](#). In *In Proc. of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24)*, pages 2983 – 2986.
- Hongru Wang, Rui Wang, Boyang Xue, Heming Xia, Jingtao Cao, Zeming Liu, Jeff Z. Pan, and Kam-Fai Wong. 2024b. [MetaBench: Planning of Multiple APIs from Various APPs for Complex User Instruction](#). In *In Proc. of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, pages 15322–15336.
- Junjie Wang, Mingyang Chen, Binbin Hu, Dan Yang, Ziqi Liu, Yue Shen, Peng Wei, Zhiqiang Zhang, Jinjie Gu, Jun Zhou, Jeff Z. Pan, Wen Zhang, and Hua-jun Chen. 2024c. [Learning to Plan for Retrieval-Augmented Large Language Models from Knowledge Graphs](#). In *In the Findings of the Association for Computational Linguistics: EMNLP 2024 (EMNLP 2024 Findings)*, pages 7813–7835.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Sean Welleck, Jiacheng Liu, Ximing Lu, Hannaneh Hajishirzi, and Yejin Choi. 2022. [Naturalprover: Grounded mathematical proof generation with language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 4913–4927. Curran Associates, Inc.
- Yike Wu, Yi Huang, Nan Hu, Yuncheng Hua, Guilin Qi, Jiaoyan Chen, and Jeff Z. Pan. 2024. [CoTKR: Chain-of-Thought Enhanced Knowledge Rewriting for Complex Knowledge Graph Question Answering](#). In *In Proc. of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, pages 3501–3520.

- Shijie Xia, Xuefeng Li, Yixin Liu, Tongshuang Wu, and Pengfei Liu. 2024. Evaluating mathematical reasoning beyond accuracy. *arXiv preprint arXiv:2404.05692*.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2024. Mm-vet: Evaluating large multimodal models for integrated capabilities. In *International conference on machine learning*. PMLR.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2024. [Multi-modal chain-of-thought reasoning in language models](#). *Preprint*, arXiv:2302.00923.
- Danna Zheng, Mirella Lapata, and Jeff Z. Pan. 2024. Archer: A Human-Labeled Text-to-SQL Dataset with Arithmetic, Commonsense and Hypothetical Reasoning. In *In Proc. of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2024)*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Preprint*, arXiv:2306.05685.
- Wendi Zhou, Tianyi Li, Pavlos Vougiouklis, Mark Steedman, and Jeff Z. Pan. 2024. A Usage-centric Take on Intent Understanding in E-Commerce. In *In Proc. of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024)*, pages 228–236.
- Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. 2016. Visual7W: Grounded Question Answering in Images. In *IEEE Conference on Computer Vision and Pattern Recognition*.

A MiCEval Dataset

A.1 Source Datasets

Detailed information of the all source datasets for MiCEval are provided in Table 13. The distribution of source datasets are shown in Figure 6, and the distribution of MCoT-generators is also provided in Figure 5. Table 14 and Table 15 shows randomly sampled questions from both MiCEval-HARD and MiCEval-NORMAL.




Figure 5: The MCoT generators distribution of each splits.

A.2 Supplement to Data Collection

Split Based on Question Difficulty. 1,500 open-ended questions were randomly sampled from 8 datasets: Visual7W dataset (Zhu et al., 2016), VSR dataset (Liu et al., 2023a), ScienceQA dataset (Lu et al., 2022), VQAv2 dataset (Goyal et al., 2017), Vizwiz dataset (Gurari et al., 2018), MMVP dataset (Tong et al., 2024), MM-Vet dataset (Yu et al., 2024), and MMstar dataset (Chen et al., 2024c). Two authors manually filtered out duplicated questions and conducted coarse-grained annotations based on the difficulty of the questions and source datasets. As a result, the dataset was divided into two splits, HARD and NORMAL, with 350 questions in each split.

MLLMs as MCoT Answer Generators. We used four representative MLLMs to generate MCoT answers: Qwen-VL-Max (Bai et al., 2023), GPT-4o (OpenAI et al., 2023), LLaVA-1.6-Mistral-7B (Liu et al., 2024), and InstructBLIP-Vicuna-7B (Dai et al., 2023). By selecting models with varying performance levels, we aimed to capture a diverse range of CoTs. Following Zhang et al. (2024), we manually crafted several MCoT answer generation prompts tailored to each dataset. Inspired by Jacovi et al. (2024), we also randomly divided each split into five parts, with four MLLMs answering one

| Complexity Error Types | Description Step (N=1,751) | Reasoning Step (N=1,057) | Both Step (N=81) |
|--------------------------------|----------------------------|--------------------------|------------------|
| Entity False | 7.2% | - | 2.5% |
| Attribute False | 10.5% | - | 9.9% |
| Spatial Relationship False | 1.7% | - | 0.0% |
| Non-spatial Relationship False | 0.9% | - | 0.0% |
| Inter-step Incorrect | - | 29.6% | 4.9% |
| Intra-step Incorrect | - | 2.5% | 13.6% |
| Image Irrelevant | 1.6% | - | 0.0% |
| Logic Irrelevant | 4.9% | 4.4% | 12.3% |
| Uninformative | - | 8.7% | 4.9% |

Table 6: The percentage of complexity error types on different MCoT steps.

part independently, and one part being answered by each MLLMs. Our aim was to enhance the dataset’s flexibility and its potential for analysis and evaluation across a wide range of MLLMs. Ultimately, we obtained 1,000 MCoT answers generated by the MLLMs.

A.3 Dataset Analysis

Visual and Language Modalities of MLLMs. In Table 6, we summarize the proportions of fine-grained error categories at the step level. For description steps, 7.2% are classified as entity false, while 10.5% are labeled as attribute false. This suggests that current MLLMs still have room for improvement in accurately extracting information from the visual modality, particularly when detecting entities and their attributes. For reasoning steps, 29.6% contain inter-step incorrect errors, indicating that MLLMs continue to struggle with reasoning, especially in maintaining logical consistency between steps. Although the number of both steps is relatively small compared to description and reasoning steps, the proportion of intra-step incorrect steps reaches 13.6%, suggesting that MLLMs are more prone to internal logical errors when generating mixed steps that combine both visual information and logical reasoning, compared to the other two types of steps.

Single-step MCoT Answers. In Figure 7a, we present the relationship between MCoT correctness and the number of steps in an MCoT. It is observed that MLLM achieves the highest quality when generating single-step MCoTs, with 81.2% being classified as High-Quality MCoTs. However, as the number of steps increases, the proportion of High-Quality MCoTs gradually declines, eventually to 44.8%.

High-quality MCoT Answers & Final Predictions. Inspired by Jung et al. (2022), we aim to analyze the relationship between MCoT correctness

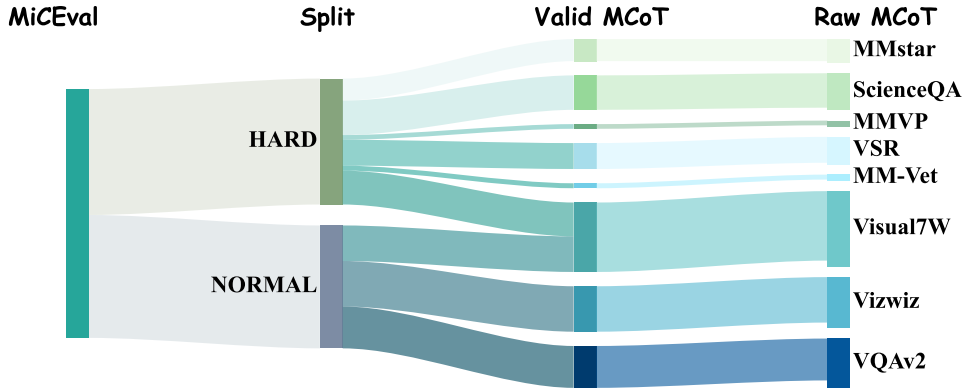


Figure 6: The MCoT generators distribution of each splits.

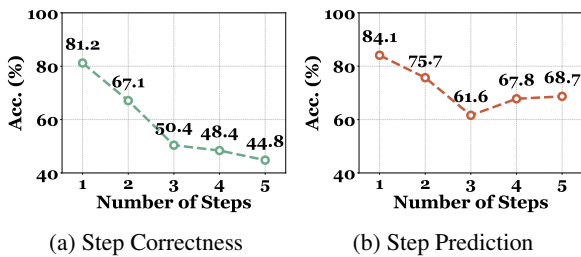


Figure 7: (a) The Relationship between MCoT answers Correctness and the number of steps. (b) The Relationship between Prediction Correctness and the number of steps.

| | Correct Prediction | Incorrect Prediction |
|--------------------------|--------------------|----------------------|
| High-quality MCoT answer | 90.4% | 9.6% |
| Low-quality MCoT answer | 42.8% | 57.2% |

Table 7: The correlation between the correctness of MCoT answer and Prediction. High-quality MCoT answer presents fully correct.

and the final prediction of the MCoT, specifically its answer to the question. To this end, we conducted additional annotations for the MCoT final predictions. For detailed information about this annotation task and inter-rater agreement, please refer to the App. B.3. Based on our annotations of MCoT answer correctness and final prediction correctness, we analyzed their relationship. As shown in Table 7, our results align with the findings of Jung et al. (2022), where high-quality MCoT answers lead to a final prediction accuracy of 90.4%, while 42.8% of MCoTs with errors still produce correct final predictions.

Errors & Final Prediction. We calculated the proportions of fine-grained error categories when the final prediction is incorrect. Table 8 shows that Inter-step Incorrect, Attribute False, and En-


| Complexity Category | Correct Prediction | Incorrect Prediction |
|--------------------------------|--------------------|----------------------|
| Entity False | 3.86% | 26.60% |
| Attribute False | 8.70% | 36.50% |
| Spatial Relationship False | 0.97% | 7.45% |
| Non-spatial Relationship False | 0.81% | 2.48% |
| Inter-step Incorrect | 7.89% | 60.30% |
| Intra-step Incorrect | 1.13% | 9.57% |
| Image Irrelevant | 0.64% | 6.74% |
| Logic Irrelevant | 4.51% | 11.70% |
| Uninformative | 10.60% | 9.57% |

Table 8: The percentage of complexity error types on correct and incorrect predictions.

tity False are more likely to lead to incorrect final predictions in MCoT answer. In particular, the proportion of wrong predictions for MCoT answers with Inter-step Incorrect reaches as high as 60.3%.

Number of Steps & Final Prediction. Figure 7b illustrates the relationship between the number of MCoT steps and the accuracy of the final predictions. We observe that MLLMs achieve their highest accuracy, 84.1%, when generating single-step MCoT answers. As the number of steps increases, the accuracy declines gradually. However, when the number of steps reaches four or more, the accuracy of the final predictions begins to improve. We attribute this improvement to the positive impact of more detailed and comprehensive reasoning, which eventually outweighs the negative effects of accumulating errors from additional steps.

MLLM generators. Figure 8 illustrates the accuracy of MCoT answers and Prediction generated by each MLLM. From these statistics, GPT-4o outperform the other three MLLMs in generating the correct MCoT answers.



(a) MCoT Answer Correctness (b) Prediction Correctness

Figure 8: MCoT answer Correctness and Prediction Correctness across different MCoT generators.

B Annotation

B.1 Detailed Definitions on Annotation Tasks

We provide the labels for each annotation task along with their detailed reference definitions as follows:

1. Step Type:

- **Description:** A step is entailed from the image.
- **Reasoning:** A step is entailed from previous steps.
- **Both:** A step is entailed from both the image and previous steps.

2. Description Correctness:

- **Fully correct:** A step without any incorrect information.
- **Partially correct:** A step contains some incorrect information, but there is still some correct information as well.
- **Unsupported:** All information is incorrect.

3. Description Relevance:

- **Image relevant:** A step is relevant to the image.
- **Logic relevant:** A step is relevant to answering the question.
- **Both:** A step is relevant to both the image and answering the question.
- **None:** A step is irrelevant to both the image and answering the question.

4. Logic Correctness:

- **Correct:** A step can be logically inferred from the previous steps and without any logical errors or conflicts between its internal clauses.

- **Incorrect:** A step cannot be logically inferred from the previous steps or contains logical errors or conflicts between its internal clauses.

5. Logic Relevance:

- **Relevant:** A step is relevant to answering the question.
- **Irrelevant:** A step is not relevant to answering the question.

6. Informativeness:

- **Uninformative:** A step is repetitive or redundant.
- **Informative:** A step without repetition or redundancy.

7. Description Error Types:

- **Entity false:** Some entities mentioned in this step do not exist in the image.
- **Attribute false:** The attributes of an entity are incorrectly described, including color, shape, texture, count, state, text recognition.
- **Spatial Relationship false:** The spatial relationship between two objects is incorrectly described.
- **Non-Spatial Relationship false:** The active, passive, or action relationship between two objects is incorrectly described.

8. Logic Error Type:

- **Inter-step Incorrect:** A step can not be logically inferred from the previous steps.
- **Intra-step Incorrect:** A step contains logical errors or conflicts between its internal clauses.
- **Both:** A step can not be logically inferred from the previous steps and contains logical errors or conflicts between its internal clauses.

9. MCoT Answer:

- **Correct:** The MCoT answer is fully correct, free of any descriptive or logical errors, fully relevant to both the image and the answer, and contains no repetition or redundancy.

- **Incorrect:** The MCoT answer is either irrelevant to the image or the answer, contains descriptive or logically incorrect, or includes repetition or redundancy.

B.2 Valid Data Filtering Process

In our valid data filtering process, we filter based on the results of all annotation tasks in Section C.1, defining invalid steps and invalid MCoTs at both the step-level and MCoT-level, respectively.

Invalid Step

- If the annotation result of a step’s step type is 3:0, we further assess the annotation results of other tasks. If there is one task with an annotation result of 1:1:1, this step is deemed an invalid step; otherwise, it is considered a valid step.
- If a step’s step type is 2:1, meaning each annotation task for step correctness has only two valid annotation results, and if one task has a result of 1:1, then this step is classified as an invalid step.

Invalid MCoT

- If an MCoT answer contains one invalid step, the MCoT is considered invalid data.
- If half or more of the steps in an MCoT answer have a step type annotation result of 2:1, this MCoT is also deemed invalid data.

After filtering, our MiCEval dataset contains 903 valid MCoT answers and 2,889 valid MCoT steps. Figure 9 provides a detailed label of the MiCEval annotated data and Figures 12 to 14 show screenshots of our annotation tool.

B.3 Additional Task: Correctness of Prediction

During the analysis process, we observed that despite providing a complete reasoning chain in the prompt, some MCoT answers generated by MLLMs were incomplete, with the final prediction missing. Additionally, since the questions are open-ended and do not have a single correct answer, we introduced an extra annotation step: We annotate the final prediction of the MCoT answer (the last step of the reasoning chain) based on whether it correctly answered the question. If the prediction is correct, it will be labeled as **Correct**; otherwise, it will be labeled as **Incorrect**. The Bennett, Alpert, and Goldstein’s S inter-rater agreement of this additional task is 0.827.

C Verifier Experiment

C.1 Detail Experimental Setting and Metric

The MLLMs evaluated include open-source MLLMs such as LLaVA-1.6-Mistral-7B, LLaVA-1.6-Yi-34B (Liu et al., 2024), MiniCPM-V-2.6 (Yao et al., 2024), and Llama-3.2-11B-Vision-Instruct (Dubey et al., 2024), as well as API-accessible models like Qwen-VL-Max, GPT-4o, and Gemini-1.5-Pro (Team et al., 2024). For each task, we conducted three trials and reported the average performance across the runs. In the few-shot evaluation, we designed nine different 4-shot prompts for all tasks, ensuring each set was label-balanced. These demonstrations were randomly drawn from the source datasets.

For the few-shot evaluation, we conducted preliminary experiments using a subset of models, specifically LLaVA-1.6-Mistral-7B, MiniCPM-V-2.6, and GPT-4o, representing both open-source and API-accessible MLLMs. We tested these models with varying shot numbers (1-4) and selected the best-performing configuration for the full few-shot evaluation. Inspired by Chen et al. (2024d), we further investigated the performance of these MLLMs across five tasks involving image inputs, examining both multimodal and textual few-shot types.

Given the limitations in MLLMs’ instruction-following capabilities, a small percentage of their predictions fell outside the predefined label set. As a result, we use accuracy as the evaluation metric, treating any predictions that do not match the defined labels as incorrect.

C.2 Prompt Template

In addition to following the prompt templates used in Jacovi et al. (2024)’s work, we also experimented with various prompt templates to identify one that optimally enhances the performance of each MLLM. The specific prompt templates are as follows:

Step Type Task

Image: [image]
 Question: [question]
 Rationale: [rationale]
 Step: [current step]
 Step Type: {Description, Reasoning, Both}

Description Correctness Task

Image: [image]
 Step: [description step]



| Question: What will happen next? | Step Type | Annotation |
|---|-----------|---|
| Step 1: The person in the image appears to be slipping on a wet surface. | Both | Both Relevant, Fully Correct |
| | | Relevant, Correct, Informative |
| Step 2: They are in mid-air, indicating a loss of balance. | Both | Both Relevant, Partially Correct (Attribute False) |
| | | Relevant, Correct, Informative |
| Step 3: Given the trajectory and the wet conditions, it is likely that the person will fall to the ground next. | Reasoning | Relevant, Correct, Informative |
| Step 4: So, the person will fall next. | Reasoning | Relevant, Correct, Informative |

Figure 9: An detailed example of annotation: The Description Correctness of Step 2 is labeled as Partially Correct. However, the subsequent Step 3 and Step 4 do not rely on the incorrect information from Step 2 for their reasoning. Since the logic in these two steps is correct, their Logic Correctness is labeled as Correct.

Output: {Fully Correct, Partially Correct, Unsupported}

Previous: [previous steps]

Step: [incorrect reasoning step]

Output: {Informative, Uninformative}

Description Relevance Task

Image: [image]

Question: [question]

Rationale: [rationale]

Step: [description step]

Output: {Both, Image Relevant, Logic Relevant, None}

MCoT Correctness Task

Image: [image]

Question: [question]

Rationale: [rationale]

Is this a good rationale or not? Output: {Yes, No}

Description Error Types Task

Image: [image]

Step: [incorrect description step]

Output: {Entity False, Attribute False, Spatial Relationship False, Non-spatial Relationship False}

Logic Correctness

Question: [question]

Premise: [previous steps]

Hypothesis: [reasoning step]

Output: {Correct, Incorrect}

Logic Relevance Task

Question: [question]

Rationale: [rationale]

Step: [reasoning step]

Output: {Relevant, Irrelevant}

Logic Error Types Task

Premise: [previous steps]

Hypothesis: [incorrect reasoning step]

Output: {Inter-step Incorrect, Intra-step Incorrect, Both}

Informativeness Task

C.3 Detailed Results

Table 9 presents the results in MiCEval-NORMAL. The results of the exploration experiments for few-shot types are illustrated in Figure 10. The F1 metrics for each class on some important tasks under two different evaluation settings, as shown in Table 10 and Table 11.

▷ *The reasoning capabilities of the MCoT verifier still need improvement.* The MCoT responses in MiCEval-HARD involve more reasoning steps and are generally more complex compared to those in MiCEval-NORMAL. MLLM performance on MiCEval-NORMAL consistently exceeds that on MiCEval-HARD, suggesting that current MLLMs align more closely with human judgments on simpler reasoning tasks. This highlights their limitations in effectively validating MCoTs that require complex reasoning.

▷ *Good performance in textual few-shot evaluation can be achieved.* GPT-4o and MiniCPM showed similar performance when handling multimodal demonstrations and textual demonstrations. In contrast, LLaVA-1.6-Mistral-7B exhibited nearly 20% improvement of accuracy in the textual few-shot evaluation compared to the multimodal few-shot evaluation.

| Evaluation Setting | Model | Step Type | Description | | | Logic | | | | MCoT | Avg. |
|--------------------|-------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|-----------------|--------------|--------------|--------------|
| | | | Relevance | Correctness | Error Types | Relevance | Correctness | Informativeness | Error Types | | |
| Zero-Shot | LLaVA-1.6-Mistral-7B | 0.719 | 0.484 | 0.144 | 0.260 | 0.866 | 0.541 | 0.860 | 0.729 | 0.574 | 0.575 |
| | MiniCPM-V-2.6 | 0.184 | 0.731 | 0.782 | 0.480 | 0.950 | 0.774 | 0.789 | 0.931 | 0.756 | 0.709 |
| | Llama-3.2-11B-Vision-Instruct | 0.380 | 0.757 | 0.137 | 0.453 | 0.842 | 0.580 | 0.345 | 0.861 | 0.632 | 0.554 |
| | LLaVA-1.6-Yi-34B | 0.580 | 0.512 | 0.804 | 0.560 | 0.962 | 0.808 | 0.849 | 0.812 | 0.729 | 0.735 |
| | GPT-4o | 0.767 | 0.763 | 0.748 | 0.650 | 0.938 | 0.753 | 0.568 | 0.812 | 0.765 | 0.752 |
| | Qwen-VL-Max | 0.420 | 0.910 | 0.714 | 0.580 | 0.952 | 0.815 | 0.805 | 0.271 | 0.740 | 0.690 |
| | Gemini-1.5-Pro | 0.838 | 0.921 | 0.778 | 0.490 | 0.935 | 0.671 | 0.291 | 0.979 | 0.646 | 0.728 |
| Few-Shot | LLaVA-1.6-Mistral-7B | 0.705 | 0.313 | 0.374 | 0.423 | 0.914 | 0.801 | 0.712 | 0.979 | 0.686 | 0.656 |
| | MiniCPM-V-2.6 | 0.277 | 0.428 | 0.795 | 0.607 | 0.879 | 0.807 | 0.717 | 0.930 | 0.724 | 0.685 |
| | Llama-3.2-11B-Vision-Instruct | 0.005 | 0.223 | 0.716 | 0.010 | 0.805 | 0.760 | 0.130 | 0.271 | 0.251 | 0.352 |
| | LLaVA-1.6-Yi-34B | 0.744 | 0.164 | 0.752 | 0.590 | 0.955 | 0.805 | 0.866 | 0.938 | 0.735 | 0.728 |
| | GPT-4o | 0.856 | 0.530 | 0.739 | 0.730 | 0.908 | 0.726 | 0.623 | 0.773 | 0.785 | 0.741 |
| | Qwen-VL-Max | 0.679 | 0.437 | 0.725 | 0.530 | 0.921 | 0.808 | 0.654 | 0.521 | 0.704 | 0.664 |
| | Gemini-1.5-Pro | - | - | - | - | - | - | - | - | - | - |

Table 9: The overall performance on Pairwise Comparison of different MLLMs in MiCEval-Normal.

| Model | Step Type | | | Description Relevance | | Description Correct | | Logic Relevance | | Logic Correct | | Informativeness | |
|-------------------------------|-------------|-----------|-------|-----------------------|--------|---------------------|-----------|-----------------|------------|---------------|-----------|-----------------|-----------|
| | Description | Reasoning | Both | Both | Others | Correct | Incorrect | Relevant | Irrelevant | Correct | Incorrect | Infor... | Uninfo... |
| Zero-shot | | | | | | | | | | | | | |
| LLaVA-1.6-Mistral-7B | 0.727 | 0.661 | 0.028 | 0.649 | 0.010 | 0.211 | 0.170 | 0.946 | 0.354 | 0.462 | 0.510 | 0.959 | 0.000 |
| MiniCPM-V-2.6 | 0.075 | 0.200 | 0.077 | 0.826 | 0.110 | 0.839 | 0.340 | 0.959 | 0.289 | 0.775 | 0.353 | 0.912 | 0.028 |
| LLaVA-1.6-Yi-34B | 0.641 | 0.181 | 0.086 | 0.716 | 0.110 | 0.884 | 0.180 | 0.974 | 0.358 | 0.791 | 0.417 | 0.956 | 0.167 |
| GPT-4o | 0.849 | 0.574 | 0.048 | 0.879 | 0.100 | 0.851 | 0.140 | 0.947 | 0.388 | 0.741 | 0.455 | 0.834 | 0.297 |
| Qwen-VL-Max | 0.589 | 0.540 | 0.075 | 0.938 | 0.050 | 0.814 | 0.240 | 0.964 | 0.396 | 0.809 | 0.226 | 0.952 | 0.255 |
| Gemini-1.5-Pro | 0.895 | 0.914 | 0.165 | 0.912 | 0.240 | 0.860 | 0.220 | 0.960 | 0.427 | 0.665 | 0.451 | 0.442 | 0.160 |
| Llama-3.2-11B-Vision-Instruct | 0.163 | 0.216 | 0.095 | 0.899 | 0.020 | 0.347 | 0.040 | 0.918 | 0.213 | 0.567 | 0.504 | 0.495 | 0.174 |
| Few-shot | | | | | | | | | | | | | |
| LLaVA-1.6-Mistral-7B | 0.349 | 0.625 | 0.035 | 0.483 | 0.060 | 0.435 | 0.090 | 0.965 | 0.364 | 0.790 | 0.389 | 0.894 | 0.168 |
| MiniCPM-V-2.6 | 0.133 | 0.124 | 0.075 | 0.592 | 0.130 | 0.848 | 0.370 | 0.926 | 0.300 | 0.790 | 0.297 | 0.865 | 0.162 |
| LLaVA-1.6-Yi-34B | 0.792 | 0.769 | 0.128 | 0.297 | 0.080 | 0.851 | 0.100 | 0.968 | 0.422 | 0.788 | 0.260 | 0.962 | 0.031 |
| GPT-4o | 0.908 | 0.897 | 0.188 | 0.625 | 0.080 | 0.843 | 0.190 | 0.940 | 0.410 | 0.757 | 0.493 | 0.871 | 0.325 |
| Qwen-VL-Max | 0.571 | 0.257 | 0.064 | 0.403 | 0.050 | 0.843 | 0.060 | 0.959 | 0.429 | 0.869 | 0.224 | 0.848 | 0.293 |
| Llama-3.2-11B-Vision-Instruct | 0.000 | 0.000 | 0.085 | 0.420 | 0.010 | 0.772 | 0.070 | 0.769 | 0.167 | 0.757 | 0.486 | 0.000 | 0.128 |

Table 10: Per-label F1 metrics of each task under zero-shot and few-shot evaluations on MiCEval-HARD.

C.4 Additional Analysis

Llama under Few-shot Evaluation. We observed that Llama-3.2-11B-Vision-Instruct performed significantly worse in few-shot evaluation compared to zero-shot evaluation. Therefore, we conducted a statistical analysis of the outputs from Llama-3.2 on several tasks where its performance was particularly poor in the few-shot evaluation. We found that the proportion of outputs labeled as “Both” for the step type was 99.9%. Additionally, in terms of informativeness, all outputs of Llama-3.2 were “Uninformative”. In the Description Error Types task, 74.0% of its outputs were invalid, meaning they fell outside the defined label range. Overall, Llama-3.2 exhibited a decline in performance across every task in the few-shot evaluation, which may be related to its training process, suggesting that Llama-3.2 might lack multi-round and multi-image training.

Description Relevance in Pairwise Comparison. We observed that on both MiCEval-HARD and MiCEval-NORMAL, all MLLMs performed worse in description relevance during few-shot evaluation compared to zero-shot evaluation. We analyzed each model’s accuracy across the labels in

the description relevance task, with results shown in Tables 10 and 11. Our analysis revealed that the prediction distribution of each model shifted in different ways during few-shot evaluation. We believe this performance drop may be due to the larger number of label classes in description relevance compared to other tasks, as well as the generally weaker in-context learning abilities of MLLMs for this task.

D Evaluator Experiment

D.1 Detail Experimental Setting and Metrics

We selected image-text matching metrics such as CLIP (Hessel et al., 2021), BLIP2-ITM, BLIP2-ITC (Li et al., 2023a), and LLM-Score (Chen et al., 2024e), along with the reasoning chain evaluation metric ReCEval (Prasad et al., 2023), as baselines. For each metric, we compute the scores corresponding to each step in the MCoT answer.

For the MiCEval evaluators, we selected LLaVA-1.6-Mistral-7B, MiniCPM-V-2.6, Llama-3.2-11B-Vision-Instruct, and GPT-4o as foundational MLLMs. These MLLMs assess all steps in the MCoT answers based on both the MiCEval description and reasoning step correctness. For

| Model | Step Type | | | Description | | Description | | Logic | | Logic | | Informativeness | |
|-------------------------------|-------------|-----------|-------|-------------|--------|-------------|-----------|----------|------------|---------|-----------|-----------------|-----------|
| | Description | Reasoning | Both | Both | Others | Correct | Incorrect | Relevant | Irrelevant | Correct | Incorrect | Infor... | Uninfo... |
| Zero-Shot | | | | | | | | | | | | | |
| LLaVA-1.6-Mistral-7B | 0.824 | 0.367 | 0.000 | 0.660 | 0.010 | 0.200 | 0.080 | 0.928 | 0.133 | 0.656 | 0.309 | 0.924 | 0.047 |
| MiniCPM-V-2.6 | 0.371 | 0.053 | 0.007 | 0.854 | 0.050 | 0.890 | 0.260 | 0.974 | 0.000 | 0.870 | 0.000 | 0.879 | 0.114 |
| LLaVA-1.6-Yi-34B | 0.774 | 0.111 | 0.025 | 0.676 | 0.110 | 0.922 | 0.100 | 0.981 | 0.154 | 0.893 | 0.097 | 0.913 | 0.450 |
| GPT-4o | 0.919 | 0.489 | 0.040 | 0.868 | 0.070 | 0.888 | 0.140 | 0.968 | 0.400 | 0.854 | 0.200 | 0.674 | 0.364 |
| Qwen-VL-Max | 0.637 | 0.242 | 0.000 | 0.960 | 0.000 | 0.861 | 0.160 | 0.975 | 0.133 | 0.898 | 0.036 | 0.878 | 0.513 |
| Gemini-1.5-Pro | 0.915 | 0.689 | 0.051 | 0.963 | 0.150 | 0.900 | 0.170 | 0.966 | 0.250 | 0.791 | 0.226 | 0.312 | 0.269 |
| Llama-3.2-11B-Vision-Instruct | 0.605 | 0.020 | 0.011 | 0.856 | 0.000 | 0.244 | 0.030 | 0.914 | 0.000 | 0.720 | 0.176 | 0.411 | 0.281 |
| Few-Shot | | | | | | | | | | | | | |
| LLaVA-1.6-Mistral-7B | 0.821 | 0.332 | 0.000 | 0.476 | 0.020 | 0.560 | 0.070 | 0.955 | 0.194 | 0.888 | 0.094 | 0.818 | 0.311 |
| MiniCPM-V-2.6 | 0.499 | 0.065 | 0.011 | 0.582 | 0.020 | 0.892 | 0.290 | 0.938 | 0.138 | 0.888 | 0.094 | 0.834 | 0.147 |
| LLaVA-1.6-Yi-34B | 0.858 | 0.529 | 0.043 | 0.268 | 0.140 | 0.893 | 0.110 | 0.977 | 0.235 | 0.891 | 0.066 | 0.927 | 0.170 |
| GPT-4o | 0.926 | 0.718 | 0.036 | 0.697 | 0.030 | 0.875 | 0.180 | 0.950 | 0.372 | 0.833 | 0.231 | 0.726 | 0.396 |
| Qwen-VL-Max | 0.813 | 0.216 | 0.032 | 0.610 | 0.010 | 0.876 | 0.060 | 0.959 | 0.207 | 0.893 | 0.067 | 0.754 | 0.416 |
| Llama-3.2-11B-Vision-Instruct | 0.005 | 0.000 | 0.007 | 0.368 | 0.000 | 0.841 | 0.060 | 0.890 | 0.123 | 0.859 | 0.186 | 0.000 | 0.230 |

Table 11: Per-label F1 metrics of each task under zero-shot and few-shot evaluations on MiCEval-NORMAL.

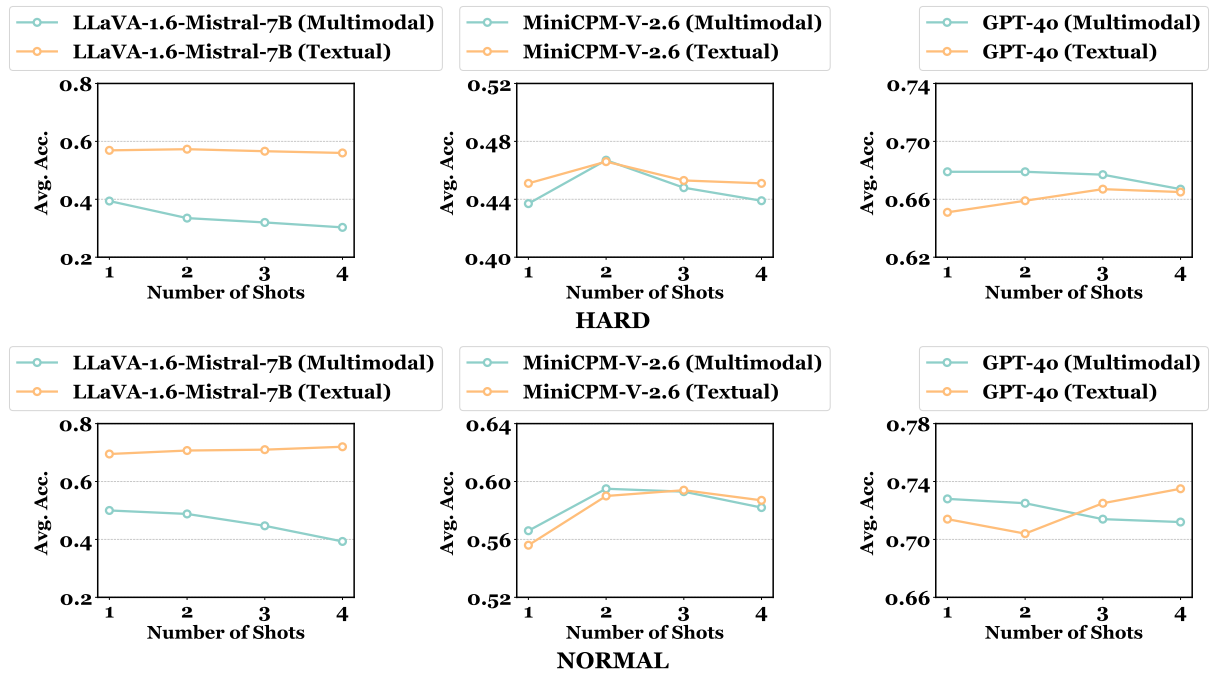


Figure 10: The relationship between the average accuracy of three MLLMs across five image-input Pairwise Comparison tasks and the number of shots. **HARD**: The performance on MiCEval-HARD; **NORMAL**: The performance on MiCEval-NORMAL.

$Correctness_{type}$, we use the results of Gemini-1.5-Pro on the step type task to determine the type label for each step (best zero-shot step type performance on Sec. 6.1.2). We then calculate two kinds of overall MCoT correctness (Step, MiCEval, Type or All). Finally, we compare the correlation between the MCoT correctness and human scores. Additionally, we evaluate the correctness scores generated by the MLLMs for each step (Step, w/o MiCEval) and for the overall MCoT (CoT, w/o MiCEval) without utilizing the MiCEval metrics. This is done to validate the effectiveness of MiCEval while minimizing the influence of the “Step-by-Step” evaluation method.

Scoring Evaluation. We use Somer’s D (Somers,

1962) as the metric for human correlation.

Choice Ranking. Each question in the dataset contains at least one fully correct MCoT option and one that is not. We calculated correctness scores for each MCoT answer and selected the option with the highest score. Accuracy was then computed based on whether the selected choice was fully correct. The MCoT evaluation metrics that achieved the highest accuracy demonstrate a stronger ability to filter high-quality MCoTs. We use accuracy as the evaluation metric. An example of Choice Ranking data as shown as Table 12.

D.2 Prompt Template

In this section, the system prompt we used in each setting are illustrated in the Figure 11, and we pro-


| Image | Question | MCoT Answer | Generator | Ground Truth |
|---|-------------------------|--|----------------------|--------------|
|  | Why is the clock there? | (1) The clock is positioned prominently in a busy street area. It is elevated on a pole, making it visible to pedestrians and drivers. The area appears to be a commercial district with shops and businesses. So the clock is there to help people keep track of time while they are out and about. | GPT-4o | Correct |
| | | (2) This is an outdoor scene during the day. A large black clock tower stands prominently against the sky. It has 2 clocks facing different directions. I can assume that its purpose might be to inform or remind passersby about the time. | Qwen-VL-Max | Incorrect |
| | | (3) The clock is on a pole in the middle of a city street. People are walking by the clock and looking at it. The clock is showing the time. Therefore, the clock is there for people to check the time. | Instructbl Vicuna-7B | Correct |
| | | (4) The clock is located on a pole on the sidewalk. It is a public clock, likely placed there for pedestrians to easily see the time. The clock is situated in a busy city street with many people walking by. The clock serves as a convenient time reference for pedestrians in the area. | LLaVA-1.6-Mistral-7B | Correct |

Table 12: An Example of Choice Ranking Data.

vide prompt templates used in the experiments as following:

MCoT Task

Image: [image]
 Question: [question]
 CoT: [MCoT answer]
 Output: [Score of MCoT]

Step Task

Image: [image]
 Question: [question]
 Previous: [previous steps]
 Current Step: [current step]
 Output: [Score of the Current Step]

D.3 Additional Results

▷ *Step-by-step evaluation does not always help MLLMs.* Step-by-step evaluation benefits MiniCPM-V-2.6 and LLaVA-1.6-Mistral but harms Llama-3.2-11B-Vision-Instruct’s performance. We observe that, compared to directly evaluating the entire MCoT without MiCEval, step-by-step eval-

uation without MiCEval improves the Somer’s D score of MiniCPM-V-2.6 from 0.130 to 0.169 and LLaVA-1.6-Mistral-7B from 0.080 to 0.140 on the whole MiCEval. In contrast, Llama’s performance drops from 0.176 to 0.075. Additionally, these scores remain lower than those achieved by the MiCEval-based step-by-step evaluation method, indicating that step-by-step evaluation alone is not the primary factor driving improvements in MiCEval MLLM-based methods.

D.4 Analysis based on Human Annotation

We use the step type results from human annotation to evaluate the correlation between the $Correctness_{type}$ -based MiCEval metrics and human judgments, focusing on the Description step, Reasoning step, and the overall MiCEval dataset. We also present the Spearman’s Rank Correlation Coefficient (Spearman, 1904) results.

The Table 16 and Table 17 show the results on human annotated step type labels. The Spearman’s Rank Correlation Coefficient (Spearman, 1904) results for different datasets, as illustrated in Table 18

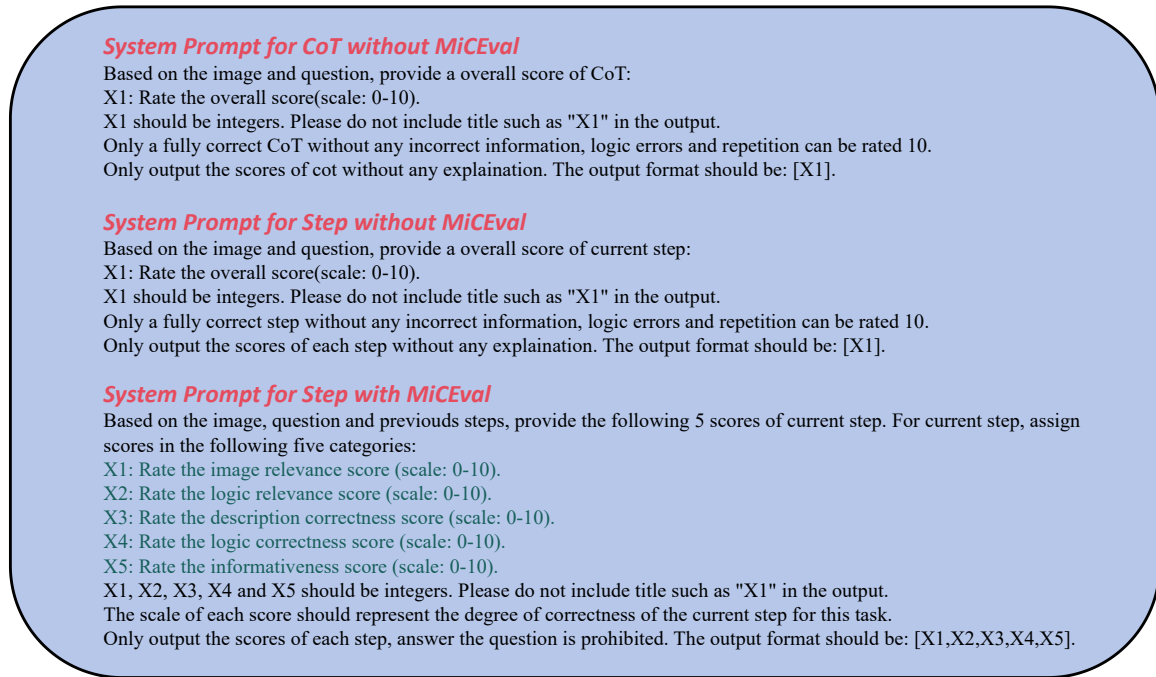


Figure 11: Different system prompts of Scoring Evaluation. **System Prompt for Step with MiCEval** is the system prompt for MLLM-based MiCEval evaluation metrics.

and Table 19. The Choice Ranking results based on human annotation is shown in Table 20.

E Case Study

In this chapter, we present examples based on two different MiCEval metrics: the high-scoring and low-scoring cases. Figures 15 to 18 show the best cases based on the $Correctness_{type}$ -based MiCEval metric, while Figures 19 to 22 illustrate the worst cases. The high-scoring cases for the $Correctness_{all}$ -based MiCEval metric are shown in Figures 23 to 26, and the low-scoring cases are depicted in Figures 27 to 30.

| Dataset | Description | Question |
|----------------|---|---|
| Visual7W | A dataset designed for grounded VQA. | Where could this be? What is in front of the building? |
| VSR | Visual spatial reasoning VQA. | Is the following statement true or false? Statement: The laptop is at the left side of the chair. Is the following statement true or false? Statement: The train is behind the motorcycle. |
| ScienceQA | VQA centers around science topics. | Which of the following fossils is oldest? Which property do these three objects have in common? |
| MMVP | A dataset aims to detect obvious visual differences in images. | Is the school bus driving towards or away from the camera? In the picture, are the elderly people moving to the left or to the right? |
| MMstar | A dataset designed with challenging VQA. | What might the woman in the image be doing? How many types of fruits are there in the image? |
| MM-Vet | A dataset featuring diverse question types and corresponding answers. | What might the woman in the image be doing? How many types of fruits are there in the image? |
| VQAv2 | A dataset with questions covering various aspects of the image. | What type of bird is on the car? What is this form of transportation? |
| Vizwiz | A dataset with photos lacking clarity and exhibiting suboptimal quality collected from users. | What is this color please? What is this box? |

Table 13: Detailed descriptions and randomly sampled questions of 8 source datasets.







| Image | Source Dataset | Question | MCoT Answer | Error |
|---|----------------|---|--|---|
|  | Visual7w | When is this room used? | In the bedroom, there is a bed covered by a striped quilt. On both sides of the headboard hang two pictures. A wooden cabinet stands next to the door. Near the window, there is an armchair placed against the wall. | Entity False (step 4) |
|  | VSR | Is the following statement true or false? Statement: The dog is at the right side of the laptop. | A small brown chihuahua sits behind an open black laptop computer. The laptop screen displays text that reads "The responsibilities include...". To the left of the laptop, there's a telephone receiver. On the desk where these objects rest, there appears to be food from take-out containers. | Image Relevant (step2, step3, step4) |
|  | ScienceQA | Which property do these two objects have in common? | Both the wooden ruler and the bowling ball are objects. They are both made of wood. The wooden ruler is a tool used for measuring. The bowling ball is a sports equipment used for playing bowling. Both the wooden ruler and the bowling ball are made of wood and have different functions. | Inter-step Incorrect (step2, step 5) |
|  | MMVP | Is the wind blowing the flag? | The flag in the image is not hanging straight down. It appears to be extended outward and has a wavy appearance. This suggests that there is some force acting on it. Therefore, the wind is blowing the flag. | Inter-step Incorrect (step 4) |
|  | MM-Vet | Are the trees taller than the giraffes? | The giraffes are taller than the trees. The trees are shorter than the giraffes. Therefore, the giraffes are taller than the trees. | Uninformative (step 3) |
|  | MMstar | What will happen next? | The man in the image is frowning and pouting. He appears to be upset and unhappy. Therefore, he might cry next. | Attribute False (step 1), Inter-step Incorrect (step2, step3) |

Table 14: Randomly sampled questions from MiCEval-HARD.



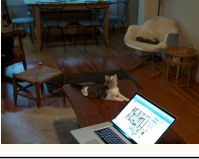

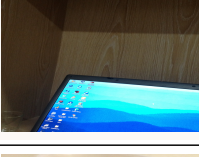

| Image | Source Dataset | Question | MCoT Answer | Error |
|---|----------------|---|---|---------------------------------------|
|  | Visual7w | What color are the middle suitcases? | In the image, there are several suitcases in the trunk of a car. The middle suitcases are red. | Attribute False (step 2) |
|  | Visual7W | What is growing on the right? | There are buildings on the right side of the image. There are trees growing next to the buildings. | Entity False (step 2) |
|  | VQAv2 | Does the cat belong there? | The image shows a cat relaxing comfortably in a home setting. The cat belongs there. | Inter-step Incorrect (step 2) |
|  | VQAv2 | What type of fruit is sliced and laying around the cheese-cake? | There's a plate containing food. On that plate, several banana slices are visible. These bananas have been cut into chunks or slices. Therefore, it appears as if they're being used to garnish or accompany the dessert. | Image Relevant (step 4) |
|  | Vizwiz | Is is a box? Is it a box? | The image shows a wooden structure with three visible sides. The structure resembles the inside of a box. This is a box. | Inter-step Incorrect (step 2, step 3) |
|  | Vizwiz | What is this? | The image shows a can of Progresso chicken and noodle soup. The can has a blue label with white text. This is a can of Progresso chicken and noodle soup. | Uninformative (step 3) |

Table 15: Randomly sampled questions from MiCEval-NORMAL.

Annotation Tool

Picture:



Question:

Is this dog ready to take a nap?

Ground Truth:

Yes.

Step1:

The dog is laying on a bed.

Step2:

The dog appears to be resting its head on a pillow.

Step3:

The dog's eyes are closed, suggesting it may be tired or ready to take a nap.

Current Annotated Step:

The dog is laying on a bed.

- Description
- Reasoning
- Both

Figure 12: The annotation tool screenshot for step type task.

Annotation Tool

Picture:



Question:

Is this dog ready to take a nap?

Ground Truth:

Yes.

Current Annotated Step:

The dog is laying on a bed.

Task1: Description Relevance Task

- Image Relevant
- Logic Relevant

Task2: Description Correctness Task

- Fully correct
- Partially correct
- Unsupported

Task3: Description Error Types Task

- Entity False
- Attribute False
- Spatial Relationship False
- Nonspatial Relationship False

Figure 13: The annotation tasks related to the Description step.

Annotation Tool

Picture:



Question:

Is this dog ready to take a nap?

Ground Truth:

Yes.

Current Annotated Step:

The dog's eyes are closed, suggesting it may be tired or ready to take a nap.

Task1: Logic Relevance Task

- Relevant
- Irrelevant

Task2: Logic Correctness Task

- Correct
- Incorrect

Task3: Logic Error Types Task

- Inter-step Incorrect
- Intra-step Incorrect

Task3: Informativeness Task

- Informative
- Uninformative

Figure 14: The annotation logic tasks related to the Both step.

| Metrics | Description Step | Reasoning Step |
|---|------------------|----------------|
| CLIP | 0.051 | 0.031 |
| BLIP2-ITM | 0.035 | -0.003 |
| BLIP2-ITC | 0.071 | 0.056 |
| ReCEval | -0.012 | -0.040 |
| LLM-Score | 0.130 | 0.141 |
| MiniCPM-V-2.6 (CoT, w/o MiCEval) | - | - |
| LLaVA-1.6-Mistral-7B (CoT, w/o MiCEval) | - | - |
| Llama-3.2-11B-Vision-Instruct (CoT, w/o MiCEval) | - | - |
| GPT-4o (CoT, w/o MiCEval) | - | - |
| MiniCPM-V-2.6 (Step, w/o MiCEval) | 0.110 | 0.103 |
| LLaVA-1.6-Mistral-7B (Step, w/o MiCEval) | 0.109 | -0.025 |
| Llama-3.2-11B-Vision-Instruct (Step, w/o MiCEval) | 0.005 | 0.126 |
| GPT-4o (Step, w/o MiCEval) | 0.168 | 0.298 |
| MiniCPM-V-2.6 (Step, MiCEval) | 0.240 | 0.183 |
| LLaVA-1.6-Mistral-7B (Step, MiCEval) | 0.098 | -0.044 |
| Llama-3.2-11B-Vision-Instruct (Step, MiCEval) | 0.115 | 0.098 |
| GPT-4o (Step, MiCEval) | 0.240 | 0.319 |

Table 16: Human-annotation-based Somer's D results on Description Step, Reasoning Step.

| Metric | NORMAL | HARD | MiCEVAL |
|---|--------------|--------------|--------------|
| CLIP | 0.079 | 0.019 | 0.060 |
| BLIP2-ITM | 0.088 | -0.065 | 0.031 |
| BLIP2-ITC | 0.072 | -0.006 | 0.075 |
| ReCEval | -0.006 | 0.015 | 0.040 |
| LLM-Score | 0.078 | 0.210 | 0.162 |
| MiniCPM-V-2.6 (CoT, w/o MiCEval) | 0.154 | 0.123 | 0.130 |
| LLaVA-1.6-Mistral-7B (CoT, w/o MiCEval) | 0.109 | -0.024 | 0.080 |
| Llama-3.2-11B-Vision-Instruct (CoT, w/o MiCEval) | 0.090 | 0.218 | 0.176 |
| GPT-4o (CoT, w/o MiCEval) | 0.154 | 0.256 | 0.208 |
| MiniCPM-V-2.6 (Step, w/o MiCEval) | 0.112 | 0.123 | 0.169 |
| LLaVA-1.6-Mistral-7B (Step, w/o MiCEval) | 0.083 | 0.178 | 0.140 |
| Llama-3.2-11B-Vision-Instruct (Step, w/o MiCEval) | 0.031 | 0.090 | 0.075 |
| GPT-4o (Step, w/o MiCEval) | 0.178 | 0.282 | 0.257 |
| MiniCPM-V-2.6 (Step, MiCEval) | 0.260 | 0.198 | 0.279 |
| LLaVA-1.6-Mistral-7B (Step, MiCEval) | 0.080 | 0.142 | 0.126 |
| Llama-3.2-11B-Vision-Instruct (Step, MiCEval) | 0.158 | 0.170 | 0.192 |
| GPT-4o (Step, MiCEval) | 0.256 | 0.264 | 0.281 |

Table 17: Human-annotation-based Somer’s D results on MiCEval dataset.

| Model | Description Step | Reasoning Step |
|---|------------------|----------------|
| CLIP | 0.105 | 0.055 |
| BLIP2-itm | 0.073 | -0.005 |
| BLIP2-itc | 0.148 | 0.100 |
| ReCEval | -0.025 | -0.072 |
| LLM-Score | 0.245 | 0.232 |
| MiniCPM-V-2.6 (CoT, w/o MiCEval) | – | – |
| LLaVA-1.6-Mistral-7B (CoT, w/o MiCEval) | – | – |
| Llama-3.2-11B-Vision-Instruct (CoT, w/o MiCEval) | – | – |
| GPT-4o (CoT, w/o MiCEval) | – | – |
| MiniCPM-V-2.6 (Step, w/o MiCEval) | 0.197 | 0.158 |
| LLaVA-1.6-Mistral-7B (Step, w/o MiCEval) | 0.178 | -0.035 |
| Llama-3.2-11B-Vision-Instruct (Step, w/o MiCEval) | 0.006 | 0.124 |
| GPT-4o (Step, w/o MiCEval) | 0.276 | 0.408 |
| MiniCPM-V-2.6 (Step, MiCEval) | 0.249 | 0.153 |
| LLaVA-1.6-Mistral-7B (Step, MiCEval) | 0.162 | -0.051 |
| Llama-3.2-11B-Vision-Instruct (Step, MiCEval) | 0.194 | 0.135 |
| GPT-4o (Step, MiCEval) | 0.365 | 0.367 |

Table 18: Spearman’s Rank Correlation Coefficient on Description Step, Reasoning Step.

| Metrix | NORMAL | HARD | MiCEval |
|---|--------------|--------------|--------------|
| CLIP | 0.148 | 0.033 | 0.105 |
| BLIP2-itm | 0.165 | -0.114 | 0.054 |
| BLIP2-itc | 0.134 | -0.011 | 0.131 |
| ReCEval | -0.011 | 0.026 | 0.069 |
| LLM-Score | 0.144 | 0.369 | 0.279 |
| MiniCPM-V-2.6 (CoT, w/o MiCEval) | 0.227 | 0.182 | 0.186 |
| LLaVA-1.6-Mistral-7B (CoT, w/o MiCEval) | 0.117 | -0.028 | 0.086 |
| Llama-3.2-11B-Vision-Instruct (CoT, w/o MiCEval) | 0.130 | 0.281 | 0.234 |
| GPT-4o (CoT, w/o MiCEval) | 0.238 | 0.384 | 0.305 |
| MiniCPM-V-2.6 (Step, w/o MiCEval) | 0.204 | 0.211 | 0.289 |
| LLaVA-1.6-Mistral-7B (Step, w/o MiCEval) | 0.146 | 0.303 | 0.232 |
| Llama-3.2-11B-Vision-Instruct (Step, w/o MiCEval) | 0.045 | 0.141 | 0.110 |
| GPT-4o (Step, w/o MiCEval) | 0.287 | 0.409 | 0.381 |
| MiniCPM-V-2.6 (Step, MiCEval) | 0.263 | 0.259 | 0.320 |
| LLaVA-1.6-Mistral-7B (Step, MiCEval) | 0.130 | 0.221 | 0.192 |
| Llama-3.2-11B-Vision-Instruct (Step, MiCEval) | 0.266 | 0.263 | 0.300 |
| GPT-4o (Step, MiCEval) | 0.408 | 0.383 | 0.413 |

Table 19: Spearman’s Rank Correlation Coefficient on MiCEval dataset.

| Model | Acc. |
|---|--------------|
| CLIP | 0.700 |
| BLIP2-ITM | 0.657 |
| BLIP2-ITC | 0.657 |
| ReCEval | 0.600 |
| LLM-Score | 0.614 |
| MiniCPM-V-2.6 (CoT, w/o MiCEval) | 0.686 |
| LLaVA-1.6-Mistral-7B (CoT, w/o MiCEval) | 0.686 |
| Llama-3.2-11B-Vision-Instruct (CoT, w/o MiCEval) | 0.729 |
| GPT-4o (CoT, w/o MiCEval) | 0.757 |
| MiniCPM-V-2.6 (Step, w/o MiCEval) | 0.700 |
| LLaVA-1.6-Mistral-7B (Step, w/o MiCEval) | 0.571 |
| Llama-3.2-11B-Vision-Instruct (Step, w/o MiCEval) | 0.757 |
| GPT-4o (Step, w/o MiCEval) | 0.771 |
| MiniCPM-V-2.6 (Step, MiCEval) | 0.800 |
| LLaVA-1.6-Mistral-7B (Step, MiCEval) | 0.586 |
| Llama-3.2-11B-Vision-Instruct (Step, MiCEval) | 0.786 |
| GPT-4o (Step, MiCEval) | 0.800 |

Table 20: Accuracy based on human-annotation labels of different evaluation metrics on Choice Ranking.



Question

Do the persons live in this house? Answer it and give the rationale.

MCoT Answer

The persons in the image are dressed in historical clothing, which suggests they are part of a historical reenactment or tour. The interior of the house is very ornate and appears to be a historical mansion or museum. The people are walking in a manner that suggests they are observing the surroundings rather than moving about their home. Therefore, it is unlikely that the persons live in this house.

| | |
|--------------------|---------|
| Description | Correct |
| Reasoning | Correct |
| MCoT | Correct |

Figure 15: A $Correctness_{type}$ -based MiCEval metric high-scoring case of MiCEval-HARD.



Question

When was this picture taken, seasonally?

MCoT Answer

The ground in the picture is covered with snow. Snow typically falls during the colder months of the year. The presence of snow suggests that the picture was taken during a cold season. So this picture may have been taken in winter.

Description

Correct

Reasoning

Correct

MCoT

Correct

Figure 16: A $Correctness_{type}$ -based MiCEval metric high-scoring case of MiCEval-HARD.



Question

What is this?

MCoT Answer

The image shows two plastic bottles on a desk. Both bottles have labels that read "Nestle Pure Life." These are water bottles.

Description

Correct

Reasoning

Correct

MCoT

Correct

Figure 17: A $Correctness_{type}$ -based MiCEval metric high-scoring case of MiCEval-NORMAL.



Question

Is the room illuminated by natural lighting?

MCoT Answer

The image shows a room with large windows and sunlight coming through. The room is illuminated by natural lighting.

Description

Correct

Reasoning

Correct

MCoT

Correct

Figure 18: A $Correctness_{type}$ -based MiCEval metric high-scoring case of MiCEval-NORMAL.



Question

What will happen next?

MCoT Answer

The image shows a playground with children playing. One child in a Batman shirt is running down a ramp. The child appears to be moving quickly and is looking ahead. Given the momentum and direction, the child will likely continue running down the ramp. So, the child will continue running down the ramp next.

Description

Correct

Reasoning

Incorrect

MCoT

Incorrect

Figure 19: A $Correctness_{type}$ -based MiCEval metric low-scoring case of MiCEval-HARD.



Question

Where is the truck going?

MCoT Answer

The truck has yellow body color. It's a Mack brand model. The license plate number can be seen. There is a man sitting inside the cabin.

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | \ |
| MCoT | Incorrect |

Figure 20: A *Correctness_{type}*-based MiCEval metric low-scoring case of MiCEval-HARD.



Question


What flavor is this?

MCoT Answer

The image shows a box of cereal on a wooden surface. The cereal appears to have a green and purple design. The flavor of the cereal is not specified in the image. Therefore, it cannot be determined what flavor the cereal is.

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | Incorrect |
| MCoT | Incorrect |

Figure 21: A *Correctness_{type}*-based MiCEval metric low-scoring case of MiCEval-NORMAL.



Question

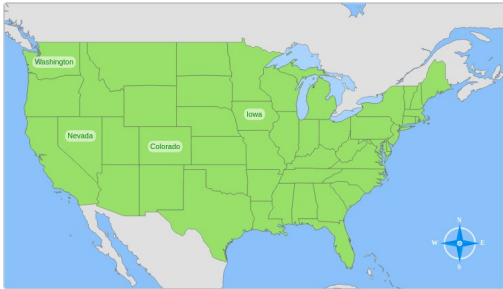
What can you see in the picture?

MCoT Answer

You can see a blue mug on a wooden table. Inside the mug, there is a dark liquid. There are red and white pieces that resemble imitation crab sticks.

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | \ |
| MCoT | Incorrect |

Figure 22: A $Correctness_{type}$ -based MiCEval metric low-scoring case of MiCEval-NORMAL.



Question

Which of these states is farthest east?

MCoT Answer

The map shows the states of Washington, Nevada, Colorado, and Iowa. Iowa is located farther east than Washington, Nevada, and Colorado. So, Iowa is the farthest east.

| | |
|--------------------|---------|
| Description | Correct |
| Reasoning | Correct |
| MCoT | Correct |

Figure 23: A $Correctness_{all}$ -based MiCEval metric high-scoring case of MiCEval-HARD.



Question

Is the school bus driving towards or away from the camera?

MCoT Answer

The image shows the back of a school bus. The rear lights and the "SCHOOL BUS" sign are visible. The bus is on a road, and the front of the bus is not visible. Therefore, the school bus is driving away from the camera.

| | |
|--------------------|---------|
| Description | Correct |
| Reasoning | Correct |
| MCoT | Correct |

Figure 24: A *Correctness_{all}*-based MiCEval metric high-scoring case of MiCEval-HARD.



Question


When will the baby have hair?

MCoT Answer

The baby in the picture has very little hair. Babies typically start to grow more hair as they get older. By the age of 6 to 12 months, many babies have a noticeable amount of hair. So the baby will likely have more hair in a few months.

| | |
|--------------------|---------|
| Description | Correct |
| Reasoning | Correct |
| MCoT | Correct |

Figure 25: A *Correctness_{all}*-based MiCEval metric high-scoring case of MiCEval-NORMAL.



Question

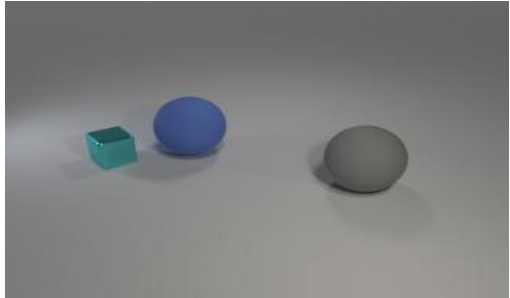
What are the men doing?

MCoT Answer

In the image, two men are holding game controllers and pointing at a screen. The men are playing a video game.

| | |
|--------------------|---------|
| Description | Correct |
| Reasoning | Correct |
| MCoT | Correct |

Figure 26: A $Correctness_{all}$ -based MiCEval metric high-scoring case of MiCEval-NORMAL.



Question

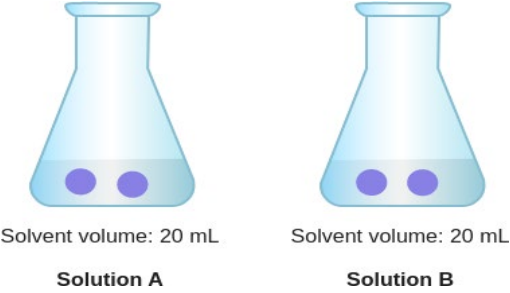
Subtract all large gray things. Subtract all small brown metallic balls. How many objects are left?

MCoT Answer

In the image, there are five large gray things and six small brown metallic balls. Five large gray things minus six small brown metallic balls equals one object remaining. Therefore, only one object remains in the scene. This could potentially be any object that was not mentioned or depicted in the original image.

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | Incorrect |
| MCoT | Incorrect |

Figure 27: A $Correctness_{all}$ -based MiCEval metric low-scoring case of MiCEval-HARD.



Solvent volume: 20 mL Solvent volume: 20 mL

Solution A **Solution B**

Question


Which solution has a higher concentration of purple particles?

MCoT Answer

Both solutions contain equal volumes (20 mL) of solvent. Solution A contains two purple particles while Solution B contains only one purple particle. Since there are twice as many particles present in Solution A compared to Solution B, it can be concluded that Solution A has a higher concentration of purple particles

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | Incorrect |
| MCoT | Incorrect |

Figure 28: A *Correctness_{all}*-based MiCEval metric low-scoring case of MiCEval-HARD.



Question

What is hanging from the oven?

MCoT Answer

In the image, there is a banana peel hanging from the oven. The banana peel is hanging from the oven.

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | Incorrect |
| MCoT | Incorrect |

Figure 29: A *Correctness_{all}*-based MiCEval metric low-scoring case of MiCEval-NORMAL.



Question

What pattern is the tie that the gentleman is wearing?

MCoT Answer

The tie the gentleman is wearing is striped.

| | |
|--------------------|-----------|
| Description | Incorrect |
| Reasoning | \ |
| MCoT | Incorrect |

Figure 30: A *Correctness_{all}*-based MiCEval metric low-scoring case of MiCEval-NORMAL.