

THROUGH THE LENS OF CONTRAST: SELF-IMPROVING VISUAL REASONING IN VLMS

000
001
002
003
004
005 **Anonymous authors**
006 Paper under double-blind review
007
008
009
010

ABSTRACT


011 Reasoning has emerged as a key capability of large language models. In linguis-
012 tic tasks, this capability can be enhanced by self-improving techniques that refine
013 reasoning paths for subsequent finetuning. However, extending these language-
014 based self-improving approaches to vision language models (VLMS) presents a
015 unique challenge: visual hallucinations in reasoning paths cannot be effectively
016 verified or rectified. Our solution starts with a key observation about visual con-
017 trast: when presented with a contrastive VQA pair, *i.e.*, two visually similar im-
018 ages with synonymous questions, VLMS identify relevant visual cues more pre-
019 cisely compared with when given a single VQA sample. Motivated by this ob-
020 servation, we propose **Visual Contrastive Self-Taught Reasoner** (VC-STaR), a
021 novel self-improving framework that leverages visual contrast to mitigate hal-
022 lucinations in model-generated rationales. We collect a diverse suite of VQA
023 datasets, curate contrastive pairs according to multi-modal similarity, and gen-
024 erate rationales using VC-STaR. Consequently, we obtain a new visual reasoning
025 dataset, VisCoR-55K, which is then used to boost the reasoning capability of var-
026 ious VLMS through supervised finetuning. Extensive experiments show that VC-
027 STaR not only outperforms existing self-improving approaches but also surpasses
028 models finetuned on the SoTA visual reasoning datasets, demonstrating that the
029 inherent contrastive ability of VLMS can bootstrap their own visual reasoning.
030 The code, dataset and trained models will be released upon acceptance.

1 INTRODUCTION


031 The scaling of large language models (LLM) has led to the emergence of reasoning capabilities (Wei
032 et al., 2022a), making a transition from *System 1* to *System 2* (Kahneman, 2011) and enabling lan-
033 guage models to tackle complex, multi-step problems (Wei et al., 2022b; Kojima et al., 2022). This
034 emergent ability can be further enhanced by various techniques (Wang et al., 2023b; Li et al., 2023b;
035 Hao et al., 2023; Gao et al., 2023; OpenAI, 2024b; Guo et al., 2025). Among them, self-improving
036 approaches (Zelikman et al., 2022; Gulcehre et al., 2023; Madaan et al., 2023; Qu et al., 2024; Ma
037 et al., 2025) form a prominent branch, mainly because they can be easily applied and extended with-
038 out external reward models (Lu et al., 2024a), predefined step decomposition (Liu et al., 2025), or
039 specially designed reasoning structures (Li et al., 2025).

040 However, it is infeasible to directly adapt such language-based self-improving methods to vision
041 language models (VLMS) (Liu et al., 2023; Bai et al., 2025). Previous self-improving approaches
042 focus on textual coherence and the quality of the final answer (Zelikman et al., 2022; Zhang et al.,
043 2024a), while they are unable to verify or rectify the visual hallucinations that persist in current
044 VLMS (Tong et al., 2024; Li et al., 2024). Even worse, they may get stuck in speculative reasoning
045 that privileges textual priors over real visual evidences (Favero et al., 2024; Wu et al., 2025). We
046 claim that the key problem for self-improving in VLMS is: *how to rectify visual hallucinations in*
047 *VLMS’ reasoning paths for high-quality visual rationale generation.*

048 Our solution is built upon an interesting observation: *VLMS can see better during contrasting*. As
049 shown in Fig. 1a, the VLMS generates a wrong rationale with visual hallucinations given a single
050 visual question answering (VQA) sample. Instead, when presented with a contrastive VQA pair,
051 *i.e.*, two similar images with synonymous questions (Setting C in sub-figure 1b), the model captures
052 fine-grained visual evidence more accurately and rectifies the erroneous rationale. Statistics of this



(a) Visual hallucinations within the reasoning paths can mislead the model. Contrasting within a contrastive VQA pair, the VLM may rectify its own hallucinations.



(b) Results and statistics of rectifying hallucinatory outputs by three settings. H: with hint; C: via contrasting.

073
074
075
076
077
078
079
080
081
082

Figure 1: **Contrasting makes the VLM see better.** (a) Contrastive VQA pairs compel a more accurate response. (b) Compared with a previous self-improving method STaR (Zelikman et al., 2022) that enhances the quality of reasoning with hints (ground-truth answers), contrasting with hints can rectify more cases. The blocks along the x -axis mark initial VLM failures. The color of each block indicates the outcome of rectifying: green for success and gray for failure. Statistical analyses suggest that contrasting with hints produces minimal additional errors and rectify more erroneous cases. Tested VLM is Qwen2.5VL-7B (Bai et al., 2025).

083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

phenomenon on a group of failure cases are shown in Fig. 1b. Compared with the hints-only (Setting H in sub-figure 1b) self-improving (provide the model with the ground-truth answers), the hints and contrasting (Setting C&H in sub-figure 1b) setting not only prevents the model from making new errors but also leads to the rectification of its original hallucinations.

Motivated by this, we propose a new self-improving framework, **Visual Contrastive Self-Taught Reasoner** (VC-STaR). VC-STaR contains three steps: (1) *think step by step* and generate a coarse rationale; (2) *compare* visual queries in a contrastive VQA pair and provide a contrastive analysis; (3) *rethink* and refine the coarse rationale via an LLM based on the contrastive analysis. In order to guarantee the scalability of VC-STaR, we also propose a task-agnostic contrastive VQA pair curation framework, which can be readily adapted to various VQA tasks, *e.g.*, reasoning (Lu et al., 2021b), math (Gao et al., 2025), chart (Liu et al. (2024a), and OCR (Yuan et al., 2022). Specifically, we curate the contrastive VQA pairs within individual datasets, based on the similarity of both images and questions. We utilize these contrastive VQA pairs to generate faithful rationales, resulting in a novel **Visual Contrastive Reasoning** dataset (VisCoR-55K) as illustrated in Fig. 2. Finetuning with VisCoR-55K enhances VLMs' visual reasoning capability.

VC-STaR achieves prominent results on a wide range of challenging benchmarks, including MMVP (Tong et al., 2024), HallusionBench (Guan et al., 2024b), MathVista (Lu et al., 2024b), MathVision (Wang et al. (2024)), and MMStar (Chen et al., 2024d). On the one hand, VC-STaR outperforms existing self-improving baselines. On the other hand, it exhibits a clear advantage over models trained on recently proposed reasoning datasets. The experimental results validate that visual reasoning capability of VLMs can be bootstrapped through the lens of contrast.

2 RELATED WORKS

Reasoning in Language. Dual-system theory (Kahneman, 2011) illustrates two systems in human cognition: a fast, intuitive *System 1* and a slow, deliberate *System 2* which is akin to emergent




Figure 2: **VisCoR-55K.** We introduce the **Visual Contrastive Reasoning** dataset (VisCoR-55K), a new collection of 55K high-quality visual reasoning samples. Spanning the domains of general VQA, reasoning, math, graph/chart, and OCR, each sample is created by leveraging a contrastive counterpart to generate a faithful rationale. Rationales are shown in the Sec. A.3.

reasoning capability of LLMs (Wei et al., 2022a). Consequently, reasoning enhancement (Wei et al., 2022b; Kojima et al., 2022) is considered a pathway to elevate LLMs’ cognitive performance. One solution involves a reward model (Li et al., 2023b; Lu et al., 2024a), often coupled with Monte Carlo tree search (Hao et al., 2023; Zhang et al., 2024a), to discover optimal reasoning paths. However, this solution is constrained by the need of an auxiliary model and the requirement for reasoning step dividing (Liu et al., 2025). Another way employs macro reasoning actions (Gao et al., 2023; Khot et al., 2023; Yang et al., 2025a) to inject human prior knowledge, however, hand-crafted macro actions struggle to adapt to diverse reasoning scenarios. While reinforcement learning (Rafailov et al., 2023; Trung et al., 2024; Guo et al., 2025) has also attracted the attention, its success relies on the data format and the design of reward functions. Self-improving methods (Zhang et al., 2024a) offer a more scalable alternative, enabling LLMs to refine its own reasoning by constructing high-quality reasoning data (Wang et al., 2023b), utilizing ground-truth answers as hints (Zelikman et al., 2022), or leveraging internal feedback (Qu et al., 2024). With fewer external constraints, self-improving methods pave the way for more flexible and general language reasoners.

Reasoning in Vision. Human reasoning is stimulated not only by textual input but also by visually-related queries. Fostering the visual reasoning ability (Zhang et al., 2024c) for VLMs (Liu et al., 2023; Li et al., 2023a) is therefore a critical frontier topic. Early attempts often rely on external scaffolding like scene graphs (Mitra et al., 2024), macro actions (Xu et al., 2025; Dong et al., 2025), or bounding boxes highlighting key region in images (Shao et al., 2024). However, such approaches suffer from fundamental limitations: they are constrained by the data structure or tend to generate stereotyped reasoning paths. Despite these advances, the self-improving paradigm which has shown its effectiveness in text-only domain is underexplored for visual reasoning. The primary obstacle is the visual hallucinations embedded in reasoning paths cannot be easily rectified by existing text-centric self-improving frameworks (Zhang et al., 2024a; Zelikman et al., 2022; Qu et al., 2024). The proposed VC-STaR attempts to bridge this gap through the lens of contrast.

Power of Contrasting. Contrasting has shown effectiveness in a wide range of machine learning topics. By comparing different *views* (Tian et al., 2020), e.g., data augmentations, of the same sample while distinguishing them from others (Wang & Isola, 2020), contrastive self-supervised learning methods (He et al., 2020; Grill et al., 2020; Radford et al., 2021; Liang et al., 2022; Pan et al., 2023) excel at learning potent feature representations. Explicitly cross-image contrasting is also studied under uni-modal setting (Pan et al., 2023; Ding et al., 2024; Chen et al., 2024a) and




Figure 3: **Contrastive VQA pair curation pipeline.** To facilitate effective contrastive analysis, we curate corresponding challenging counterparts for VQA samples from a pool of diverse datasets. Each curated pair consists of two samples that share a synonymous question but feature distinct yet semantically similar images. Collected pairs are filtered by a difficulty-based sampling procedure.

multi-modal setting (Park et al., 2019; Kim et al., 2021; Yao et al., 2022; Dunlap et al., 2024). Based on these advancements, VLMs are endowed with robust capabilities for multi-image comprehension and comparison (Alayrac et al., 2022; Bai et al., 2025; Chameleon, 2025; Lin et al., 2025). Some prior works have leveraged contrasting to create better instruction-tuning data (Jiao et al., 2025; Ma et al., 2024). However, how contrasting can help visual reasoning remains an open question. We observe that VLMs’ inherent comparative ability can be repurposed to actively suppress its own visual hallucinations, bootstrapping their visual reasoning capability. This discovery offers a new perspective about the power of contrasting in reasoning.

3 VISUAL CONTRASTIVE SELF-TAUGHT REASONER (VC-STaR)

Let θ be a VLM and $\mathcal{D} = \{(v_i, q_i, a_i)\}_{i=1}^N$ be a set of visual question answering (VQA). The VQA set consists N triplets, where v_i , q_i , and a_i represent the i -th image, question, and corresponding ground-truth answer, respectively. Following previous self-taught reasoners (Zelikman et al., 2022; Madaan et al., 2023), the original VQA dataset \mathcal{D} can be enriched by generating a rationale r with θ for each triplet, which transforms \mathcal{D} into a visual reasoning dataset $\mathcal{R} = \{(v_i, q_i, a_i, r_i)\}_{i=1}^M$. However, as mentioned in Sec. 1, rationale r_i may be contaminated by visual hallucinations. Motivated by the observation illustrated in Fig. 1, VC-STaR aims to refine rationale r_i into a more faithful one \hat{r}_i by contrasting the (v_i, q_i, a_i) with a contrastive VQA counterpart sample $(\hat{v}_i, \hat{q}_i, \hat{a}_i)$ where q_i is synonymous with \hat{q}_i and v_i shares similar context with \hat{v}_i . The contrastive VQA pairs $\mathcal{P} = \{((v_i, q_i, a_i), (\hat{v}_i, \hat{q}_i, \hat{a}_i))\}_{i=1}^K$ support the contrasting and rationale refining process. The contrastive VQA pairs are curated by searching $(\hat{v}_i, \hat{q}_i, \hat{a}_i)$ for (v_i, q_i, a_i) within diverse data groups in \mathcal{D} for different VQA tasks, ensuring the generalization of VC-STaR. The VC-STaR is designed to address two key challenges: (1) how to curate meaningful contrastive VQA pairs; (2) how to transfer the fine-grained discriminative ability from dual-image contrasting to refine the single-image reasoning. Sec. 3.1 elaborates on the pipeline for curating contrastive VQA pairs. Building upon this foundation, Sec. 3.2 introduces our *contrasting and rethinking* procedure which embeds the dual-image comparison into a new reasoning path, guided by an LLM, to produce a more faithful rationale. The refined rationales are then used to construct a new reasoning dataset $\tilde{\mathcal{R}} = \{(v_i, q_i, a_i, \hat{r}_i)\}_{i=1}^L$, which we name the **Visual Contrastive Reasoning** dataset (VisCoR-55K). The VLM θ is updated to a new version $\tilde{\theta}$ with improved reasoning capability by finetuning on VisCoR-55K.

3.1 CONTRASTIVE VQA PAIR CURATION

To ensure the generalization of VC-STaR, the contrastive VQA pair curation pipeline should be flexible enough across a wide spectrum of VQA tasks. For better contrasting, each contrastive VQA pair $((v_i, q_i, a_i), (\hat{v}_i, \hat{q}_i, \hat{a}_i))$ should possess three key properties: (1) q_i and \hat{q}_i are synonymous. This shared question acts as a semantic anchor, grounding the two images v_i and \hat{v}_i at the same point in the semantic space. The images thus represent different manifestations of this anchor, providing a solid basis for contrasting; (2) v_i and \hat{v}_i are visually similar. v_i and \hat{v}_i should not be trivially distinct but exhibit visual similarity, creating a confusing contrasting. This visual proximity compels VLMs to engage in fine-grained contrasting to discriminate subtle differences; (3) q_i is reasoning dependent.




Figure 4: Faithful rationale generation pipeline. A contrastive analysis can be obtained based on the curated contrastive VQA pair. Leveraging the property of VLMs illustrated in Fig. 1, the contrastive analysis is then used to trigger a rethinking procedure, which refines the naive rationale into a more faithful one. This pipeline is designed to generate rationales for supervised finetuning.

q_i should be reasoning-provoking rather than one that can be solved by a straightforward answer. To achieve these requirements, as illustrated in Fig. 3, we propose a three-stage curation pipeline:

Data Collection. We collect 21 VQA datasets spanning five categories: reasoning (Zhang et al., 2019; Lu et al., 2021b), graph/chart (Kembhavi et al., 2016; Mathew et al., 2022; Masry et al., 2022; Tang et al., 2023; Lu et al., 2023; Liu et al., 2024a), math (Lu et al., 2021a; Cao & Xiao, 2022; Gao et al., 2025), general (Zhu et al., 2016; Johnson et al., 2017; Acharya et al., 2019; Schwenk et al., 2022; Wang et al., 2023a; Chen et al., 2024b), and OCR (ICDAR, 2019; Kiela et al., 2020; Yuan et al., 2022; Zhang et al., 2024b). This broad collection enriches the diversity of our curated pairs, which ensures the generalization ability of the finetuned model.

Contrastive VQA Pair Hunting. In order to compute the similarity of VQA pairs, we first represent the question q_i and the image v_i by high-dimensional embeddings, denoted as e_i^q and e_i^v respectively. We use GTE (Li et al., 2023c) text embeddings to represent the questions. In terms of image embedding, existing models fall into two types, *i.e.*, vision-language contrastive learning approaches (Radford et al., 2021; Tschannen et al., 2025) while vision-only self-supervised learning methods (Zhang et al., 2023; Oquab et al., 2024). The former ones mainly capture at global semantic information, while the later ones are good at instance discrimination. Neither of them are generic enough to adapt to the diverse domains. To tackle the dilemma, we build a versatile visual embedding model based on ID-based visual metric learning (Ypsilantis et al., 2024; An et al., 2023). Hunting for a counterpart $(\hat{v}_i, \hat{q}_i, \hat{a}_i)$ is then performed dataset-by-dataset. A sample (v_j, q_j, a_j) is recalled as a valid counterpart if it satisfies: $\gamma(e_i^v, e_j^v) < \phi_v$ and $\gamma(e_i^q, e_j^q) < \phi_q$, where $\gamma(\cdot, \cdot)$ is the cosine distance, and ϕ_v and ϕ_q are pre-defined thresholds for visual and question similarity, respectively. Any sample that fails to meet both conditions is dropped.

Difficulty-Based Data Sampling. For the goal of developing visual reasoning capability, q_i should be a difficult question requires reasoning rather than a straightforward one. We define the levels of difficulty based on the performance of VLM θ : (1) *easy samples* are with the simple q_i which can be correctly answered by θ without any auxiliary help; (2) *median samples* are with q_i which makes θ initially fails but succeeds when contrasting with $(\hat{v}_i, \hat{q}_i, \hat{a}_i)$ based on provided hint a_i (the $C\&H$ setting introduced in Fig. 1); (3) *hard samples* are the ones with q_i that cannot be correctly addressed by θ even with the help of contrasting. We only keep median-difficult contrastive VQA pairs for the rationale generating.

3.2 CONTRASTING AND RETHINKING

Rationales in the reasoning dataset $\mathcal{R} = \{(v_i, q_i, a_i, r_i)\}_{i=1}^M$ generated by the VLM θ itself include visual hallucinations. To achieve the goal:

$$\mathcal{R} = \{(v_i, q_i, a_i, \textcolor{red}{r}_i)\}_{i=1}^M \rightarrow \tilde{\mathcal{R}} = \{(v_i, q_i, a_i, \textcolor{blue}{\tilde{r}}_i)\}_{i=1}^M, \quad (1)$$

270 where \tilde{r}_i is the rectified rationale, we use the contrastive VQA counterpart $(\hat{v}_i, \hat{q}_i, \hat{a}_i)$ to provoke a
 271 rethinking action to refine r_i into \tilde{r}_i . As illustrated in Fig. 4, this pipeline includes three steps:
 272

273 **Thinking step.** Following the design of Zelikman et al. (2022) to provide the VLM θ with ground-
 274 truth answer a_i as hints, we prompt the VLM θ to generate the coarse rationale r_i for the target VQA
 275 sample (v_i, q_i, a_i) as follows:

$$r_i = f(v_i, q_i, a_i | \theta, \delta^t), \quad (2)$$

276 where f is a inference process with a “thinking prompt” δ^t . Details of δ^t are in the Sec. A.2.
 277

278 **Contrasting step.** Asking the VLM θ to compare the target VQA sample (v_i, q_i, a_i) with its
 279 contrastive counterpart $(\hat{v}_i, \hat{q}_i, \hat{a}_i)$ results a contrastive analysis c_i which may provide more faithful
 280 visual information:

$$c_i = f(((v_i, q_i, a_i), (\hat{v}_i, \hat{q}_i, \hat{a}_i)) | \theta, \delta^c), \quad (3)$$

281 where the δ^c is the “contrasting prompt”. When a_i has the same meaning with \hat{a}_i , δ^c requires
 282 summarizing the common patterns of v_i and \hat{v}_i ; When a_i is different from \hat{a}_i , δ^c expects the analysis
 283 about the fine-grained differences between v_i and \hat{v}_i . Details of δ^c are in the Sec. A.2.
 284

285 **Rethinking step.** As demonstrated in Fig. 1, c_i is more trustworthy than r_i . Hence, we adopt a
 286 LLM ψ to transfer the information from c_i to a new reasoning path according to r_i :

$$\tilde{r}_i = f(r_i, c_i | \psi, \delta^r), \quad (4)$$

287 where δ^r is the “rethinking prompt” which asks the LLM ψ to rectify the visual hallucinations in r_i
 288 according to the visual information from c_i . δ^r requires LLM ψ responding like directly answer the
 289 question q_i , details are in the Sec. A.2.
 290

291 To ensure the quality of the $\tilde{\mathcal{R}}$, we finalize the visual reasoning dataset by employing a text-matching
 292 post-processing to filter out samples that contain incorrect reasoning patterns. The final visual rea-
 293 soning dataset contains 55K VQA samples with corresponding rationales, *a.k.a.*, the VisCoR-55K.
 294

295 4 EXPERIMENTS

296 Section 4.1 details our experimental setup, including the supervised finetuning process and the
 297 benchmarks used to evaluate the effectiveness of the VC-STaR. In Section 4.2, we present a com-
 298 prehensive performance comparison. As a self-improving method for visual reasoning, we benchmark
 299 VC-STaR against two primary groups: (1) other self-improving baselines adaptable to visual reason-
 300 ing, and (2) models trained on off-the-shelf visual reasoning datasets. Finally, Section 4.3 provides
 301 in-depth ablation studies on designs of our method, including the contrastive VQA pair construction,
 302 the generalization on other base models, the difficulty sampling strategy, and the effect of the types
 303 of contrastive VQA counterpart.
 304

305 4.1 SETUP

306 **Implementation Details.** Using the LLaMA-factory framework (Zheng et al., 2024), we finetune
 307 the model for 3 epochs via full-parameter supervised finetuning (SFT), with the vision tower’s pa-
 308 rameters frozen. The SFT utilizes a learning rate of $1e-5$, a batch size of 256. The inference process
 309 of the finetuned model does not require such a contrastive pipeline illustrated in Fig. 4, and it fol-
 310 lows the standard inference paradigm of VLMs. As for the curation of contrastive VQA pair, the
 311 question similarity threshold ϕ_q is set to 0.15 and the visual similarity threshold ϕ_v is set as 0.5 for
 312 datasets of general images. For the datasets including icon, geometry, chart or graph images, the
 313 visual similarity threshold ϕ_v is set as 0.3. The LLM ψ used in the rethinking step of our rationale
 314 generation pipeline is the open-sourced Qwen2.5-72B.
 315

316 **Evaluation Benchmarks.** We employ 6 benchmarks designed to assess its robustness against hal-
 317 lucination, mathematical reasoning, and general abilities. The MMVP (Tong et al., 2024) and Hal-
 318 lusion (Guan et al., 2024a) benchmarks focus on visual hallucination, and the MathVista (Lu et al.,
 319 2024b) and MathVision (Wang et al., 2024) benchmarks are about the mathematical reasoning. The
 320 MMStar (Chen et al., 2024c) is a highly curated benchmark, composed of purified samples from
 321 multiple benchmarks, *e.g.*, MMMU (Yue et al., 2024) and MMBench Liu et al. (2024b). The MME-
 322 RealWorld benchmark Zhang et al. (2025b) is a large-scale, human-annotated benchmark for diffi-
 323 cult, real-world tasks. Therefore, MMStar and MME-RealWorld are suitable to evaluate the general
 324 perceptual and cognitive abilities under varied scenarios.
 325

324
 325 Table 1: Performance comparison with self-improving baselines and the models trained on off-
 326 the-shelf visual reasoning datasets on hallucination, math, and general benchmarks. We adopt the
 327 Qwen2.5VL-7B as our base model, and report its reasoning performance as a baseline. MME-
 328 RW is short for MME-RealWorld Zhang et al. (2025b); R1-OV is short for R1-Onevision (Yang
 329 et al., 2025b). **Blue** (**red**) numbers in parentheses represent performance **gains** (**drops**) relative to the
 329 baseline. The best performance is in **boldface**, and the second best is underlined.

Method	Bench.	Hallucination		Math		General		Avg.
		MMVP	Hallusion	MathVista	MathVision	MMStar	MME-RW	
Base Model		70.0	53.1	68.4	24.0	61.8	55.9	55.5
VQA SFT		74.3 _(+4.3)	<u>54.2</u> _(+1.1)	65.4 _(-3.0)	19.4 _(-4.6)	59.7 _(-2.1)	56.8 _(+0.9)	55.0 _(-0.5)
<i>Self-Improving Approaches</i>								
STaR(2022)		73.0 _(+3.0)	55.9 _(+2.8)	66.9 _(-1.5)	19.8 _(-4.2)	58.9 _(-2.9)	58.1 _(+2.2)	55.4 _(-0.1)
Verifier(2024a)		73.7 _(+3.7)	53.2 _(+0.1)	67.0 _(-1.4)	20.3 _(-3.7)	58.2 _(-3.6)	56.7 _(+0.8)	54.9 _(-0.6)
Feedback(2024)		75.0 _(+5.0)	53.4 _(+0.3)	68.8 _(+0.4)	22.1 _(-1.9)	63.2 _(+1.4)	56.0 _(+0.1)	56.4 _(+0.9)
<i>Off-the-Shelf Visual Reasoning Datasets</i>								
Virgo(2025)		68.0 _(-2.0)	47.2 _(-5.9)	63.5 _(-4.9)	21.5 _(-2.5)	59.7 _(-2.1)	29.4 _(-26.5)	48.2 _(-7.3)
LLaVA-CoT(2025)		71.7 _(+1.7)	50.3 _(-2.8)	68.4 _(+0.0)	24.4 _(+0.4)	63.1 _(+1.3)	59.3 _(+3.4)	56.2 _(+0.7)
R1-OV(2025b)		68.0 _(-2.0)	55.8 _(+2.7)	68.2 _(-0.2)	25.4 _(+1.4)	53.2 _(-8.6)	46.3 _(-9.6)	52.8 _(-2.7)
LPT(2025)		74.0 _(+4.0)	53.4 _(+0.3)	69.2 _(+0.8)	24.2 _(+0.2)	64.3 _(+2.5)	56.1 _(+0.2)	56.9 _(+1.4)
VC-STaR(Ours)		75.7 _(+5.7)	56.3 _(+3.2)	69.7 _(+1.3)	25.3 _(+1.3)	62.4 _(+0.6)	59.3 _(+3.4)	58.1 _(+2.6)

4.2 MAIN RESULTS

352 **Comparison with the base model.** To evaluate the effectiveness of our approach, we employ
 353 Qwen2.5VL-7B as the base model and adopt the "think step by step" prompt to enable chain-of-
 354 thought reasoning. We compare our method against this baseline, with results summarized in Ta-
 355 ble 1. VC-STaR demonstrates consistent performance gains across diverse challenging benchmarks,
 356 achieving an average improvement of 2.4%. Notably, it yields substantial improvements of 5.7%
 357 and 3.2% on MMVP and the Hallusion Benchmark, respectively, validating its efficacy in mitigating
 358 hallucinations within the reasoning process. Our approach also shows its enhanced reasoning capa-
 359 bilities on mathematical benchmark, *i.e.*, MathVista and MathVision. Furthermore, the improvement
 360 on the MMStar and MME-RealWorld underscore the generalizability of the VC-STaR under varied
 361 challenging general-purpose scenarios.


362 For qualitative validation, Figure 5 provides visual comparisons that offer deeper insights. The visu-
 363 alizations reveal that our model excels at grounding its textual rationales in the corresponding visual
 364 evidence. This capability remains robust even when confronted with visually complex patterns,
 365 thereby effectively mitigating hallucinations.

366 **Comparison with self-improving baselines.** We reproduce three self-improving baselines and
 367 compare VC-STaR against them. Each baseline is applied to the Qwen2.5VL-7B base model and
 368 generates rationales on VisCoR-55K for finetuning, differing in their core improvement mechanism:
 369 (1) STaR (Zelikman et al., 2022): Leverages ground-truth answers to regenerate rationales for in-
 370 correct predictions. (2) Verifier (Lu et al., 2024a): Filters out visually hallucinated rationales via
 371 a self-verification step (Zhang et al., 2025a) to ensure visual grounding. (3) Feedback (Qu et al.,
 372 2024): Refines rationales based on self-generated feedback in a recursive manner. Table 1 reveals a
 373 critical trade-off: existing self-improving methods boost performance on hallucination benchmarks
 374 at the expense of math and general capabilities. Our approach mitigates this pattern and achieves
 375 robust, consistent performance gains.

376 **Comparison with off-the-shelf visual reasoning datasets.** We also evaluate VC-STaR against base
 377 model finetuned on four off-the-shelf visual reasoning datasets. These datasets represent diverse
 378 strategies for rationale generation. For instance, Virgo (Du et al., 2025) makes the VLM think slowly

378	Question		What is the meaning of the white sign?		How many forks are visible in the image?
379	Qwen2.5VL-7B	The white sign in the picture has a red circle with a line through it, which is a universal symbol indicating "no parking." This sign is typically used to prohibit vehicles from parking in a specific area for a certain period of time. The exact duration or conditions under which parking is prohibited would be specified by additional text on the sign, but the symbol itself clearly indicates that parking is not allowed.			There are two forks visible in the image, one on each plate with the cake slices.
380	Qwen2.5VL-7B "think step by step"	The white sign in the image has a red circle with a line through it, which is a universal symbol indicating "no" or "prohibited." Specifically, this sign is a "No Parking" sign. It means that parking is not allowed in the area where the sign is posted. This type of sign is typically used to prevent vehicles from blocking traffic or causing congestion in certain areas.			To determine how many forks are visible in the image, let's follow these steps: 1. Identify the objects: Look closely at the image and identify all the objects that could be forks. 2. Count the forks: - There is one fork placed on the napkin next to the slice of cake on the left plate. - There is another fork partially visible on the right side of the image, resting on the plate with the chocolate cake. 3. Verify the count: Double-check to ensure no additional forks are hidden or partially obscured. After following these steps, we can conclude that there are two forks visible in the image.
381	Qwen2.5VL-7B + VC-STaR	<think>The white sign in the image contains a symbol that indicates "No Right Turn." This is evident because the sign features an arrow pointing to the right, which is universally recognized as a prohibition against turning right at that location. The absence of any conflicting symbols or additional text confirms this interpretation. Furthermore, the consistency in traffic signage conventions reinforces the conclusion that the sign's purpose is to prohibit right turns.</think><answer>The white sign means "No Right Turn." It prohibits vehicles from making a right turn at that particular location.</answer>			<think>The white sign in the image contains a symbol that indicates "No Right Turn." This is evident because the sign features an arrow pointing to the right, which is universally recognized as a prohibition against turning right at that location. The absence of any conflicting symbols or additional text confirms this interpretation. Furthermore, the consistency in traffic signage conventions reinforces the conclusion that the sign's purpose is to prohibit right turns.</think><answer>The white sign means "No Right Turn." It prohibits vehicles from making a right turn at that particular location.</answer>

395 **Figure 5: Qualitative Comparison with base model.** The second row shows the directly response
396 from the base model, the third row shows the response when the base model is prompted to “think
397 stey by step”, the last row shows the model improved with our VC-STaR. We highlight the key visual
398 evidences with red boxes for clarity of visualization. More results are in Sec. A.3.



411 **Figure 6: Performance comparison with other contrastive VQA pair construction strategies.**
412 Rationales in all settings are generated from the proposed VC-STaR. The red dashed line represents
413 the base model (Qwen2.5VL-7B) performance.

416 with purely textual rationales. In contrast, LLaVA-CoT (Xu et al., 2025) leverages hand-crafted
417 templates filled by the powerful GPT-4o (OpenAI, 2024a). Other approaches first convert visual
418 information into text; R1-Onevision (Yang et al., 2025b) generates rationales from image captions
419 using the DeepSeek-R1 model (Guo et al., 2025), while Long Perceptual Thought (LPT) (Liao
420 et al., 2025) extends this by using dense captions (Onoe et al., 2024) and keywords like “wait”
421 to elicit more detailed outputs from a similar LLM. In our experiments, we directly finetune the
422 base model on each of these datasets. Based on the results shown in Table 1, we can draw the
423 following conclusion: (a) Enhancing visual reasoning with purely textual rationales from Virgo is
424 ineffective. This strongly indicates that visual modality matters. (b) The model trained on LLaVA-
425 CoT suffers limited improvement, which demonstrates that the hand-crafted template struggle to
426 generalize across diverse VQA tasks. (c) The models trained on datasets generated by DeepSeek-R1
427 based on captions achieves notable improvements. However, the performance gap between them and
428 ours highlights the clear advantage of our visually-native approach over relying on textual captions.

4.3 ANALYSIS

431 **Can contrastive VQA pairs constructed in other ways?** To answer this, we explore alternative
432 strategies for curating contrastive VQA pairs. The first strategy is editing-based, utilizing the HQ-

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

Table 2: Evaluation of the effect of VC-STaR on other base models. **Blue** numbers in parentheses represent performance **gains**.

Model	VC-STaR	Hallusion	MathVision	MMStar
Qwen2.5VL	✗	46.9	18.4	55.0
3B	✓	53.2 _(+6.3)	21.9 _(+3.5)	55.7 _(+0.7)
InternVL2.5	✗	48.2	21.3	61.1
8B	✓	55.4 _(+7.2)	23.4 _(+2.1)	62.5 _(+1.4)

Table 3: Effect of the easy samples adding to VisCoR-55K. **Red** numbers in parentheses represent performance **drops**.

N_{easy}	Hallusion	MathVision	MMStar
0k	56.3	25.3	62.4
+20k	52.2 _(-4.1)	23.3 _(-2.0)	61.3 _(-1.1)
+40k	55.7 _(-0.6)	21.9 _(-3.4)	59.5 _(-2.9)

Table 4: Analysis about the effect of positive and negative contrastive VQA counterparts on GQA benchmark. We adopt the Qwen2.5VL-7B as our base model, and report its reasoning performance as a baseline. QR: query for relationships; QA: query for attributes; QG: query for global information; QC: query for category; CA: comparing of attribute; CC: choosing the object of one certain category; CAt: choosing the object of one certain attribute. **Blue** (**red**) numbers in parentheses represent performance **gains** (**drops**) relative to the baseline.

Setting	Pos.	Neg.	QR	QA	QG	QC	CA	CC	CAt	Total
Base Model	-	-	43.8	51.4	31.5	44.4	30.2	60.6	44.4	45.4
VC-STaR	✓	✗	48.3 _(+4.5)	52.8 _(+1.4)	46.8 _(+15.3)	57.1 _(+12.7)	42.9 _(+12.7)	60.1 _(-0.5)	44.4 _(+0.0)	50.6 _(+5.2)
VC-STaR	✗	✓	51.6 _(+7.8)	56.8 _(+5.4)	33.9 _(+2.4)	59.2 _(+14.8)	46.0 _(+15.8)	70.8 _(+10.2)	66.7 _(+22.3)	53.7 _(+8.3)
VC-STaR	✓	✓	53.5 _(+9.7)	57.3 _(+5.9)	46.3 _(+14.8)	55.6 _(+11.2)	36.5 _(+6.3)	72.7 _(+12.1)	55.6 _(+11.2)	54.7 _(+9.3)

Edit dataset (Hui et al., 2025). By prompting an LLM to create questions from editing instructions, we generate pairs where an original and an edited image yield different answers. The second strategy is caption-based, leveraging a dense caption dataset, *i.e.* DOCCI (Onoe et al., 2024). For this, we instruct an LLM to parse dense captions of visually similar images and generate a question that hinges on their subtle differences. For both strategies, we generate rationales for these newly created contrastive pairs using our proposed VC-STaR and finetune the Qwen2.5VL-7B. The results, presented in Fig. 6, lead to several observations: (a) VC-STaR is broadly effective, but performance is data-dependent. This is attributable to the biased data distribution of HQ-Edit and DOCCI, highlighting a key limitation of their curation scope. (b) VisCoR-55K includes contrastive pairs from a broader range of reasoning tasks, resulting in a more balanced performance.

Does VC-STaR generalize to other base models? We conduct experiments on Qwen2.5VL-3B and InternVL2.5-8B (Chen et al., 2025). Following the same self-improving procedure, we use VC-STaR to generate visual reasoning datasets from our VisCoR contrastive pairs, specifically for the two base models. We then finetune the Qwen2.5VL-3B and InternVL2.5-8B via the LLaMA-factory and SWIFT (Zhao et al., 2025), respectively. The results, presented in Table 2, demonstrate the model-agnostic effectiveness of our approach. These consistent and significant gains confirm that VC-STaR is a versatile and broadly applicable strategy for enhancing the visual reasoning ability.

What is the effect of easy samples on visual reasoning? Starting with our VisCoR-55K datasets, we incrementally add easy samples of two batches with 20K each. As illustrated in Table 3, we observe that the inclusion of easy samples is harmful. Specifically, when the number of easy samples increases, performance decreases. Therefore, we do not use the easy samples to avoid the potential “overthink” for straightforward problems.

How the contrastive VQA pairs of different types contribute? A contrastive VQA pair can be categorized as “positive” if both samples yield the same answer, and “negative” if their answers differ. To investigate the respective contributions of these two types of counterparts to our method’s performance, we conducted a controlled experiment on the GQA dataset (Hudson & Manning, 2019). The structured nature of GQA allows for the reliable curation of both positive and negative pairs via simple text matching. We applied VC-STaR to three distinct training sets: one generated from only positive contrastive pairs, one from only negative pairs, and a combined set including both. The results, detailed in Table 4, reveal a clear and significant trend. While both types of pairs are beneficial, negative counterparts are substantially more effective than positive ones, and their combination

486 yields the optimal total gain, highlighting their complementary roles. We attribute the superior effi-
 487 cacy of negative counterparts to their ability to induce stronger semantic contrast. Accordingly, our
 488 approach incorporates both positive and negative pairs without restriction to achieve optimal gain.
 489

490 5 DISCUSSION

493 **Rethinking VC-STaR from a cognitive perspective.** Learning and reasoning are inherently com-
 494 parative and contrastive processes. Humans rarely learn concepts in isolation. Instead, our human-
 495 beings refine our understanding by comparing examples, identifying distinguishing features, and
 496 reasoning through analogies and differences. The prototype theory concludes this cognitive be-
 497 havior as that our human-beings identify new identities by comparing them with the prototype
 498 concept (Rosch, 1975). Besides, the structure-mapping theory says that analogical reasoning can
 499 recognize the relationships shared by two domains (Gentner, 1983). This mapping can be treated as
 500 a fine-grained contrasting process. In our work, the contrasting process provide an opportunity to
 501 learn visual concept from the prototype, and our rethinking strategy reinforce the structure-mapping
 502 by generating new reasoning paths via contrasting. Through this work, we hope to highlight the
 503 potential of porting such human-like cognitive behaviors to the domain of reasoning.

504 **Computational cost of the VC-STaR.** Our approach introduces a modest cost in data generation,
 505 while maintaining a fine-tuning cost comparable to existing self-improving frameworks like STaR.
 506 We train the 7B model using the LLaMAFactory framework once the VisCoR-55K dataset is gen-
 507 erated. This process takes approximately 4 hours on a single node with 8×A800 GPUs. This cost
 508 is nearly identical to the standard fine-tuning procedure in approach like STaR. The additional cost
 509 stems from our data generation pipeline. The one-time contrastive VQA pair mining step processes
 510 approximately 7 samples per second, accelerated by our implementation featuring multi-processing
 511 concurrency and batch-wise computation. The rethinking step generates reasoning paths at about 1
 512 sample per second, thanks to the advanced vLLM framework Kwon et al. (2023) for fast inference.
 513 We argue this is a reasonable and acceptable investment to create our reasoning dataset.

514 6 CONCLUSION

516 In this work, we demonstrate that visual hallucination, a critical bottleneck for VLM reasoning, can
 517 be effectively mitigated to achieve better visual reasoning ability through the lens of contrast. Our
 518 key finding is that VLMs can see better by contrast. This phenomenon motivates us to propose
 519 the VC-STaR. The VC-STaR refines hallucinatory reasoning paths through analysis over curated
 520 contrastive VQA pairs, which yields our high-quality VisCoR-55K. Finetuning on VisCoR-55K
 521 delivers a consistent performance gain across five benchmarks targeting hallucination, math, and
 522 general reasoning, significantly surpassing other self-improving baselines and models trained on
 523 state-of-the-art visual reasoning datasets.

524 ETHICS STATEMENT

526 We confirm that this research adheres to the ICLR Code of Ethics. The datasets utilized for this
 527 work are all publicly accessible and intended for academic research. No commercial or private data
 528 was used, ensuring that our study is free from associated privacy concerns. The authors take full
 529 responsibility for the ethical conduct of this research.

531 532 REPRODUCIBILITY STATEMENT

533 To ensure the reproducibility of our results and to foster future research, we commit to making our
 534 resources publicly available upon the acceptance of this paper. These resources will include: our
 535 newly curated VisCoR-55K dataset, and the final finetuned model weights. Crucially, our dataset
 536 release will not only contain the final rationales generated by VC-STaR but also the corresponding
 537 contrastive VQA counterparts used to produce them. We believe this unique data will be a valuable
 538 resource for the community. We hope that providing this data will enable researchers to explore the
 539 role of contrast in promoting visual reasoning and other cognitive abilities across a wider range of
 settings, such as reinforcement learning.

540 REFERENCES
541

542 Manoj Acharya, Kushal Kafle, and Christopher Kanan. Tallyqa: Answering complex counting
543 questions. *AAAI*, pp. 8076–8084, 2019.

544 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel
545 Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language
546 model for few-shot learning. *NeurIPS*, pp. 23716–23736, 2022.

547 Xiang An, Jiankang Deng, Kaicheng Yang, Jiawei Li, Ziyong Feng, Jia Guo, Jing Yang, and
548 Tongliang Liu. Unicom: Universal and compact representation learning for image retrieval. In
549 *ICLR*, 2023.

550 Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
551 Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan,
552 Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng,
553 Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025.
554 URL <https://arxiv.org/abs/2502.13923>.

555 Jie Cao and Jing Xiao. An augmented benchmark dataset for geometric question answering through
556 dual parallel text encoding. In *COLING*, pp. 1511–1520, 2022.

557 Chameleon. Chameleon: Mixed-modal early-fusion foundation models, 2025. URL <https://arxiv.org/abs/2405.09818>.

558 Hongruixuan Chen, Jian Song, Chengxi Han, Junshi Xia, and Naoto Yokoya. Changemamba: Re-
559 mote sensing change detection with spatiotemporal state space model. *TGRS*, pp. 1–20, 2024a.

560 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahu
561 Lin. Sharegpt4v: Improving large multi-modal models with better captions. In *ECCV*, pp. 370–
562 387, 2024b.

563 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi
564 Wang, Yu Qiao, Dahu Lin, and Feng Zhao. Are we on the right way for evaluating large vision-
565 language models? In *NeurIPS*, pp. 27056–27087, 2024c.

566 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi
567 Wang, Yu Qiao, Dahu Lin, et al. Are we on the right way for evaluating large vision-language
568 models? In *NeurIPS*, 2024d.

569 Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shen-
570 glong Ye, Hao Tian, Zhaoyang Liu, Lixin Gu, Xuehui Wang, Qingyun Li, Yimin Ren, Zixuan
571 Chen, Jiapeng Luo, Jiahao Wang, Tan Jiang, Bo Wang, Conghui He, Botian Shi, Xingcheng
572 Zhang, Han Lv, Yi Wang, Wenqi Shao, Pei Chu, Zhongying Tu, Tong He, Zhiyong Wu, Huipeng
573 Deng, Jiaye Ge, Kai Chen, Kaipeng Zhang, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu,
574 Tong Lu, Dahu Lin, Yu Qiao, Jifeng Dai, and Wenhui Wang. Expanding performance bound-
575 aries of open-source multimodal models with model, data, and test-time scaling, 2025. URL
576 <https://arxiv.org/abs/2412.05271>.

577 Lei Ding, Jing Zhang, Haitao Guo, Kai Zhang, Bing Liu, and Lorenzo Bruzzone. Joint spatio-
578 temporal modeling for semantic change detection in remote sensing images. *TGRS*, 62:1–14,
579 2024.

580 Yuhao Dong, Zuyan Liu, Hai-Long Sun, Jingkang Yang, Winston Hu, Yongming Rao, and Ziwei
581 Liu. Insight-v: Exploring long-chain visual reasoning with multimodal large language models. In
582 *CVPR*, pp. 9062–9072, 2025.

583 Yifan Du, Zikang Liu, Yifan Li, Wayne Xin Zhao, Yuqi Huo, Bingning Wang, Weipeng Chen, Zheng
584 Liu, Zhongyuan Wang, and Ji-Rong Wen. Virgo: A preliminary exploration on reproducing o1-
585 like mllm, 2025. URL <https://arxiv.org/abs/2501.01904>.

586 Lisa Dunlap, Yuhui Zhang, Xiaohan Wang, Ruiqi Zhong, Trevor Darrell, Jacob Steinhardt, Joseph E
587 Gonzalez, and Serena Yeung-Levy. Describing differences in image sets with natural language.
588 In *CVPR*, pp. 24199–24208, 2024.

594 Alessandro Favero, Luca Zancato, Matthew Trager, Siddharth Choudhary, Pramuditha Perera,
 595 Alessandro Achille, Ashwin Swaminathan, and Stefano Soatto. Multi-modal hallucination control
 596 by visual information grounding. In *CVPR*, pp. 14303–14312, 2024.

597 Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wanjun Zhong, Yufei Wang, Lanqing Hong,
 598 Jianhua Han, Hang Xu, Zhenguo Li, et al. G-llava: Solving geometric problem with multi-modal
 599 large language model. *ICLR*, 2025.

600 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and
 601 Graham Neubig. Pal: Program-aided language models. In *ICML*, pp. 10764–10799, 2023.

602 Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive science*, 7(2):
 603 155–170, 1983.

604 Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena
 605 Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar,
 606 Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a
 607 new approach to self-supervised learning. In *NeurIPS*, pp. 21271–21284, 2020.

608 Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang
 609 Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An
 610 advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-
 611 language models. In *CVPR*, pp. 14375–14385, 2024a.

612 Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang
 613 Chen, Furong Huang, Yaser Yacoob, et al. Hallusionbench: an advanced diagnostic suite for
 614 entangled language hallucination and visual illusion in large vision-language models. In *CVPR*,
 615 pp. 14375–14385, 2024b.

616 Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek
 617 Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, Wolfgang Macherey, Arnaud
 618 Doucet, Orhan Firat, and Nando de Freitas. Reinforced self-training (rest) for language modeling,
 619 2023. URL <https://arxiv.org/abs/2308.08998>.

620 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu
 621 Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhi-
 622 hong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng,
 623 Chengda Lu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Deli Chen, Dongjie
 624 Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li,
 625 H. Zhang, Hanwei Xu, Honghui Ding, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li,
 626 Jingchang Chen, Jingyang Yuan, Jinhao Tu, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang,
 627 Jin Chen, Kai Dong, Kai Hu, Kaichao You, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean
 628 Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan
 629 Zhang, Minghua Zhang, Minghui Tang, Mingxu Zhou, Meng Li, Miaojun Wang, Mingming
 630 Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi
 631 Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu,
 632 Shangyan Zhou, Shanhua Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou,
 633 Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang,
 634 Wangding Zeng, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao,
 635 Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin
 636 Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue
 637 Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou,
 638 Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao
 639 Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying
 640 He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou,
 641 Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You,
 642 Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunx-
 643 ian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean
 644 Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui
 645 Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng
 646 Xu, Zhongyu Zhang, and Zhen Zhang. DeepSeek-R1 incentivizes reasoning in LLMs through
 647 reinforcement learning. *Nature*, 645(8081):633–638, 2025.

648 Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhitong Hu. Reasoning
 649 with language model is planning with world model. In *EMNLP*, pp. 8154–8173, 2023.
 650

651 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for
 652 unsupervised visual representation learning. In *CVPR*, pp. 9729–9738, 2020.

653 Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning
 654 and compositional question answering. In *CVPR*, pp. 6700–6709, 2019.

655

656 Mude Hui, Siwei Yang, Bingchen Zhao, Yichun Shi, Heng Wang, Peng Wang, Cihang Xie, and
 657 Yuyin Zhou. HQ-edit: A high-quality dataset for instruction-based image editing. In *ICLR*, 2025.

658 ICDAR. Overview - icdar 2019 robust reading challenge on scanned receipts ocr and information
 659 extraction, 2019. URL <https://rrc.cvc.uab.es/?ch=13>.

660

661 Qirui Jiao, Dao yuan Chen, Yilun Huang, Bolin Ding, Yaliang Li, and Ying Shen. Img-diff: Con-
 662 trastive data synthesis for multimodal large language models. In *CVPR*, pp. 9296–9307, 2025.

663 Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and
 664 Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual
 665 reasoning. In *CVPR*, pp. 2901–2910, 2017.

666 Daniel Kahneman. *Thinking, fast and slow*. macmillan, 2011.

667

668 Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali
 669 Farhadi. A diagram is worth a dozen images. In *ECCV*, pp. 235–251, 2016.

670

671 Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish
 672 Sabharwal. Decomposed prompting: A modular approach for solving complex tasks. *ICLR*, 2023.

673 Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ring-
 674 shia, and Davide Testuggine. The hateful memes challenge: Detecting hate speech in multimodal
 675 memes. *NeurIPS*, pp. 2611–2624, 2020.

676

677 Hoeseong Kim, Jongseok Kim, Hyungseok Lee, Hyun a Park, and Gunhee Kim. Viewpoint-agnostic
 678 change captioning with cycle consistency. 2021 ieee. In *ICCV*, pp. 2075–2084, 2021.

679

680 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large
 681 language models are zero-shot reasoners. *NeurIPS*, pp. 22199–22213, 2022.

682

683 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph
 684 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 685 serving with pagedattention. In *SOSP*, pp. 611–626, 2023.

686

687 Baiqi Li, Zhiqiu Lin, Wenxuan Peng, Jean de Dieu Nyandwi, Daniel Jiang, Zixian Ma, Simran
 688 Khanuja, Ranjay Krishna, Graham Neubig, and Deva Ramanan. Naturalbench: Evaluating vision-
 689 language models on natural adversarial samples. In *NeurIPS*, pp. 17044–17068, 2024.

690

691 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: bootstrapping language-image
 692 pre-training with frozen image encoders and large language models. In *ICML*, 2023a.

693

694 Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. Making
 695 language models better reasoners with step-aware verifier. In *ACL*, pp. 5315–5333, 2023b.

696

697 Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards
 698 general text embeddings with multi-stage contrastive learning, 2023c. URL <https://arxiv.org/abs/2308.03281>.

699

700 Zhuoqun Li, Haiyang Yu, Xuanang Chen, Hongyu Lin, Yaojie Lu, Fei Huang, Xianpei Han, Yongbin
 701 Li, and Le Sun. DeepSolution: Boosting complex engineering solution design via tree-based
 702 exploration and bi-point thinking. In *ACL*, pp. 4380–4396, 2025.

703

704 Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. Mind
 705 the gap: Understanding the modality gap in multi-modal contrastive representation learning. In
 706 *NeurIPS*, pp. 17612–17625, 2022.

702 Yuan-Hong Liao, Sven Elflein, Liu He, Laura Leal-Taixé, Yejin Choi, Sanja Fidler, and David
 703 Acuna. Longperceptualthoughts: Distilling system-2 reasoning for system-1 perception. In
 704 *COLM*, 2025.

705 Wei Lin, Muhammad Jehanzeb Mirza, Sivan Doveh, Rogerio Feris, Raja Giryes, Sepp Hochreiter,
 706 and Leonid Karlinsky. Comparison visual instruction tuning. In *CVPR*, pp. 2973–2983, 2025.

708 Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Aligning large
 709 multi-modal model with robust instruction tuning. *ICLR*, 2024a.

710 Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*,
 711 pp. 34892–34916, 2023.

713 Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan,
 714 Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around
 715 player? In *ECCV*, pp. 216–233, 2024b.

716 Yuliang Liu, Junjie Lu, Zhaoling Chen, Chaofeng Qu, Jason Klein Liu, Chonghan Liu, Zefan Cai,
 717 Yunhui Xia, Li Zhao, Jiang Bian, et al. Adaptivestep: Automatically dividing reasoning step
 718 through model confidence. *ICML*, 2025.

720 Jianqiao Lu, Zhiyang Dou, Hongru Wang, Zeyu Cao, Jianbo Dai, Yunlong Feng, and Zhijiang Guo.
 721 Autopsv: Automated process-supervised verifier. *NeurIPS*, pp. 79935–79962, 2024a.

722 Pan Lu, Ran Gong, Shibiao Jiang, Liang Qiu, Siyuan Huang, Xiaodan Liang, and Song-Chun Zhu.
 723 Inter-gps: Interpretable geometry problem solving with formal language and symbolic reasoning.
 724 *ACL*, 2021a.

726 Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang,
 727 and Song-Chun Zhu. Iconqa: A new benchmark for abstract diagram understanding and visual
 728 language reasoning. *NeurIPS DB*, 2021b.

729 Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter
 730 Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured
 731 mathematical reasoning. *ICLR*, 2023.

733 Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-
 734 Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of
 735 foundation models in visual contexts. In *ICLR*, 2024b.

736 Ji Ma, Wei Suo, Peng Wang, and Yanning Zhang. C3 l: content correlated vision-language instruc-
 737 tion tuning data generation via contrastive learning. In *IJCAI*, 2024.

739 Ruotian Ma, Peisong Wang, Cheng Liu, Xingyan Liu, Jiaqi Chen, Bang Zhang, Xin Zhou, Nan Du,
 740 and Jia Li. S²R: Teaching LLMs to self-verify and self-correct via reinforcement learning. In
 741 *ACL*, pp. 22632–22654, 2025.

742 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri
 743 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement
 744 with self-feedback. *NeurIPS*, pp. 46534–46594, 2023.

745 Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A bench-
 746 mark for question answering about charts with visual and logical reasoning. *ACL findings*, pp.
 747 2263–2279, 2022.

749 Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar.
 750 Infographiccvqa. In *WACV*, pp. 1697–1706, 2022.

751 Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. Compositional chain-of-
 752 thought prompting for large multimodal models. In *CVPR*, pp. 14420–14431, 2024.

754 Yasumasa Onoe, Sunayana Rane, Zachary Berger, Yonatan Bitton, Jaemin Cho, Roopal Garg,
 755 Alexander Ku, Zarana Parekh, Jordi Pont-Tuset, Garrett Tanzer, Su Wang, and Jason Baldridge.
 DOCCI: Descriptions of Connected and Contrasting Images. In *ECCV*, 2024.

756 OpenAI. Gpt-4o system card, 2024a. URL <https://arxiv.org/abs/2410.21276>.
 757

758 OpenAI. Openai o1 system card, 2024b. URL <https://arxiv.org/abs/2412.16720>.
 759

760 Maxime Oquab, Timothée Darzet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov,
 761 Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaeldin El-Nouby, Mido Assran, Nicolas
 762 Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael
 763 Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut,
 764 Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without super-
 765 vision. *TMLR*, 2024. ISSN 2835-8856.
 766

766 Zhiyu Pan, Yinpeng Chen, Jiale Zhang, Hao Lu, Zhiguo Cao, and Weicai Zhong. Find beauty in the
 767 rare: Contrastive composition feature clustering for nontrivial cropping box regression. In *AAAI*,
 768 pp. 2011–2019, 2023.

769 Dong Huk Park, Trevor Darrell, and Anna Rohrbach. Robust change captioning. In *ICCV*, pp.
 770 4624–4633, 2019.
 771

772 Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching
 773 language model agents how to self-improve. *NeurIPS*, pp. 55249–55285, 2024.

774 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agar-
 775 wal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya
 776 Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pp.
 777 8748–8763, 2021.
 778

779 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 780 Finn. Direct preference optimization: Your language model is secretly a reward model. *NeurIPS*,
 781 pp. 53728–53741, 2023.

782 Eleanor Rosch. Cognitive representations of semantic categories. *Journal of experimental psychol-
 783 ogy: General*, 104(3):192, 1975.
 784

785 Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi.
 786 A-okvqa: A benchmark for visual question answering using world knowledge. In *ECCV*, pp.
 787 146–162, 2022.

788 Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hong-
 789 sheng Li. Visual cot: Advancing multi-modal language models with a comprehensive dataset and
 790 benchmark for chain-of-thought reasoning. *NeurIPS*, pp. 8612–8642, 2024.
 791

792 Benny Tang, Angie Boggust, and Arvind Satyanarayan. Vistext: A benchmark for semantically rich
 793 chart captioning. In *ACL*, pp. 7268–7298, 2023.

794

795 Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In Andrea Vedaldi,
 796 Horst Bischof, Thomas Brox, and Jan-Michael Frahm (eds.), *ECCV*, pp. 776–794, 2020.

797 Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide
 798 shut? exploring the visual shortcomings of multimodal llms. In *CVPR*, pp. 9568–9578, 2024.
 799

800 Luong Trung, Xinbo Zhang, Zhanming Jie, Peng Sun, Xiaoran Jin, and Hang Li. ReFT: Reasoning
 801 with reinforced fine-tuning. In *ACL*, pp. 7601–7614, 2024.

802

803 Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdul-
 804 mohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff,
 805 Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. Siglip 2: Multilingual vision-language
 806 encoders with improved semantic understanding, localization, and dense features, 2025. URL
 807 <https://arxiv.org/abs/2502.14786>.
 808

809 Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. To see is to
 believe: Prompting gpt-4v for better visual instruction tuning, 2023a. URL <https://arxiv.org/abs/2311.07574>.

810 Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and
 811 Hongsheng Li. Measuring multimodal mathematical reasoning with MATH-vision dataset. In
 812 *NeurIPS DB*, 2024.

813

814 Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through align-
 815 ment and uniformity on the hypersphere. In *ICML*, pp. 9929–9939, 2020.

816 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 817 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 818 In *ICLR*, 2023b.

819

820 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yo-
 821 gatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language
 822 models. *TMLR*, 2022a.

823

824 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 825 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*,
 826 pp. 24824–24837, 2022b.

827 Shengqiong Wu, Hao Fei, Liangming Pan, William Yang Wang, Shuicheng Yan, and Tat-Seng Chua.
 828 Combating multimodal lilm hallucination via bottom-up holistic reasoning. *AAAI*, pp. 8460–8468,
 829 2025.

830 Guowei Xu, Peng Jin, Ziang Wu, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. Llava-cot: Let
 831 vision language models reason step-by-step. In *ICCV*, 2025.

832

833 Ling Yang, Zhaochen Yu, Tianjun Zhang, Minkai Xu, Joseph E Gonzalez, Bin Cui, and Shuicheng
 834 Yan. Supercorrect: Advancing small lilm reasoning with thought template distillation and self-
 835 correction. *ICLR*, 2025a.

836

837 Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng
 838 Yin, Fengyun Rao, Minfeng Zhu, Bo Zhang, and Wei Chen. R1-onevision: Advancing general-
 839 ized multimodal reasoning through cross-modal formalization, 2025b. URL <https://arxiv.org/abs/2503.10615>.

840

841 Linli Yao, Weiyang Wang, and Qin Jin. Image difference captioning with pre-training and contrastive
 842 learning. In *AAAI*, pp. 3108–3116, 2022.

843

844 Nikolaos-Antonios Ypsilantis, Kaifeng Chen, André Araujo, and Ondřej Chum. Udon: Universal
 845 dynamic online distillation for generic image representations. *NeurIPS*, pp. 86836–86859, 2024.

846

847 Ye Yuan, Xiao Liu, Wondimu Dikubab, Hui Liu, Zhilong Ji, Zhongqin Wu, and Xiang Bai. Syntax-
 848 aware network for handwritten mathematical expression recognition. In *CVPR*, pp. 4553–4562,
 849 2022.

850

851 Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens,
 852 Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun,
 853 Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and
 854 Wenhui Chen. Mmmu: A massive multi-discipline multimodal understanding and reasoning
 855 benchmark for expert agi. In *CVPR*, 2024.

856

857 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. Star: Bootstrapping reasoning with
 858 reasoning. In *NeurIPS*, pp. 15476–15488, 2022.

859

860 Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. Raven: A dataset for relational
 861 and analogical visual reasoning. In *CVPR*, pp. 5317–5327, 2019.

862

863 Dan Zhang, Sining Zhoubian, Ziniu Hu, Yisong Yue, Yuxiao Dong, and Jie Tang. Rest-mcts*: Llm
 864 self-training via process reward guided tree search. In *NeurIPS*, pp. 64735–64772, 2024a.

865

866 Fuxiang Zhang, Jiacheng Xu, Chaojie Wang, Ce Cui, Yang Liu, and Bo An. Incentivizing llms to
 867 self-verify their answers, 2025a. URL <https://arxiv.org/abs/2506.01369>.

864 Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung
 865 Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *ICLR*,
 866 2023.

867 Yanzhe Zhang, Ruiyi Zhang, Jiuxiang Gu, Yufan Zhou, Nedim Lipka, Dyi Yang, and Tong Sun.
 868 Llavar: Enhanced visual instruction tuning for text-rich image understanding, 2024b. URL
 869 <https://arxiv.org/abs/2306.17107>.

870 Yi-Fan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Junfei Wu, Feng
 871 Li, Kun Wang, Qingsong Wen, Zhang Zhang, et al. Mme-realworld: Could your multimodal llm
 872 challenge high-resolution real-world scenarios that are difficult for humans? In *ICLR*, 2025b.

873 Zhiusheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. Multimodal
 874 chain-of-thought reasoning in language models. *TMLR*, 2024c.

875 Y. Zhao, J. Huang, J. Hu, X. Wang, Y. Mao, D. Zhang, et al. Swift: a scalable lightweight infras-
 876 tructure for fine-tuning. In *AAAI*, pp. 29733–29735, 2025.


877 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, and Zheyan Luo. LlamaFactory: Unified
 878 efficient fine-tuning of 100+ language models. In *ACL*, pp. 400–410, 2024.

879 Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answer-
 880 ing in images. In *CVPR*, pp. 4995–5004, 2016.

881

A APPENDIX

A.1 DETAILS ABOUT VisCoR-55K



900 Figure 7: **Statistics of the contrastive VQA pair curation.** The bar chart (left y-axis) displays
 901 the total number of contrastive VQA pairs found in each dataset, with colors indicating the data
 902 category. The line graph (right y-axis) plots the ratio of hard samples identified within those
 903 pairs for each dataset. In the upper right, the pie charts provide a categorical breakdown of the sample
 904 distribution.
 905

906 The construction of our VisCoR-55K dataset is a multi-stage process involving efficient pair
 907 curation, difficulty-based filtering, and quality-controlled rationale generation. The entire pipeline
 908 is designed to produce a high-quality, challenging visual reasoning dataset. Our curation process for
 909 contrastive VQA pairs begins with a dataset-by-dataset, divide-and-conquer strategy. To maintain
 910 computational tractability and avoid a costly $O(n^2)$ search complexity across the entire data pool,
 911 we implement a greedy, first-match-exit search algorithm: for each sample within a given source
 912 dataset, the search for a contrastive VQA counterpart terminates as soon as the first valid match is
 913 identified. This efficient approach allows us to scale the curation process effectively. Following this
 914 procedure, we initially curated a large pool of 240k raw contrastive VQA pairs. The distribution is
 915 visualized by the bar chart in Fig. 7, with the left y-axis indicating the number of samples.
 916

This initial pool of 240k pairs then undergoes a rigorous filtering and refinement pipeline. First, we apply the difficulty-based sampling strategy (as detailed in Sec. 3.1) to select only the hard samples, which are most effective for enhancing the model’s reasoning capabilities. The proportion of hard samples varies significantly across datasets, and is illustrated by the line graph in Fig. 7 (plotted against the right y-axis). This critical filtering step narrows our collection down to 86k challenging contrastive pairs. Subsequently, we leverage the contrasting and rethinking pipeline to generate a high-quality rationale for each of these 86k samples. As a final quality control measure, we employ a text-matching-based post-processing step to automatically filter out any rationales containing unexpected or erroneous reasoning patterns. This process culminates in our final VisCoR-55K dataset, a collection of high-fidelity visual reasoning samples ready for finetuning. The pie charts in Fig. 7 provide a categorical overview of the data composition throughout this pipeline.

A.2 PROMPTS FOR THINKING, CONTRASTING, AND RETHINKING

As introduced in Sec. 3.2, 3 steps lead to the final rationales. We design 3 prompts for the thinking, contrasting, and rethinking steps. The thinking prompt is:

You are a helpful assistant to answer the question by thinking step by step.
INPUT
- Image: The image that serves as the basis for answering the question.
- Question: The question pertains to the content of the image.
- Answer: The correct answer for the question about the image.
INSTRUCTION
- You should analyze the question and decide to focus on which visual content.
- You should parse the details of visual content based on the question.
- You should conclude the visual evidence to answer the question.
OUTPUT
- The returned content MUST be in the natural flow.
<Image><Question><Answer>

The contrasting prompt is:

You are a helpful assistant to think step by step for discriminating between two images to answer two synonymous questions.
INPUT
- First Image: One image that serves as the basis for answering the question.
- Second Image: The other image that serves as the basis for answering the question.
- First Question: The question pertains to the content of the First Image.
- Second Question: The question pertains to the content of the Second Image.
- Answer: The correct answer for the question about the images.
INSTRUCTION
- When the correct answers for the two images are the same, you should summarize the common patterns in the visual content of the two images.
- When the correct answers for the two images are different, you should identify the differences in visual content between two images.
- Conclude the visual evidence to answer the questions respectively.
OUTPUT
- Return in the natural flow.
<FirstImage><SecondImage>
<FirstQuestion><SecondQuestion><Answer>

The “Answer” here is the concatenation of both samples. The rethinking prompt is:

972
 973 You are a helpful assistant to rewrite the coarse rationale into a more correct and more
 974 logical one based on a contrastive analysis.
 975 **### INPUT ###**
 976 - Question: The question to be answered based one given target image.
 977 - Answer: The correct answer to answer the question.
 978 - Coarse Rationale: The naive reasoning process answering the question.
 979 - Contrastive Analysis: The reasoning process when comparing the first image with the
 980 second image for synonymous questions.
 981 **### INSTRUCTION ###**
 982 - The contrastive analysis is more reliable than the coarse rationale.
 983 - If the answers in the contrastive analysis are the same for the two images, the model
 984 should formulate a summary reasoning schema. This schema must summarize the key visual
 985 features and confirm that the provided visual evidence aligns with this schema to derive the
 986 conclusion.
 987 - If the answers in the contrastive analysis are different for the two images, you can employ
 988 backward chaining hypothesizing the visual cues that would be present if the alternative
 989 answer were correct, and then highlighting the critical distinctions between this hypothetical
 990 scenario and the actual visual evidence.
 991 **### OUTPUT ###**
 992 - The output MUST be in the format of '<think>the thinking content</think><answer>the answering content</answer>'.
 993 - The content of thinking content MUST be between the special token of '<think>' and
 994 '</think>'
 995 - The content of answering content MUST be between the special token of '<answer>'
 996 and '</answer>'.
 997 <Question><Answer><CoarseRationale><ContrastiveAnalysis>

998 A.3 ADDITIONAL QUALITATIVE RESULTS

1000 Examples of rationales generated by VC-STaR in VisCoR-55K are illustrated in Fig. 8. After ob-
 1001 taining the model finetuned with VisCoR-55K on Qwen2.5VL-7B, we test it on some customized
 1002 visual question answering cases and observe the interesting results shown in Fig. 9.
 1003

1004 A.4 STATEMENT ON THE USE OF LARGE MODELS

1005 During the preparation of this paper, we utilized the large language model to assist with grammar
 1006 checking and language polishing. The outputs of the model were used solely as a reference to
 1007 improve the clarity and readability of our writing. All final content presented in this paper was
 1008 written and verified by the human authors, who take full responsibility for its accuracy and integrity.
 1009 The core intellectual contributions of this work, *e.g.*, initial motivation, experimental design, and
 1010 execution of the experiments, were developed exclusively through discussions among the authors,
 1011 with no involvement from any large language model.
 1012

1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025


1026	Question	Rationales
1027		<think>To determine the figure that logically completes the sequence, let's carefully analyze the target image: 1. **Shape Analysis**: - The shapes in the sequence are pentagons, and they remain consistent in form across all positions. 2. **Size Progression**: - Observing the figures, there is a clear trend of increasing size as the sequence progresses from left to right and top to bottom. Each subsequent shape grows larger than the previous one. 3. **Color/Pattern Observation**: - There is no apparent change in color or shading patterns influencing the placement of the figures. The consistency in color suggests that size is the primary factor driving the sequence. 4. **Conclusion Based on Size Increase**: - Given the established pattern of gradual size expansion, the next figure in the sequence must continue this trend by being larger than the preceding shapes while maintaining the same form. The details in the image confirm that the progression relies solely on incremental growth in size without introducing additional complexities such as color changes or new shapes. Therefore, the figure that aligns perfectly with this principle is option D.</think><answer>D</answer>
1028	Choose the figure that would logically complete the sequence.	
1029		<think>The image shows two red London United buses parked side by side, with their front ends visible. Therefore, based on this observation: There are **two** buses shown in the picture.</think><answer>2.</answer>
1030	How many buses are there?	
1031		<think>The image shows an animal positioned in such a way that only one horn is clearly visible, as the rest of its head is partially obscured by an archway. This indicates that while there may be another horn present, it is not visible from this angle or due to the obstruction. The visual evidence supports the conclusion that only one horn is identifiable in the given view.</think><answer>1</answer>
1032	How many horns are visible?	
1033		<think>To determine the measure of angle A, we start by analyzing the given information and leveraging geometric properties. Triangle ABC is inscribed in circle O, with D as the midpoint of side BC. Since D is the midpoint, OD is perpendicular to BC, and OD extended intersects the circle at point E. Connecting EC, we are given that angle OEC is 65°. From the property of isosceles triangles, since OE and OC are radii of the circle, triangle OEC is isosceles, and thus angles OEC and OCE are equal. Therefore, angle EOC can be calculated as: $\text{Angle EOC} = 180^\circ - 2 \times 65^\circ = 50^\circ$. Next, because D is the midpoint of BC and OD is perpendicular to BC, it follows that angle EOB is also 50°. Consequently, the central angle BOC, which subtends arc BC, is twice angle EOB: $\text{Angle BOC} = 2 \times 50^\circ = 100^\circ$. Finally, angle A, being an inscribed angle subtending the same arc BC, is half of the central angle BOC: $\text{Angle A} = \frac{1}{2} \times 100^\circ = 50^\circ$. This reasoning aligns with the specific details provided in the problem and confirms the measure of angle A.</think><answer>A: 50°</answer>
1034	In the given diagram, triangle ABC is inscribed within circle O, and D is the midpoint of side BC. If OD is extended to intersect circle O at point E, and EC is then connected, what is the measure of angle A if angle OEC is 65°? Choices: A: 50° B: 55° C: 60° D: 65°	

Figure 8: Examples of rationales in VisCoR-55K.

1035	Question	Qwen2.5VL-7B	Qwen2.5VL-7B "think step by step"	Qwen2.5VL-7B+VC-STaR
1036		The image shows a group of people in what appears to be a sports stadium, reacting with shock and surprise. The person in the center is holding a baseball bat , which could suggest that something unexpected or dramatic has occurred on the field, such as a foul ball hitting someone or a player making an extraordinary play. The expressions and body language of the spectators indicate a sudden and intense reaction to an event that has just taken place. Without additional context, it's not possible to determine the exact cause of their shock , but it is clear that whatever happened was surprising and possibly dangerous.	To determine what makes the people in the image so shocked, let's analyze the context and details: 1. Setting : The image appears to be taken at a baseball game, as indicated by the presence of a baseball bat and the crowd's attire, which includes team colors and logos. 2. Action : One person is holding a baseball bat , which suggests that something unexpected or dramatic has occurred on the field. This could be a play that resulted in a collision or an unusual hit, such as a home run, a foul ball, or an unusual play. 3. Reaction : The people in the image are displaying strong reactions—some are covering their faces with their hands, while others are looking toward the field with expressions of shock or surprise. This indicates that whatever happened was unexpected and likely had a significant impact on the game. 4. Conclusion : The fact that the bat is being held up suggests that it might have been used in a way that surprised the spectators, possibly due to a foul ball or a hit that went into the stands, leading to a potential injury or disruption. Given these observations, the people are likely shocked because of a surprising event on the field, such as a foul ball or something that started a controversial play or an unexpected turn of events during the game. The exact cause would depend on the specific circumstances of the game, but the general reaction points to a moment of high drama or surprise.	<think>The image depicts a scene at what appears to be a baseball game, with spectators showing expressions of surprise or shock. The key detail is the presence of a bat in mid-air, which suggests an unexpected event such as a foul ball or a player's swing. This aligns with the idea that something sudden and surprising has occurred during the game, causing the audience's reaction. While other scenarios might involve different causes for shock—such as accidents or unrelated incidents—the visual evidence here points directly to the bat as the source of their astonishment.</think><answer>The people are shocked because a bat has been thrown into the crowd, likely during a moment of excitement or surprise at the baseball game.</answer>
1037	What makes the people so shocked?			
1038		The car parked in the parking spot labeled "33" is blue.	The car on the parking spot labeled "33" is blue.	<think>The question asks about the color of the car in parking spot 33. The analysis shows that there is no car occupying parking spot 33 in the image, as it appears empty. Since there is no vehicle present in this specific spot, it is impossible to determine any color associated with it. This conclusion aligns with the observation that the absence of a car directly answers the question.</think><answer>There is no car parked in parking spot 33, so I cannot determine its color.</answer>
1039	What color is the car on the parking spot 33?			
1040		The girl with a single braid is working on the problem "11 - 2". The answer to this problem is 9.	To determine what answer the girl with a single braid will write on the board, we need to solve the mathematical expression she is working on: The expression on the board for the girl with a single braid is: 7×2 . Let's break it down step by step: 1. Identify the operation: the operation here is multiplication. 2. Multiply the numbers: $7 \times 2 = 14$. So, the girl with a single braid will write the number 14 on the board.	<think>To answer the question "What answer will the girl with a single braid write on the board?", let's break down the analysis step by step: 1. Identify the Girl with Single Braid : The girl with a single braid is located on the far left side of the image. 2. Position on the Board : The girl is positioned directly below the equation $7 \times 2 =$. 3. Action with Chalk : The girl is holding a piece of chalk, indicating that she is ready to write. 4. Conclusion from Visual Evidence : The equation $7 \times 2 = 14$. The girl is positioned directly below this equation, suggesting she will write the result. Based on these observations, it is clear that the girl with a single braid will write the number 14 on the board as the answer to the problem 7×2 . Therefore, the answer is: 14 .</think><answer>The girl with a single braid will write the number 14 on the board.</answer>
1041	What answer will the girl with a single braid write on the board?			

Figure 9: Additional qualitative comparison.