

Faithful-First Reasoning, Planning, and Acting for Multimodal LLMs

Junxian Li^{1*} Xinyue Xu^{2*} Sai Ma^{3*} Di Zhang⁴ Sichao Li^{5†}

¹Shanghai Jiao Tong University

²The Hong Kong University of Science and Technology

³The Australian National University ⁴Fudan University ⁵University of Sydney
lijunxian0531@sjtu.edu.cn, xinyue.xu@connect.ust.hk, sai.ma@anu.edu.au
di.zhang@ustc.edu, sichao.li@sydney.edu.au

Abstract

Multimodal Large Language Models (MLLMs) frequently suffer from unfaithfulness, generating reasoning chains that drift from visual evidence or contradict final predictions. We propose *Faithful-First Reasoning, Planning, and Acting (RPA) framework* in which FAITHEVI provides step-wise and chain-level supervision by evaluating the faithfulness of intermediate *reasoning*, and FAITHACT uses these signals to plan and execute faithfulness-aware actions during inference. Experiments across multiple multimodal reasoning benchmarks show that faithful-first RPA improves perceptual faithfulness by up to **24%** over *prompt-based* and *tool-augmented* reasoning frameworks, without degrading task accuracy. Our analysis shows that treating *faithfulness as a guiding principle* perceptually faithful reasoning trajectories and mitigates hallucination behavior. This work thereby establishes a unified framework for both evaluating and enforcing faithfulness in multimodal reasoning. Code is at <https://github.com/lijunxian111/Faithful-First-RPA>.

1 Introduction

Despite rapid progress in multimodal large language models (MLLMs) (Sun et al., 2023; Chen et al., 2024b; An et al., 2025b; Wang et al., 2024), their reasoning traces *remain unfaithful*: models frequently produce persuasive explanations that conflict with perceptual evidence, or utilize post-hoc rationalizations fabricate their reasoning progress (Arcuschin et al., 2025; Barez et al., 2025). This gap poses a central challenge for trustworthy reasoning. Existing efforts typically focus on improving task accuracy or enriching CHAIN-OF-THOUGHT (CoT) generation (Zhang et al., 2023), yet the unfaithful reasoning remains unaddressed.

Motivation. We are motivated by the following principle and observations (Goyal et al., 2017).

A perceptually faithful model reasons only over what is visually observable; it does not “see” what the image does not reveal.

This principle echoes long-standing findings in Visual Question Answering (VQA): systems should avoid answering beyond available evidence and resist over-reliance on language priors (Antol et al., 2015; Agrawal et al., 2018; Bender et al., 2021).

As illustrated by examples in Fig. 1, we observe that MLLMs can generate plausible explanations that are **perceptually inconsistent** with the underlying visual input, *regardless of whether the final prediction is correct*. To formalize this observation, we distinguish perceptual faithfulness (reasoning steps align with the model’s input) from the predominant focus in prior work on behavioral faithfulness (reasoning steps align with the model’s output) (Arcuschin et al., 2025; Matton et al., 2025; Ming et al., 2024; Li et al., 2023b).

Our answer. We argue that FAITHFULNESS SHOULD BE A DESIGN PRINCIPLE, NOT MERELY A POST-HOC EVALUATION OBJECTIVE. Reasoning frameworks should *explicitly verify* the evidential grounding of each step before it is admitted into the chain, ensuring that reasoning remains both perceptually grounded and behaviorally aligned.

Contribution. This motivates the *Faithful-First Reasoning, Planning, and Acting (RPA) framework*, which operationalizes the principle as a design constraint, enforcing faithfulness throughout multimodal reasoning. Within this framework, we make the following contributions:

- We introduce FAITHEVI, an evidence-based perceptual faithfulness evaluation pipeline via evidence that extracts claimed objects from each reasoning step and verifies their existence at both local

*Equal contribution.

†Correspondence: sichao.li@sydney.edu.au

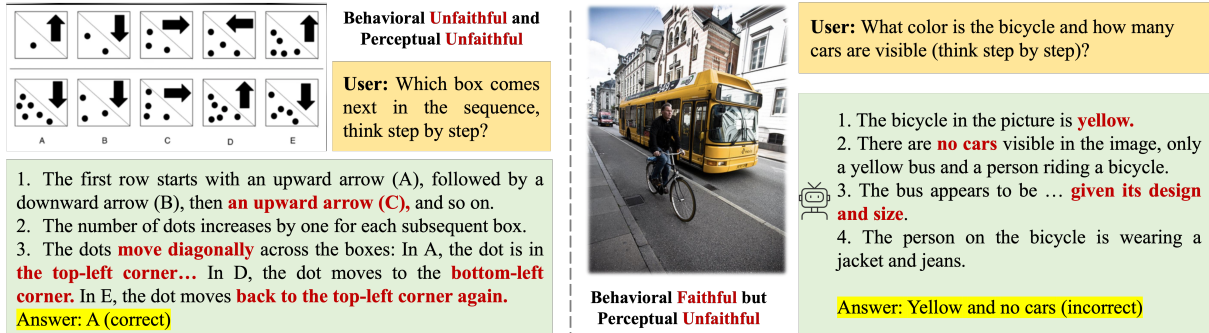


Figure 1: Perceptually and behaviorally unfaithful examples. *Left: Perceptually unfaithful and behaviorally unfaithful.* The model presents a step-by-step reasoning trace describing alternating arrow directions and increasing dot counts, yet such reasoning does not reflect its actual decision process. The final choice (A) is likely made through pattern association, with the explanation generated post hoc to rationalize it. *Right: Perceptually unfaithful but behaviorally faithful.* The reasoning aligns with its final answer but incorrectly describes the bicycle as yellow, influenced by a nearby yellow bus. This illustrates a visually plausible but perception-unfaithful reasoning step, where linguistic association overrides perceptual grounding.

and global levels through preference polling and visual grounding. It assigns a step-level faithfulness score ($F_{\text{step},t}$) to each reasoning step and aggregates them into a chain-level score (F_{chain}), providing a general quantitative measure of perceptual faithfulness with respect to the input evidence.

- We propose FAITHACT, a faithful-first planning and acting mechanism that enforces evidential grounding during inference. It operates through two sequential phases: (1) *Planning*, which leverages FAITHEVI signals to assess the current chain’s faithfulness (F_{chain}) and determine admissible next steps; and (2) *Acting*, which executes faithfulness-aware actions subject to dynamic thresholds. The mechanism is realized via a lightweight, extensible interface of callable functions.
- Extensive experiments across multiple benchmarks reveal that current MLLMs consistently underestimate evidential support in reasoning process. We demonstrate that enforcing step-wise verification significantly enhances perceptual faithfulness without compromising task accuracy. Compared to standard prompt-based and tool-augmented frameworks, *Faithful-First RPA* achieves the highest overall perceptual faithfulness and effectively mitigates content hallucination.

2 Perceptual and Behavioral Faithfulness

Building on the motivating observations (Fig. 1) and prior analyses of explanation faithfulness (Appx. A), we distinguish two “unfaithful” notions and clarify our focus of this work.

Previous studies mainly concern whether a reasoning trace accurately reflects the decision process

that led to the model’s final output (Lyu et al., 2023; Matton et al., 2025; Barez et al., 2025; Arcuschin et al., 2025; Li et al., 2026; Ma et al., 2026a), which we refer to as **Behavioral faithfulness (BF)**. While important, it underestimates the importance of input evidence in the reasoning process. We therefore propose **Perceptual faithfulness (PF)**, which concerns whether individual reasoning steps are grounded in verifiable input evidence, such as entities and attributes that are perceptually present in the input. A perceptually unfaithful explanation may invoke objects and properties that are unsupported or contradicted by the visual input. Importantly, the model is not required to output the correct answer to be faithful (Jacovi and Goldberg, 2020; Dasgupta et al., 2022).

For instance, the left example in Fig. 1 shows a correct prediction accompanied by reasoning that is neither perceptually grounded nor behaviorally aligned. The right example illustrates reasoning that aligns with the model’s prediction but remains perceptually unfaithful, as it relies on visually unsupported attributes (e.g., bicycle color). While ideal reasoning would be both perceptually and behaviorally faithful and yield a correct prediction, optimizing all three remains challenging. In this work, we focus on *perceptual faithfulness*.

This choice is motivated by the hypothesis that when reasoning remains perceptually grounded, behavioral consistency tends to follow, without introducing additional optimization objectives such as output correctness. Practically, perceptual faithfulness offers a more actionable and operationally measurable objective, each reasoning step can be

directly validated against the available multimodal evidence by inference. In contrast, behavioral faithfulness depends on inaccessible internal dynamics of MLLMs and is therefore treated as a consequence rather than a controllable design target.

3 FAITHEVI: Perceptual Faithfulness Evaluation Pipeline

We operationalize perceptual faithfulness by first evaluating how well each reasoning step aligns with the input verifiable evidence. To this end, we design a general-purpose pipeline that systematically evaluates the degree of perceptual grounding across reasoning process at step-wise and chain-level. The pipeline consists of three key stages: (i) *Extraction of Claimed Objects* (ii) *Preference Polling and Grounding* (iii) *Faithfulness Scoring*.

3.1 Extraction of Claimed Objects

We begin by defining the setting for perceptual faithfulness verification. Given a multimodal input consisting of an image I and a corresponding textual query Q , an MLLM produces a reasoning chain by prompting “think step by step” (Wei et al., 2022): $R_{raw} = \{s_1, s_2, \dots, s_T\}$, where each s_t represents an intermediate reasoning step generated by the model. To quantify perceptual faithfulness, we must identify all *claimed objects* across both the question and the reasoning steps.

We therefore employ a structured extraction process to isolate visually grounded claims from textual noise. Specifically, we process the concatenated text of Q and each reasoning step s_t using a helper LLM to extract meaningful object mentions. Each reasoning step is independently analyzed, excluding introductory or concluding phrases (e.g., “Let’s begin reasoning” or “Therefore, the answer is...”) to focus solely on evidence-bearing content.

Concretely, we query Qwen2.5-7B-Instruct (Qwen) (Yang et al., 2024) with a structured prompt (see Appx. B) to extract, for each step s_t , a set O_t of m_t claimed objects, where $O_t = \{O_t^1, O_t^2, \dots, O_t^{m_t}\}$. The union of all extracted objects across the reasoning chain is denoted as

$$O = \bigcup_{t=1}^T O_t. \quad (1)$$

Each O_t^i represents a semantically meaningful object or concept that may correspond to perceivable entities in the input image. This step ensures that the subsequent grounding and polling stages

operate only on visually relevant content, filtering out abstract or non-visual reasoning tokens. The extracted object sets $\{O_t\}_{t=1}^T$ thus serve as the foundation for our evidence verification pipeline, linking textual reasoning to visual perception.

3.2 Preference Polling and Grounding

Once the claimed objects are extracted, the next step is to verify their perceptual validity in the input image. This verification involves evaluating both the *existence* and *localization* of each claimed object through a two-stage evidence assessment pipeline: (i) **Preference Polling**, which estimates the likelihood that a claimed object is perceptually verifiable in the image, and (ii) **Grounding**, which localizes verified objects to specific visual regions for spatial confirmation.

3.2.1 Preference Polling

While object detectors can localize objects spatially, their confidence scores are often unreliable when visual cues are weak, such as in cases of occlusion, clutter, or low salience. To address this limitation, we introduce a lightweight *preference polling* module that performs an initial verification of object existence prior to grounding. This module serves as an evidence gate, polling the visual scene to estimate whether an object is perceptually verifiable.

We employ a frozen CLIP-ViT-Large (CLIP) model (Radford et al., 2021) as the multimodal encoder pair (f_{img}, f_{txt}) and attach a lightweight polling head, a two-layer MLP with GELU activation, to predict whether a given object exists in the image. The polling model is trained on the POPE dataset (Li et al., 2023b), which provides binary image-object existence labels (see Appx. C).

Given an image I and an object X , CLIP encoders produce visual and textual embeddings:

$$\mathbf{v} = f_{img}(I) \in \mathbb{R}^d, \quad \mathbf{t} = f_{txt}(X) \in \mathbb{R}^d. \quad (2)$$

The element-wise product $\mathbf{v} \odot \mathbf{t}$ captures cross-modal interaction, which is passed through the polling head to yield a confidence score:

$$c_p = \sigma(\mathbf{W}_2 \text{GELU}(\mathbf{W}_1(\mathbf{v} \odot \mathbf{t}))), \quad (3)$$

where $\mathbf{W}_1 \in \mathbb{R}^{512 \times d}$, $\mathbf{W}_2 \in \mathbb{R}^{2 \times 512}$, and σ denotes the sigmoid function. The resulting confidence $c_p \in [0, 1]$ is interpreted as the probability that object X exists in the image.

For each reasoning step s_t , polling produces a set of existence confidence scores:

$$C_{p,t} = \{c_{p,t}^1, c_{p,t}^2, \dots, c_{p,t}^{m_t}\}, \quad (4)$$

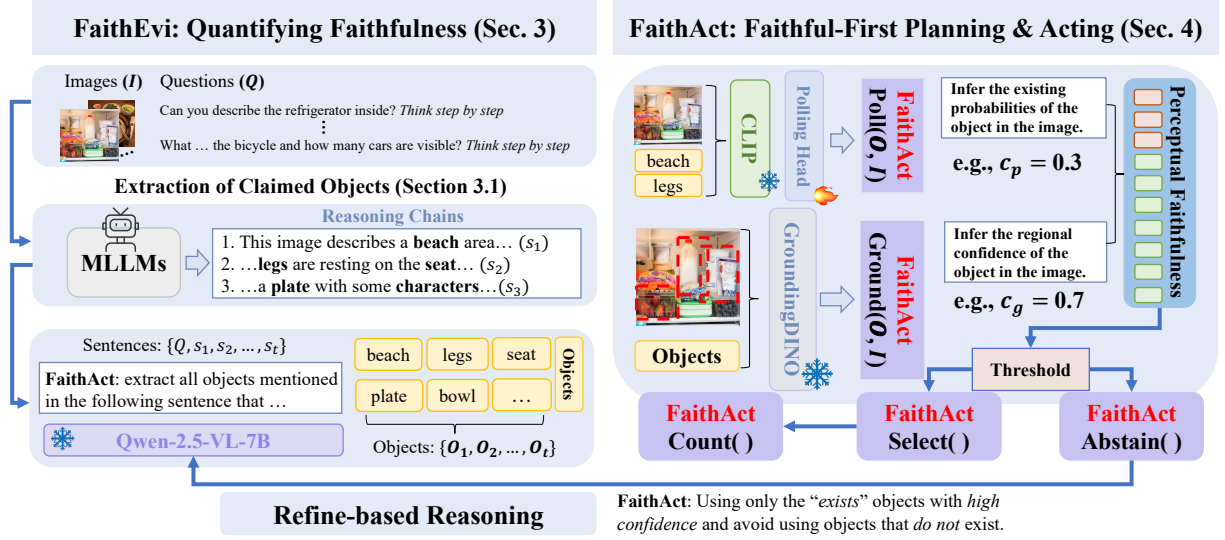


Figure 2: **Faithful-first reasoning, planning, and acting framework.** Given an image-question pair, FAITH_{EVI} evaluates the perceptual faithfulness of intermediate **reasoning**, producing step- and chain-level faithfulness scores. Guided by these signals, FAITH_{ACT} *plans* and *acts* faithfulness-aware actions during inference.

where $c_{p,t}^i$ corresponds to the predicted existence probability for each claimed object O_t^i . This early verification step thus acts as a perceptual verifier, aligning with our evidence-first philosophy.

3.2.2 Grounding

After preference polling confirms the likely existence of an object, we proceed to localize it within the image using a grounded detection model. This step provides region-level visual evidence and confidence scores for objects.

For each object O_t^i we employ a frozen GroundingDINO-base model (Liu et al., 2024) to detect and ground the object in the input image I . GroundingDINO returns a set of bounding boxes and associated confidence scores:

$$\begin{aligned} B_t^i &= \{b_t^{i,1}, b_t^{i,2}, \dots, b_t^{i,k_i}\}, \\ C_t^i &= \{c_t^{i,1}, c_t^{i,2}, \dots, c_t^{i,k_i}\}. \end{aligned} \quad (5)$$

Each $b_t^{i,j}$ represents a candidate region corresponding to object O_t^i , and $c_t^{i,j} \in [0, 1]$ denotes the model’s confidence that the region indeed contains the object. We retain the most confident detection:

$$c_{g,t}^i = \max_j c_t^{i,j}. \quad (6)$$

The resulting set of grounded confidence scores for step s_t is:

$$C_{g,t} = \{c_{g,t}^1, c_{g,t}^2, \dots, c_{g,t}^{m_t}\}. \quad (7)$$

Remark. While preference polling provides global existence verification, grounding supplies fine-grained spatial evidence. Combining these complementary signals yields a more reliable perceptual faithfulness measurement.

3.3 Faithfulness Scoring

After obtaining the *polling confidence* set $C_{p,t}$ and *grounding confidence* set $C_{g,t}$ for each reasoning step s_t , we combine them to compute the perceptual faithfulness of the reasoning process.

Object-Level Confidence. For each claimed object O_t^i within step s_t , we compute an overall confidence score c_t^i by fusing the existence confidence from preference polling and the spatial confidence from grounding:

$$\begin{aligned} c_t^i &= \alpha c_{p,t}^i + (1 - \alpha) c_{g,t}^i, \\ c_{p,t}^i &\in C_{p,t}, \quad c_{g,t}^i \in C_{g,t}, \end{aligned} \quad (8)$$

where $\alpha \in [0, 1]$ controls the relative importance of existence versus localization confidence. Since the polling confidence directly reflects perceptual existence, we empirically set $\alpha = 0.7$.

To interpret the confidence c_t^i as a discrete faithfulness score, we define a three-level mapping function f_t^i as:

$$f_t^i = \begin{cases} 0, & c_t^i < 0.4, \quad (\text{confidently absent}) \\ c_t^i, & 0.4 \leq c_t^i \leq 0.6, \quad (\text{uncertain existence}) \\ 1, & c_t^i > 0.6, \quad (\text{confidently present}). \end{cases}$$

Step-Level and Chain-Level Faithfulness. We then aggregate the object-level confidence scores to obtain a *step-level* faithfulness score for each reasoning step s_t :

$$F_{\text{step},t} = \frac{1}{m_t} \sum_{i=1}^{m_t} f_t^i, \quad (9)$$

which quantifies how faithfully the specific step’s visual claims align with the image. Finally, we compute the *chain-level* perceptual faithfulness of the entire reasoning process as the mean of the verified step-level scores:

$$F_{\text{chain}} = \frac{1}{n} \sum_{t=1}^n F_{\text{step},t}, \quad (10)$$

where n denotes the total number of reasoning steps (excluding introductory or concluding phrases).

A higher F_{chain} indicates that the reasoning chain is more consistently grounded in visual evidence, while lower scores suggest the presence of unsupported object references. This quantitative formulation allows us to evaluate perceptual faithfulness both locally (per step) and globally (across the reasoning trace), serving as a *general* perceptual faithfulness evaluation pipeline.

4 FAITHACT: Faithful-First Planning and Acting

Having established how to quantify perceptual faithfulness, we now turn to the problem of *how to integrate faithfulness into the reasoning process itself*. We propose FAITHACT, a ***Faithfulness-First Planning and Acting*** framework that enforces evidential verification during reasoning generation. Unlike conventional *generate-then-verify* paradigm, FAITHACT follows a *verify-as-you-generate* principle, where each step is explicitly checked for perceptual support before being admitted into the reasoning chain.

4.1 Planning Objective

We explicitly formulate reasoning process as a faithfulness-constrained planning problem: Given a query and image, the planner seeks a sequence of faithful reasoning steps:

$$S^* = \arg \max F_{\text{step}}(s_t) \quad \text{s.t.} \forall t \quad F_{\text{step}}(s_t) \geq c, \quad (11)$$

where s_t is the t -th reasoning step, $F_{\text{step}}(s_t)$ is its faithfulness score from Sec 3.3, and c is a minimum evidential confidence threshold. This converts reasoning into a faithfulness-regularized plan

process: only steps sufficiently supported by evidence are eligible to advance the reasoning chain. When a proposed step s_t fails to meet the threshold ($F_{\text{step}}(s_t) < c$), the planner either refines or regenerates the step before proceeding, thereby enforcing the *faithfulness desideratum*.

4.2 Planning-and-Acting Loop

FAITHACT executes a *faithfulness-first planning loop* at each iteration. A lightweight instruction-tuned MLLM serves as the reasoning controller. This planner composes executable steps such as *SELECT*, *ABSTAIN*, or *COUNT*, based on both textual inputs and verified visual states. Each proposed step is immediately validated by perceptual faithfulness metrics before execution, ensuring that the reasoning process evolves only evidence-supported actions. To facilitate this, FAITHACT provides an extensible interface composed of callable functions that serve as structured APIs. These functions supply multimodal evidence signals and can be invoked by the MLLM during planning, enabling the model to retrieve, verify, and reason over perceptual information in a unified and controlled manner.

`Poll()`: Returns the probability of objects’ existence based on the polling model (Sec. 3.2.1).

`Ground()`: Returns the bounding boxes and confidence scores of a claimed object detected by GroundingDINO (Sec. 3.2.2).

`Select()`: Selects an object as *existent* if its confidence score exceeds the threshold in Sec. 3.3.

`Abstain()`: Abstains from selecting an object if its overall confidence is below the threshold.

`Count()`: Counts the number of *reliably grounded* bounding boxes returned by `Ground()`, yielding the object count for quantitative reasoning.

Among these, `COUNT()` is a functional reasoning operation, while the remaining calls are verification functions that enforce perceptual faithfulness constraints. This modular design allows new functions, such as `Attribute()` or `Relate()`, to be incorporated, and supports further optimization or refinement strategies following `Abstain()`.

4.3 Action-Guided Reasoning Refinement

The overall reasoning process of FAITHACT is summarized in Algorithm 1 in Appx. E, which follows a *refine-based* procedure. After verification, any reasoning step that fails to meet the perceptual faithfulness threshold is not discarded outright but revised through an *action-guided refinement* process. In this phase, the planner re-invokes the MLLM with

Baseline Methods & Datasets	LLaVA-BENCH (%)	REALWORLDQA (%)	POPE (%)	MMHAL (%)	Average (%)
Qwen + CoT	46.05±19.58	48.11±27.04	45.21±24.87	53.34±24.02	48.18
+ VAT	51.59±21.37	50.13±26.43	21.46±19.20	55.32±28.58	44.62
InternVL + CoT	45.63±16.60	44.23±25.43	43.25±23.27	53.17±23.64	46.57
+ VAT	48.97±17.22	45.31±28.19	40.26±22.47	<u>54.51±26.99</u>	47.26
LLaVA + CoT	47.56±23.35	52.31±28.44	52.28±25.66	<u>30.63±28.56</u>	45.70
+ VAT	46.16±19.46	50.15±30.66	<u>52.59±27.15</u>	30.30±28.96	44.80
Qwen + Grounded-CoT	50.04±17.54	53.35±26.68	53.49±22.47	56.77±25.86	<u>53.41</u>
+ ReAct	54.82±26.53	<u>56.82±31.71</u>	45.02±25.04	33.76±28.43	47.61
InternVL + Grounded-CoT	48.35±18.05	47.94±19.36	17.44±19.01	18.10±14.96	32.96
+ ReAct	51.97±24.08	<u>56.56±31.30</u>	52.32±24.57	31.61±29.61	<u>48.11</u>
LLaVA + Grounded-CoT	50.62±18.74	52.30±28.89	<u>50.56±25.86</u>	31.69±27.00	46.29
+ ReAct	59.20±27.18	<u>56.82±31.71</u>	46.09±34.04	<u>32.23±31.73</u>	<u>48.59</u>
Qwen + FaithAct	55.10±20.14	57.22±27.85	56.87±24.29	66.45±27.87	58.91
InternVL + FaithAct	52.64±17.75	57.35±29.40	56.01±21.76	61.71±27.01	56.93
LLaVA + FaithAct	<u>52.82±22.77</u>	58.11±30.37	56.09±27.71	39.91±27.92	51.73

Table 1: **Faithfulness evaluation across reasoning paradigms.** We report the mean and standard deviation of the chain-level faithfulness score F_{chain} (in %) on four benchmarks. Methods are organized by *reasoning paradigm*: prompt-based reasoning (gray shading), tool-augmented reasoning (blue shading), and faithfulness-first planning (green shading), while holding backbone models fixed to enable controlled comparison. Best and second-best results within each backbone model are highlighted in **bold** and underline, respectively.

updated evidence from the evaluator, such as object existence labels, bounding boxes, and counts, to regenerate or adjust the unfaithful steps. The refinement prompt guides the model to produce a reasoning trace that remains consistent with perceptual evidence while maintaining logical continuity. This action-guided correction mechanism ensures that FAITHACT preserves both interpretability and grounding, progressively improving the reasoning chain through verified feedback. Further implementation and prompt details are provided in Appx. D.

5 Experiments

We conduct an extensive experimental study demonstrating the effectiveness of Faithful-First RPA framework in measuring and improving faithfulness in MLLM reasoning. Unless otherwise noted, all results are reported as percentages (%).

5.1 Experimental Setup

Datasets. We evaluate our framework on widely used multimodal benchmarks covering object recognition, visual grounding, and hallucination-sensitive question answering, including LLaVA-BENCH (Liu et al., 2023), REALWORLDQA (xAI, 2024), POPE (Li et al., 2023b), and MMHAL-BENCH (Sun et al., 2023) (MMHal). These datasets feature images with rich real-world objects and are well suited for evaluating perceptual faithfulness. Specifically, the queried objects in POPE are removed to ensure fair assessment.

Baselines. We compare our framework with representative training-free reasoning frameworks, in-

cluding CoT (Wei et al., 2022; Zhang et al., 2023), REACT (Yao et al., 2022), GROUNDED-CoT (Wu et al., 2025a), and VISUAL ABSTRACT THINKING (VAT) (Liu et al., 2025a). These methods are evaluated on widely used MLLMs, namely Qwen-2.5-VL-7B (Qwen) (Wang et al., 2024), InternVL3-8B (InternVL) (Zhu et al., 2025), and LLaVA-one-Vision-1.5-8B (LLaVA) (An et al., 2025b).

Hyper-parameters setup. The box threshold of GroundingDINO is 0.35 and the text threshold is 0.25. We keep the default generation settings for helper LLM and MLLMs.

5.2 Experimental Results and Analysis

Table 1 shows the performance of FAITHACT compared with selected baseline methods. The experimental results highlight several key findings.

Perceptual faithfulness is broadly underestimated. Across all evaluated models, faithfulness remains far from ideal. For instance, MLLMs such as Qwen achieve faithfulness scores of only around ~ 50 across datasets and evaluation metrics, substantially below the desired level (near 100). This highlights that current reasoning models often generate partially ungrounded or inconsistent reasoning traces, indicating that perceptual verification and evidential grounding are still open challenges.

FAITHACT improves faithfulness across models. Introducing the FAITHACT framework consistently enhances reasoning faithfulness across three tested models. For example, InternVL

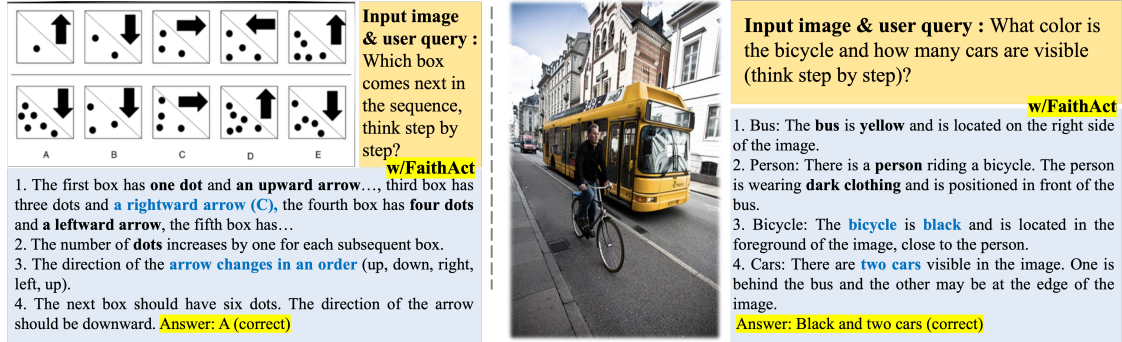


Figure 3: Qualitative comparison of reasoning chains generated with and without FAITHACT on two illustrative cases. In both tasks, FAITHACT enforces step-level perceptual verification, correcting hallucinated descriptions (red colored in Fig. 1) and producing more structured, visually grounded reasoning (colored in blue).

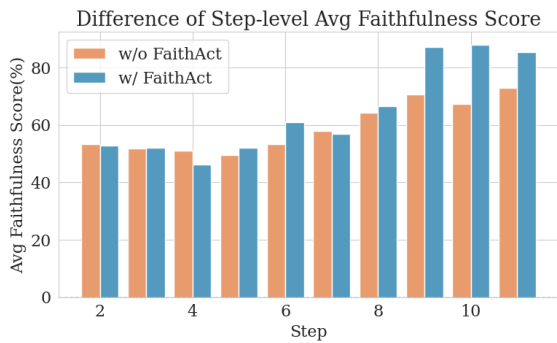


Figure 4: Distribution of average F_{step} difference across reasoning steps. The x-axis are reasoning steps and y-axis represents the F_{step} averaged difference between QWEN with and without FAITHACT.

achieves 57.35 ± 29.40 on REALWORLDQA with FAITHACT, compared to 44.23 ± 25.43 without it. Generally, our method attains the **highest** faithfulness in **11 out of 12** evaluated settings, demonstrating its effectiveness across different architectures and datasets. Averaged across models, FAITHACT achieves a mean score of 55.86%, outperforming the strongest baseline reasoning paradigm (ReAct, 48.10%) by 7.76 percentage points and CoT by 9.04%. We show that F_{chain} of REACT is theoretically bounded by that of FAITHACT (in Appx. F). We also note that the reported standard deviations are relatively large, which is expected given the instance-level nature of perceptual faithfulness.

FAITHACT mitigates hallucination. Although FAITHACT is not explicitly designed to mitigate hallucinations, it yields substantial gains on hallucination-focused benchmarks. In particular, on MMHAL, models with FAITHACT exhibit a marked reduction in hallucinated reasoning steps, achieving an average improvement of 21.99% and

9.81% over tool-augmented prompt-based reasoning frameworks, respectively. For example, Qwen with FAITHACT achieves 66.45 ± 27.87 , surpassing its second-best score. This notable gain indicates that the principle enforced by FAITHACT effectively constrains perceptual unfaithfulness during intermediate reasoning steps, thereby enhancing faithfulness of model outputs.

Improving faithfulness does not degrade task performance. Table 2 reports the performance of three models with and without FAITHACT across the two benchmarks. The results show that integrating FAITHACT preserves the model’s ability to generate correct final answers to multimodal questions, with slight improvements observed in two out of three datasets. These findings indicate that FAITHACT enhances the faithfulness of reasoning steps without harming the model’s original performance. Moreover, although behavioral faithfulness is not explicitly measured, the empirical evidence supports our hypothesis that perceptually grounded reasoning encourages behavioral consistency.

Faithfulness increases across reasoning steps.

Fig. 4 reports the distribution of F_{step} , when comparing FAITHACT and raw CoT without FAITHACT. We observe that the benefit of FAITHACT becomes particularly pronounced in the later reasoning steps, suggesting that its intervention is most effective when the model engages in deeper chains of reasoning. This observation is consistent with prior work (Wu et al., 2025b) which points out that excessively CoT increases susceptibility to noise in their later steps and thereby leads to more unfaithfulness.

Datasets & Models	REALWORLDQA (%)	MMHAL (rating)
Qwen + CoT	70.1	3.40
+FaithAct (ours)	74.5	3.48
InternVL + CoT	70.8	3.61
+FaithAct (ours)	71.2	3.58
LLaVA + CoT	68.1	3.41
+FaithAct (ours)	67.8	3.46

Table 2: Task performance with CoT and FAITHACT.

5.2.1 Qualitative Analysis and Case Studies.

We conduct a qualitative comparison to illustrate how FAITHACT improves reasoning faithfulness in the motivating cases shown in Fig. 1. In the first example (Fig. 3 top), the baseline model hallucinates *a yellow bicycle* and *no cars* by relying on language priors or implication (e.g., yellow bus) rather than image evidence. With FAITHACT, each reasoning step is perceptually verified, leading to the correct identification of a *black bicycle* and *two cars*. In the second example (bottom), both models correctly predict the next box in a visual sequence, but with FAITHACT guidance, reasoning chain is more structured and explicitly justifies each visual attribute transition. For example, the direction of upward arrow in (C) is corrected to rightward. These cases demonstrate that FAITHACT enhances perceptual grounding and visually consistent reasoning processes, empirically supporting and improving behavioral consistency.

5.3 Human Validation on Extracted Objects

To evaluate whether Qwen extracted objects (Sec. 3.1) accurately capture what is indeed stated in the text, we conducted a human validation study using 50 snippet and 25 annotators, producing a total of 7,550 object-level labels, detailed in Appx. (G). Compared with human judgments, the LLM shows near-perfect consistency, achieving a precision of **99.42%** with only 44 false positives (over-extraction rate = 0.58%). At snippet level, we compute *snippet validity*, defined as the probability that all extracted objects for a snippet are judged correct by a human annotator. Among 1,250 annotator–snippet pairs, the LLM attains a mean snippet validity of **0.9680**, showing extracted object sets are almost fully aligned with human.

5.4 Ablation and Sensitivity Analysis

We conduct ablation and sensitivity analyses to examine the contribution of individual components in the proposed framework. In particular, we focus on the two core functions `Poll()` and `Ground()`.




Figure 5: Comparative performance of the FAITHACT framework using SAM3 vs. GroundingDINO as the `Ground()` function. Results show mean accuracy (%) and standard deviation (error bars) across three MLLMs on the REALWORLDQA and MMHAL datasets.

Datasets & Models	REALWORLDQA (%)	MMHAL (%)
FaithAct	57.22±27.85	66.45±27.87
FaithAct (w/o Poll)	54.24±28.13	63.25±26.75
FaithAct (w/o Ground)	53.16±29.12	62.47±28.83

Table 3: Ablation study of the two core components.

Throughout these experiments, the faithfulness evaluation protocol remains unchanged and we disable one module at a time to isolate its effect.

Results on Qwen are reported in Table 3. Removing either `Poll()` or `Ground()` leads to a noticeable decrease in faithfulness (approximately 5%), with `Ground()` having a slightly larger impact. This suggests that object localization provides essential visual evidence for perceptually grounded reasoning. The results indicate that the two modules are complementary, and that FAITHACT achieves the greatest improvement in faithfulness when both are jointly enabled.

We additionally conduct a study replacing GroundingDINO with a recently published alternative, SAM3 (Carion et al., 2025). Equipped with SAM3, the performance of FAITHACT drops obviously (5% on REALWORLDQA and more on MMHAL). This may suggest that FAITHACT needs localization-specific models (like GroundingDINO) to provide grounding results more accurate with proper confidence.

we also choose Qwen on RealWorldQA for an additional ablation analysis over various threshold values. Results are reported in Table 4. Notably, the changes in hyperparameters do not influence the final performance of FaithAct too much. The threshold 0.5-0.5 will lead to a slight decline in performance.

α & threshold	0.3-0.7	0.4-0.6	0.5-0.5
0.5	55.76±27.55	54.15±26.09	52.52±26.63
0.6	55.84±27.59	55.96±23.73	52.58±24.62
0.7	57.44±26.24	57.22±27.85	52.72±21.47

Table 4: Sensitivity analysis on α and threshold.

5.5 Computational Overhead

We report the average time per chain generation (in seconds) using Qwen-2.5-VL-7B under identical decoding settings in Table 5. Importantly, our method does not modify the underlying generation function or decoding procedure of the MLLM. Consequently, the token generation rate (tokens/second) remains effectively the same as that of raw CoT. We also report reasoning-chain lengths in Appendix H.

5.6 Perception-light Benchmarks

we have additionally evaluated FaithAct on MathVista (Lu et al., 2024) and ScienceQA (Lu et al., 2022), where visual evidence is necessary but not sufficient and abstract or external knowledge is often required. Table 6 reports the F_{chain} of FaithAct and raw CoT (%) on these datasets. FaithAct performs well on ScienceQA and slightly better than CoT on MathVista.

6 Discussion: Behavioral Faithfulness

To provide more analysis and insights between behavioral faithfulness and perceptual faithfulness, we added a qualitative taxonomy to showing that perceptual faithfulness (PF) and behavioral faithfulness (BF) can dissociate across all four regimes (BF+ PF+; BF+ PF-; BF- PF+; BF- PF-, '+' and '-' mean faithful and unfaithful). Importantly, PF+ does not guarantee BF+, where a model may correctly ground intermediate observations but still produce an incorrect final decision due to downstream reasoning errors. Conversely, BF+ does not imply PF+, where a model can arrive at the correct answer while relying on perception-unfaithful rationales. We will also clarify a standard operationalization of behavioral faithfulness from prior works in the main text. For example, if the model truly uses the cited evidence, then removing or perturbing that evidence should systematically alter the model’s decision.

7 Conclusion

We introduced Faithful-first RPA, a framework that enforces perceptual grounding throughout the reasoning process. Within this framework, FAITHEVI

Datasets	LLAVA-BENCH	REALWORLDQA	MMHAL	POPE
FaithAct (s)	18.85	14.04	16.48	14.72
CoT (s)	10.73	4.65	5.59	3.41
δ Time	8.12	9.39	10.89	11.31

Table 5: Inference time comparison.

Dataset & Methods	MATHVISTA	SCIENCEQA
FaithAct	47.65±24.73	52.97±32.63
CoT	47.00±23.91	40.14±32.88

Table 6: Faithfulness evaluation on perception-light benchmarks.

provides a principled and fine-grained evaluation of perceptual faithfulness, while FAITHACT operationalizes faithfulness-first reasoning through planning and acting. Experiments across benchmarks show that our framework improves perceptual faithfulness by up to 24% without compromising accuracy, effectively mitigating hallucinations. These results underscore the value of faithfulness as a core design principle. Future work will extend to behavioral faithfulness and more challenging open-ended reasoning settings.

Limitations

This work primarily focuses on perceptual faithfulness and does not directly evaluate behavioral faithfulness, i.e., the alignment between reasoning traces and the model’s final decision process. While our empirical results suggest that enforcing perceptual grounding may be associated with more behaviorally consistent outputs, this relationship is not explicitly measured or guaranteed.

In addition, although we conduct a human validation study indicating that LLM-extracted objects are generally accurate, we do not perform large-scale human evaluations of step-level or chain-level perceptual faithfulness. Complementary human studies at different levels of granularity could help further contextualize and validate our findings.

Finally, our current implementation verifies perceptual faithfulness primarily at the level of object existence. Extending the Faithful-first RPA framework to incorporate attribute- and relation-level verification remains an important direction for future work, and may further reduce perceptual unfaithfulness in cases where objects are present but their properties or relations are mischaracterized.

Acknowledgements

This work was supported in part by computing resources provided by Lambda AI.

References

- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. Don't just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4971–4980.
- Ruichuan An, Sihan Yang, Ming Lu, Renrui Zhang, Kai Zeng, Yulin Luo, Jiajun Cao, Hao Liang, Ying Chen, Qi She, and 1 others. 2024. Mc-llava: Multi-concept personalized vision-language model. *arXiv preprint arXiv:2411.11706*.
- Ruichuan An, Sihan Yang, Renrui Zhang, Zijun Shen, Ming Lu, Gaole Dai, Hao Liang, Ziyu Guo, Shilin Yan, Yulin Luo, and 1 others. 2025a. Unictokens: Boosting personalized understanding and generation via unified concept tokens. *arXiv preprint arXiv:2505.14671*.
- Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Didi Zhu, and 1 others. 2025b. Llava-onevision-1.5: Fully open framework for democratized multimodal training. *arXiv preprint arXiv:2509.23661*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2425–2433.
- Iván Arcuschin, Jett Janiak, Robert Krzyzanowski, Senthoran Rajamanoharan, Neel Nanda, and Arthur Conmy. 2025. Chain-of-thought reasoning in the wild is not always faithful. *arXiv preprint arXiv:2503.08679*.
- Fazl Barez, Tung-Yu Wu, Iván Arcuschin, Michael Lan, Vincent Wang, Noah Siegel, Nicolas Collignon, Clement Neo, Isabelle Lee, Alasdair Paren, and 1 others. 2025. Chain-of-thought is not explainability. *Preprint, alphaXiv*, page v2.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623. ACM.
- Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, and 1 others. 2025. Sam 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*.
- Qiguang Chen, Libo Qin, Jin Zhang, Zhi Chen, Xiao Xu, and Wanxiang Che. 2024a. M3cot: A novel benchmark for multi-domain multi-step multi-modal chain-of-thought. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8199–8221.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, and 1 others. 2024b. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Sanjoy Dasgupta, Nave Frost, and Michal Moshkovitz. 2022. Framework for evaluating faithfulness of local explanations. In *International Conference on Machine Learning*, pages 4794–4815. PMLR.
- Timin Gao, Peixian Chen, Mengdan Zhang, Chaoyou Fu, Yunhang Shen, Yan Zhang, Shengchuan Zhang, Xiawu Zheng, Xing Sun, Liujuan Cao, and 1 others. 2024. Cantor: Inspiring multimodal chain-of-thought of mllm. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 9096–9105.
- Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Utkarsh Tyagi, Oriol Nieto, Zeyu Jin, and Dinesh Manocha. 2025. Visual description grounding reduces hallucinations and boosts reasoning in lvlms. In *The Thirteenth 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Wentao Hu, Wengyu Zhang, Yiyang Jiang, Chen Jason Zhang, Xiaoyong Wei, and Li Qing. 2025. Removal of hallucination on hallucination: Debate-augmented rag. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15839–15853.
- Yushi Hu, Benlin Liu, Jungo Kasai, Yizhong Wang, Mari Ostendorf, Ranjay Krishna, and Noah A. Smith. 2023. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. *arXiv preprint arXiv:2303.11897*.
- Alon Jacovi and Yoav Goldberg. 2020. Towards faithfully interpretable nlp systems: How should we define and evaluate faithfulness? In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 4198–4205.
- Liqiang Jing, Ruosen Li, Yunmo Chen, and Xinya Du. 2023. Faithscore: Fine-grained evaluations of hallucinations in large vision-language models. *arXiv preprint arXiv:2311.01477*.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, and 1 others. 2023. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*.

- Junxian Li, Di Zhang, Xunzhi Wang, Zeying Hao, Jingdi Lei, Qian Tan, Cai Zhou, Wei Liu, Yaotian Yang, Xinrui Xiong, and 1 others. 2025a. Chemvlm: Exploring the power of multimodal large language models in chemistry area. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 415–423.
- Pengyu Li, Lingling Zhang, Zhitao Gao, Yanrui Wu, Yuxuan Dong, Huan Liu, Bifan Wei, and Jun Liu. 2026. AGT^{AO}: Robust and stabilized llm unlearning via adversarial gating training with adaptive orthogonality. *arXiv preprint arXiv:2602.01703*.
- Xuhong Li, Mengnan Du, Jiamin Chen, Yekun Chai, Himabindu Lakkaraju, and Haoyi Xiong. 2023a. M⁴: A unified xai benchmark for faithfulness evaluation of feature attribution methods across metrics, modalities and models. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*.
- Yanshu Li, Yi Cao, Hongyang He, Qisen Cheng, Xiang Fu, Xi Xiao, Tianyang Wang, and Ruixiang Tang. 2025b. M²IV: Towards efficient and fine-grained multimodal in-context learning via representation engineering. In *Second Conference on Language Modeling*.
- Yanshu Li, Tian Yun, Jianjiang Yang, Pinyuan Feng, Jinfa Huang, and Ruixiang Tang. 2025c. Taco: Enhancing multimodal in-context learning via task mapping-guided sequence configuration. *arXiv preprint arXiv:2505.17098*.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Honglin Lin, Zheng Liu, Yun Zhu, Chonghan Qin, Juekai Lin, Xiaoran Shang, Conghui He, Wentao Zhang, and Lijun Wu. 2026a. Mmfineason: Closing the multimodal reasoning gap via open data-centric methods. *arXiv preprint arXiv:2601.21821*.
- Honglin Lin, Chonghan Qin, Zheng Liu, Qizhi Pei, Yu Li, Zhanping Zhong, Xin Gao, Yanfeng Wang, Conghui He, and Lijun Wu. 2026b. Scientific image synthesis: Benchmarking, methodologies, and downstream utility. *arXiv preprint arXiv:2601.17027*.
- Weifeng Lin, Xinyu Wei, Ruichuan An, Tianhe Ren, Tingwei Chen, Renrui Zhang, Ziyu Guo, Wentao Zhang, Lei Zhang, and Hongsheng Li. 2025. Perceive anything: Recognize, explain, caption, and segment anything in images and videos. *arXiv preprint arXiv:2506.05302*.
- Dairu Liu, Ziyue Wang, Minyuan Ruan, Fuwen Luo, Chi Chen, Peng Li, and Yang Liu. 2025a. Visual abstract thinking empowers multimodal reasoning. *arXiv preprint arXiv:2505.20164*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *Advances in neural information processing systems*, volume 36, pages 34892–34916.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, and 1 others. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer.
- Zheng Liu, Mengjie Liu, Jingzhou Chen, Jingwei Xu, Bin Cui, Conghui He, and Wentao Zhang. 2025b. Fusion: Fully integration of vision-language representations for deep cross-modal understanding. *arXiv preprint arXiv:2504.09925*.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *The Twelfth International Conference on Learning Representations*.
- Pan Lu, Swaroop Mishra, Tanglin 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. *Advances in neural information processing systems*, 35:2507–2521.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP-AACL 2023)*.
- Weijian Ma, Shizhao Sun, Tianyu Yu, Ruiyu Wang, Tat-Seng Chua, and Jiang Bian. 2026a. Thinking with blueprints: Assisting vision-language models in spatial reasoning via structured object representation. *arXiv preprint arXiv:2601.01984*.
- Xiaoxu Ma, Xiangbo Zhang, and Zhenyu Weng. 2026b. Stable and explainable personality trait evaluation in large language models with internal activations. *arXiv preprint arXiv:2601.09833*.
- Katie Matton, Robert Osazuwa Ness, John Guttag, and Emre Kiciman. 2025. Walk the talk? measuring the faithfulness of large language model explanations. *arXiv preprint arXiv:2504.14150*.
- Yifei Ming, Senthil Purushwalkam, Shrey Pandit, Zixuan Ke, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. 2024. Faitheval: Can your language model stay faithful to context, even if “the moon is made of marshmallows”. *arXiv preprint arXiv:2410.03727*.

- Letitia Parcalabescu and Anette Frank. 2023. On measuring faithfulness or self-consistency of natural language explanations. *arXiv preprint arXiv:2311.07466*.
- Jingyu Peng, Maolin Wang, Nan Wang, Jiatong Li, Yuchen Li, Yuyang Ye, Wanyu Wang, Pengyue Jia, Kai Zhang, and Xiangyu Zhao. 2025a. Logic jailbreak: Efficiently unlocking llm safety restrictions through formal logical expression. *arXiv preprint arXiv:2505.13527*.
- Jingyu Peng, Maolin Wang, Xiangyu Zhao, Kai Zhang, Wanyu Wang, Pengyue Jia, Qidong Liu, Ruocheng Guo, and Qi Liu. 2025b. Stepwise reasoning disruption attack of llms. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5040–5058.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmlR.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. *Advances in Neural Information Processing Systems*, 37:8612–8642.
- Linzhuang Sun, Hao Liang, Jingxuan Wei, Bihui Yu, Tianpeng Li, Fan Yang, Zenan Zhou, and Wentao Zhang. 2025. Mm-verify: Enhancing multimodal reasoning with chain-of-thought verification. *arXiv preprint arXiv:2502.13383*.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liangyan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. 2023. Aligning large multimodal models with factually augmented rlhf. *arXiv:2309.14525*.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liangyan Gui, Yu-Xiong Wang, Yiming Yang, and 1 others. 2024. Aligning large multimodal models with factually augmented rlhf. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13088–13110.
- 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. *Advances in Neural Information Processing Systems*, 36:74952–74965.
- Haozhe Wang, Chao Qu, Zuming Huang, Wei Chu, Fangzhen Lin, and Wenhui Chen. 2025. Vl-rethinker: Incentivizing self-reflection of vision-language models with reinforcement learning. *arXiv preprint arXiv:2504.08837*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Xi Wang, Songlei Jian, Shasha Li, Xiaopeng Li, Zhaoye Li, Bin Ji, Baosheng Wang, and Jie Yu. 2026. Jpu: Bridging jailbreak defense and unlearning via on-policy path rectification. *arXiv preprint arXiv:2601.03005*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Zheng Weihua, Roy Ka-Wei Lee, Zhengyuan Liu, Wu Kui, Aiti Aw, and Bowei Zou. 2025. Ccl-xcot: An efficient cross-lingual knowledge transfer method for mitigating hallucination generation. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 1768–1788.
- Qiong Wu, Xiangcong Yang, Yiyi Zhou, Chenxin Fang, Baiyang Song, Xiaoshuai Sun, and Rongrong Ji. 2025a. Grounded chain-of-thought for multimodal large language models. *ArXiv*, abs/2503.12799.
- Yuyang Wu, Yifei Wang, Ziyu Ye, Tianqi Du, Stefanie Jegelka, and Yisen Wang. 2025b. When more is less: Understanding chain-of-thought length in llms. *arXiv preprint arXiv:2502.07266*.
- xAI. 2024. Announcing grok-1.5. <https://x.ai/news/grok-1.5>.
- Xi Xiao, Yunbei Zhang, Xingjian Li, Tianyang Wang, Xiao Wang, Yuxiang Wei, Jihun Hamm, and Min Xu. 2025a. Visual instance-aware prompt tuning. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 2880–2889.
- Xi Xiao, Yunbei Zhang, Lin Zhao, Yiyang Liu, Xiaoying Liao, Zheda Mai, Xingjian Li, Xiao Wang, Hao Xu, Jihun Hamm, and 1 others. 2025b. Prompt-based adaptation in large-scale vision models: A survey. *arXiv preprint arXiv:2510.13219*.
- Zhenhua Xu, Dongsheng Chen, Shuo Wang, Jian Li, Chengjie Wang, Meng Han, and Yabiao Wang. 2026. Adamarp: An adaptive multi-agent interaction framework for general immersive role-playing. *arXiv preprint arXiv:2601.11007*.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, and 1 others. 2024. Qwen2 technical report. arxiv 2024. *arXiv preprint arXiv:2407.10671*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language

models. In *The eleventh international conference on learning representations*.

T. Yu, Y. Zhang, and J. Li. 2024. RLhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yuanjun Zhang, Fuzel Ahamed Shaik, Suvojit Achary, Fahad Khalid, and Mourad Oussalah. 2026. Towards reliable multimodal disaster severity assessment through preference optimization and explainable vision-language reasoning. *Reliability Engineering & System Safety*, page 112674.

Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alexander J. Smola. 2023. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.

Yuchen Zhou, Jiayu Tang, Shuo Yang, Xiaoyan Xiao, Yuqin Dai, Wenhao Yang, Chao Gou, Xiaobo Xia, and Tat-Seng Chua. 2025. Logic unseen: Revealing the logical blindspots of vision-language models. *arXiv preprint arXiv:2508.11317*.

Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, and 1 others. 2025. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.

A Related Work

Reasoning Frameworks for MLLMs. MLLMs have developed surprising reasoning capabilities in multimodal problem solving, task planning, scientific discovery and so on (Lu et al., 2024; Li et al., 2025a; Gao et al., 2024; An et al., 2024, 2025a; Lin et al., 2025; Li et al., 2025b,c; Xiao et al., 2025b; Xu et al., 2026; Zhang et al., 2026; Wang et al., 2026; Lin et al., 2026b; Liu et al., 2025b). To enhance systematic reasoning, several frameworks (Chen et al., 2024a; Sun et al., 2025; Shao et al., 2024; Zhou et al., 2025; Wang et al., 2025; Xiao et al., 2025a; Ma et al., 2026b; Lin et al., 2026a) have been proposed to decompose multimodal problems into interpretable steps. Among them, CoT (Chen et al., 2024a), GROUNDED-CoT (Wu et al., 2025a) and REACT frameworks (Yao et al., 2022) are famous and insightful ones. More recently, VAT (Liu et al., 2025a) explores hierarchical and compositional reasoning, encouraging models to form abstract visual concepts that support complex decision-making. Such generated reasoning chains, though often fluent and logically structured, may still include steps unsupported by

visual evidence or inconsistent with the model’s actual decision process (Yu et al., 2024).

Faithfulness in Multimodal Reasoning. Despite the demonstrated effectiveness of reasoning in improving the performance of MLLMs, an unfaithful problem emerges: the reasoning traces generated by these models do **not always behave faithfully** to the internal processes that produce their final answers (Lyu et al., 2023; Turpin et al., 2023; Lanham et al., 2023). Models often rely on latent knowledge or shortcut associations that are not explicitly expressed in their reasoning chains. As a result, the generated reasoning steps may read as a plausible but **untrustworthy explanation**, motivating the need for explicit faithfulness assessment (Jacovi and Goldberg, 2020; Parcalabescu and Frank, 2023; Matton et al., 2025; Barez et al., 2025; Arcuschin et al., 2025; Peng et al., 2025a,b; Weihua et al., 2025). Several works have introduced benchmarks or metrics to evaluate faithfulness explicitly. M⁴ (Li et al., 2023a), FAITHSCORE (Jing et al., 2023) and TIFA (Hu et al., 2023) propose metrics for evaluating the faithfulness for vision-language models. However, we find that behavioral alignment (behavioral faithfulness) does not guarantee the correctness of the final output (see *left panel* in Fig.1).

Object Hallucination as Unfaithful Consequences. Object hallucination is identified as a common challenge, as a consequence, in Large Vision-Language Models (LVLMs), where the model describes or reasons about objects absent in the input image. Several benchmarks, such as POPE (Li et al., 2023b) and MMHAL-BENCH (Sun et al., 2024), have been developed to systematically evaluate this phenomenon. However, these efforts typically treat hallucination as an isolated failure mode rather than as a manifestation of broader **unfaithful reasoning**.

Recent studies address hallucination through different mechanisms: (i) *training-time alignment*, such as hallucination-aware REINFORCEMENT LEARNING FROM HUMAN FEEDBACK or preference optimization (Sun et al., 2024; Hu et al., 2025); (ii) *decoding-time constraints*, including grounded or contrastive decoding (Li et al., 2023b); and (iii) *feature-level grounding strategies* that enhance cross-modal alignment (Ghosh et al., 2025). Together, these advances underscore that faithful reasoning requires not only linguistic coherence but also perceptual accountability, further motivating to verify evidential grounding before inference.

B Object Extraction Prompts and Examples

Extract all objects mentioned in the following sentence that may occur in an image. Only extract nouns meaning objects, not abstract adjectives, concepts, actions, general nouns or locations. Do not include non-object nouns or words like “Image”, “Object”, “Feature”, or “Photo”. \n \n####{One Reasoning Step}####\n \nReturn only a list of nouns like [“xxx”, “xxx”, “xxx”] and do not include any other things. If no available nouns, return an empty list [].

Example 1:

Text: ****Location Context****: The presence of a coastal area with a beach and a city in the background suggests a location near the ocean.

Extract Result: [“coastal area”, “beach”, “city”]

Example 2:

Text: ****Setting****: The image appears to be taken on a city street, likely in an urban area given the presence of taxis and buildings in the background.

Extract Result: [“taxis”, “buildings”]

C Training Details on POPE

The training set of POPE released on Huggingface is [here](#). It consists of three parts, random, popular and adversarial, all with image-object existence labels. The total training size is 9000 items. We utilize the released test set for testing model performance.

The CLIP+polling head model is trained on two NVIDIA RTX4090 48GB GPUs. During training, we freeze the backbone CLIP, and only set the head trainable. We set the batch size to 32 each GPU, learning rate to $1e - 3$, and train 50 epochs with early stop. The final test accuracy of the model is 99.80%, with 9000 real-world test examples in POPE. We test the trained model in the wild, too. And we discover that it can reliably tell whether an object exists in an image. Thus, we can safely use it in our preference polling task.

Algorithm 1 Faithfulness-First Planner (FAITHACT)

Require: Image I , textual query Q , MLLM planner M , helper LLM H_M

Ensure: Faithful reasoning output R

```

1: // Step 1: Initial Reasoning
2:  $R_{\text{raw}} \leftarrow M.\text{Reason}(I, Q)$ 
3: // Step 2: Extract Claimed Objects
4:  $O_{\text{raw}} \leftarrow H_M.\text{ExtractObj}(R_{\text{raw}})$ 
5: // Step 3: Verification and Function Calls
6: for each object  $o_i \in O_{\text{raw}}$  do
7:    $c_p^i \leftarrow \text{Poll}(o_i)$  {Existence confidence}
8:    $B_i, c_g^i \leftarrow \text{Ground}(o_i)$  {Bounding boxes and spatial scores}
9:    $f^i \leftarrow \text{Select}(c_p^i, c_g^i)$  or  $\text{Abstain}(c_p^i, c_g^i)$ 
   {Faithfulness threshold check}
10:  if  $\text{Select}$  then
11:     $n_i \leftarrow \text{Count}(B_i, c_g^i)$  {Count reliably grounded instances}
12:    record  $(o_i, f^i, n_i, B_i)$ 
13:  end if
14: end for
15: // Step 4: Refine-Based Faithful Reasoning
16:  $R_{\text{new}} \leftarrow M.\text{Reason}(I, Q, O_{\text{raw}}, \{f^i, B_i, n_i\})$ 
17: return  $R_{\text{new}}$ 

```

D Prompts for FAITHACT

Question: {your original question}.\n \n Model Response: {MLLM’s original response}\n \n Additional location information:\n \n {Information from the functions}\n \n Using only the “exists” objects with high confidence and avoid using objects that do not exist. Do not include new objects or descriptions. Do not repeat the evidences, confidence scores and bounding boxes in your reasoning. Think step by step. Steps should be like: 1.<object1>:xxx\n \n 2.<object2>:xxx\n \n ... \n \n ..., .

E Algorithm of FAITHACT

Here we list the algorithm process of FAITHACT in Algorithm 1.

F Proof of faithfulness by FaithAct v.s. ReAct

Lemma 1 (Faithfulness Dominance of FaithAct over ReAct). Let $R^{\text{ReAct}} = \{s_t^{\text{ReAct}}\}_{t=1}^T$ and $R^{\text{FaithAct}} = \{s_t^{\text{FaithAct}}\}_{t=1}^{T'}$ denote the reasoning chains generated by ReAct and FaithAct, respectively. Let $F_{\text{step}}(s_t)$ be the perceptual faithfulness

Dataset & Metrics	LLAVA-BENCH	REALWORLDQA	MMHAL	POPE
Avg token length (FaithAct)	270.88	171.43	112.85	149.11
Avg token length (CoT)	357.52	162.78	200.06	148.23
Min token length (FaithAct)	30	17	12	28
Min token length (CoT)	98	64	45	43
Max token length (FaithAct)	512	512	510	512
Max token length (CoT)	512	432	502	512

Table 7: Token length statistics for FaithAct and CoT across different benchmarks.

of step s_t , and define the chain-level faithfulness as

$$F_{\text{chain}}(R) = \frac{1}{|R|} \sum_{t=1}^{|R|} F_{\text{step}}(s_t). \quad (12)$$

Assume that FaithAct refines each candidate step $s_t^{(k)}$ using verified evidence such that

$$F_{\text{step}}(s_t^{(k+1)}) \geq F_{\text{step}}(s_t^{(k)}), \quad (13)$$

and accepts only refined steps satisfying $F_{\text{step}}(s_t) \geq c$ for some threshold $c \in [0, 1]$. Then

$$F_{\text{chain}}(R^{\text{FaithAct}}) \geq F_{\text{chain}}(R^{\text{ReAct}}). \quad (14)$$

Proof. For each semantic subgoal g , let s_g^{ReAct} denote the step generated by ReAct and $s_g^{(0)}$ the initial unverified step proposed by the same MLLM within FaithAct. FaithAct refines $s_g^{(0)}$ through iterative verification:

$$s_g^{(0)} \rightarrow s_g^{(1)} \rightarrow \dots \rightarrow s_g^{(K_g)} = s_g^{\text{FaithAct}}.$$

By monotonicity in (13),

$$F_{\text{step}}(s_g^{\text{FaithAct}}) \geq F_{\text{step}}(s_g^{(0)}) = F_{\text{step}}(s_g^{\text{ReAct}}), \forall g. \quad (15)$$

If FaithAct drops an unverified claim (via `Abstain()`), it effectively removes a low-faithfulness step, which cannot decrease the average of the remaining step scores. Let A and B denote the multisets of step scores in FaithAct and ReAct, respectively. Then every element in A dominates or replaces an element in B with greater or equal score. Removing low-valued elements weakly increases the mean, hence

$$\frac{1}{|A|} \sum_{a \in A} a \geq \frac{1}{|B|} \sum_{b \in B} b.$$

By definition, this is equivalent to

$$F_{\text{chain}}(R^{\text{FaithAct}}) \geq F_{\text{chain}}(R^{\text{ReAct}}).$$

□

This result follows directly from FaithAct’s *verify-and-refine* constraint: each reasoning step is either (i) retained and refined until it is perceptually grounded, or (ii) rejected through `Abstain()` if unsupported, ensuring that no unverified or hallucinated step reduces overall faithfulness.

Corollary 1 (Strict Improvement Under Unfaithful Steps). *Under the assumptions of Lemma 1, suppose there exists at least one subgoal g^* such that the ReAct step $s_{g^*}^{\text{ReAct}}$ is perceptually unfaithful, i.e.,*

$$F_{\text{step}}(s_{g^*}^{\text{ReAct}}) < 1. \quad (16)$$

Assume further that FaithAct either (i) refines this step into a perceptually grounded step $s_{g^*}^{\text{FaithAct}}$ with

$$F_{\text{step}}(s_{g^*}^{\text{FaithAct}}) > F_{\text{step}}(s_{g^*}^{\text{ReAct}}), \quad (17)$$

or (ii) rejects the claim via `Abstain()`, thereby removing $s_{g^*}^{\text{ReAct}}$ entirely from its chain. Then

$$F_{\text{chain}}(R^{\text{FaithAct}}) > F_{\text{chain}}(R^{\text{ReAct}}). \quad (18)$$

Proof. Case (i): If FaithAct refines $s_{g^*}^{\text{ReAct}}$ into $s_{g^*}^{\text{FaithAct}}$ with strictly higher step-level faithfulness, then by the stepwise dominance in (15), FaithAct strictly improves at least one term in the average while leaving all other terms weakly improved. The mean of a set of real numbers strictly increases if at least one element increases and no element decreases. Hence $F_{\text{chain}}(R^{\text{FaithAct}}) > F_{\text{chain}}(R^{\text{ReAct}})$.

Case (ii): If FaithAct abstains on g^* , then ReAct includes a low-faithfulness step $s_{g^*}^{\text{ReAct}}$ in its average, while FaithAct omits it. Removing a strictly sub-maximal element from an arithmetic mean strictly increases that mean, provided the remaining elements are not all equal to that element. Since $F_{\text{step}}(s_{g^*}^{\text{ReAct}}) < 1$ by assumption, this condition holds. Therefore the average step score of FaithAct is strictly higher than that of ReAct.

In both cases,

$$F_{\text{chain}}(R^{\text{FaithAct}}) > F_{\text{chain}}(R^{\text{ReAct}}).$$

□

Empirically (Table 1), we observe that the inequality is typically strict, consistent with Corollary 1: whenever ReAct produces at least one perceptually ungrounded step, FaithAct either corrects it using verified evidence or removes it rather than propagating hallucinated content.

G Human Validation of LLM-Extracted Objects

This section provides detailed methodology and analysis for the human validation study used to assess the accuracy of LLM-extracted objects described in Sec. 3.1.

G.1 Study Design and Data Collection

We randomly sampled 50 text snippets from our evaluation corpus. Each snippet was processed by the Qwen to extract a set of candidate objects that were intended to represent entities explicitly or implicitly stated in the text. To validate the correctness of these extracted objects, we conducted a human annotation study involving 25 annotators.

Each annotator independently evaluated all extracted objects for each snippet. For every object, annotators answered the binary question: “*whether the extracted object explicitly exists in the text.*” Objects were labeled as *supported* (1) if they were explicitly mentioned or unambiguously implied by the text, and as *unsupported* (0) otherwise. Annotators were instructed to rely solely on the provided text and to avoid using external world knowledge. When uncertain, they were instructed to mark the object as unsupported. Notably, each annotator will receive about 10 dollars for participating. We have explained that their annotations will be used only for evaluating the correctness of LLM outputs, and all of them agree with this.

In total, the study yielded 7,550 object-level labels, corresponding to 1,250 annotator–snippet pairs (50 snippets \times 25 annotators).

G.2 Evaluation Metrics

We first evaluate precision at the object level by comparing LLM-extracted objects against human judgments. An extracted object is considered a false positive if it is labeled as unsupported by a human annotator. We compute object-level precision as

$$\text{Precision} = \frac{\# \text{ supported objects}}{\# \text{ extracted objects}}.$$

Across all object-level annotations, the LLM achieves a precision of 99.42%, with only 44 false

positives, corresponding to an over-extraction rate of 0.58%.

While precision metrics capture local correctness, they do not reflect whether an entire set of extracted objects for a snippet is jointly accurate. To assess holistic correctness, we introduce *snippet validity*, defined as an indicator function over annotator–snippet pairs:

$$\text{SV}(a, s) = \mathbf{1}[\forall o \in \mathcal{O}_s, o \text{ is supported by annotator } a],$$

where \mathcal{O}_s denotes the set of objects extracted for snippet s .

We report mean snippet validity by averaging $\text{SV}(a, s)$ across all annotator–snippet pairs:

$$\text{Mean SV} = \frac{1}{|\mathcal{A}||\mathcal{S}|} \sum_{a \in \mathcal{A}} \sum_{s \in \mathcal{S}} \text{SV}(a, s).$$

Across 1,250 annotator–snippet pairs, the LLM achieves a mean snippet validity of 0.97. This indicates that for nearly all annotator–snippet evaluations, *all* extracted objects for a snippet are judged correct.

Summary The combined object-level and snippet-level analyses provide complementary views of the choice of Qwen as an object extractor. High object-level precision and high snippet validity demonstrates that extracted object sets are almost always entirely accurate, serving as reliable inputs for downstream faithfulness evaluation in our framework.

H Reasoning-Chain Length

We report the reasoning chain length of our method of Qwen-2.5-VL-7B in Table 7, compared with raw CoT. The official tokenizer of the model is used. Here we report the metrics calculated from the final output of FaithAct. We find that FaithAct does not produce longer reasoning chains than raw CoT. In fact, on certain metrics, it yields shorter chains. A plausible explanation is that FaithAct constrains the model to reason only over visually verified objects with acceptable faithfulness scores. By filtering out unsupported or hallucinated objects, the framework thereby leads to more concise and focused chains.

Current id: 1 / 50
Complete ids: 0

Task Instructions

In this study, you will be shown a text snippet and a list of objects automatically extracted from that text by a language model. Your task is to decide, for each object, whether it explicitly exists in the text.

- Select **Supported (1)**: if the object is explicitly mentioned in the text or its presence is clearly implied.
- Select **Unsupported (0)**: if the object is not mentioned, only weakly or ambiguously implied from general world knowledge.

Text snippet

1. **Artistic Appreciation**

- **Classical Art Enthusiasts**: These viewers might appreciate the cleverness of the artist in reimagining a masterpiece. They could see it as a nod to the original painting while also celebrating the unique charm of the dachshund. The blending of classical art techniques with modern pop culture can be seen as a form of artistic innovation.
- **Critics of Modern Art**: Some might view this as a disrespectful or irreverent alteration of a revered piece of art. They could argue that such modifications undermine the integrity and historical significance of the original work.

Object list

Object	Supported (1)	Unsupported (0)
dachshund	<input type="radio"/>	<input type="radio"/>
classical art	<input type="radio"/>	<input type="radio"/>
modern pop culture	<input type="radio"/>	<input type="radio"/>
artistic innovation	<input type="radio"/>	<input type="radio"/>

Figure 6: **Human annotation interface for object-existence validation.** For each snippet, annotators are presented with the original text and a set of objects automatically extracted by the LLM. Annotators judge whether each object is explicitly mentioned or unambiguously implied by the text, producing binary Supported (1) or Unsupported (0) labels used for evaluating extraction faithfulness.