

ViSRL: Intention-Driven Visual Perception via Reinforced Reasoning

Zhangquan Chen^{1*} Xufang Luo^{2†} Dongsheng Li²
¹Tsinghua University, China ²Microsoft Research Asia, China

Abstract

Visual understanding is inherently intention-driven—humans selectively focus on different regions of a scene based on their goals. Recent advances in large multimodal models (LMMs) enable flexible expression of such intentions through natural language, allowing queries to guide visual reasoning processes. Frameworks like Visual Chain-of-Thought have demonstrated the benefit of incorporating explicit reasoning steps, where the model predicts a focus region before answering a query. However, existing approaches rely heavily on supervised training with annotated intermediate bounding boxes, which severely limits scalability due to the combinatorial explosion of intention-region pairs. To overcome this limitation, we propose ViSRL, the first framework that applies reinforcement learning (RL) to the problem of intention-driven visual perception. ViSRL optimizes the entire visual reasoning process using only reward signals. By treating intermediate focus selection as an internal decision optimized through trial-and-error, our method eliminates the need for costly region annotations while aligning more closely with how humans learn to perceive the world. Extensive experiments across multiple benchmarks show that ViSRL consistently outperforms strong baselines, demonstrating both its effectiveness and its strong generalization across different LMMs. Our code is available at <https://github.com/zhangquanchen/ViSRL>.

1. Introduction

Visual understanding is a fundamental problem in computer vision, enabling machines to interpret and interact with their surroundings [22, 41, 59]. Traditional methods of visual perception often process entire scenes uniformly [9, 11, 25, 47, 69], without considering the intent behind a given task. However, human perception is inherently intention-driven — people focus on different aspects of a scene depending on their goals. For example, when entering a room, a person searching for a television remote will

*The work was conducted during the internship of Zhangquan Chen (czq23@mails.tsinghua.edu.cn) at Microsoft Research Asia.

†Corresponding author (xuflo@microsoft.com)

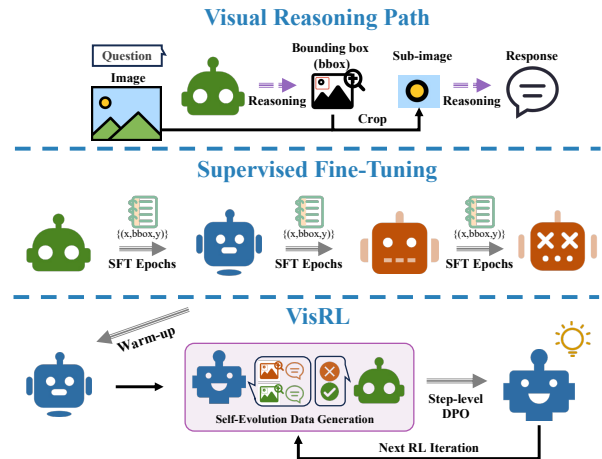


Figure 1. Illustration of using RL to optimize the visual reasoning process. SFT trains with densely annotated training data for several epochs. ViSRL leverages self-generated data and self-provided rewards to iteratively update the model using step-level DPO. This RL process removes the need for bounding box annotations, enabling a more human-like, intention-driven visual perception.

scan tables and couches, while someone checking the time will look at the walls for a clock. This context-dependent approach to visual perception suggests that intelligent models should also adapt their focus based on the task at hand. This leads to the problem of intention-driven visual perception, where the goal is to dynamically determine the most relevant regions of an image based on a given query or task [51].

With the advent of large multimodal models, the intention in perception tasks can now be expressed in a highly flexible way, which is natural language. Common LMMs, such as LLaVA [34] and Qwen-VL [56], first encode visual signals into tokens, and then both visual and text tokens are jointly processed by the large language models to produce final outputs. Despite showing powerful ability in many tasks [12, 26, 32, 57, 68, 70], this kind of one-pass end-to-end LMMs still suffer from hallucinations [20] and do not explicitly address the intention-driven visual perception problem. Further extending this paradigm, instead of treating the entire process as a single black-box inference, recent works have proposed frameworks like Visual Chain-of-Thought (Visual CoT [48]). These methods introduce an explicit reasoning step where the model first predicts a bounding box

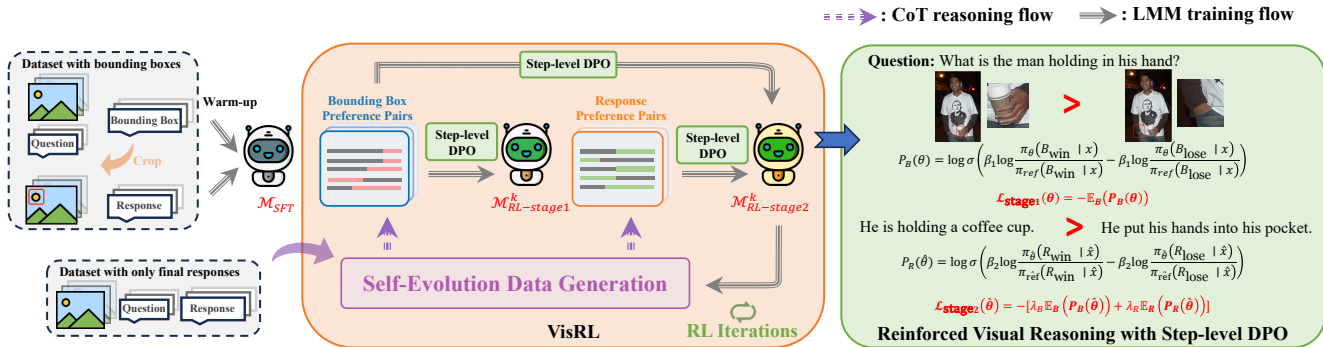


Figure 2. The schematic illustration of our VisRL framework. VisRL first utilizes a small amount of data for SFT warm-up, but in the subsequent RL training phase, it can leverage large-scale data without bounding box annotations. The RL phase of VisRL consists of iterative cycles of data generation and optimization, and k in the figure indicates the iteration index. The data generation process does not rely on external models or annotations; instead, it employs the model itself for data synthesis and scoring. The optimization step adopts step-level DPO to ensure the model learns each step of the reasoning process. In summary, VisRL enables intention-driven visual perception by leveraging RL to learn from task rewards without requiring annotations and external helps.

representing the critical region to focus on, crops the image to extract this region, and then feeds the cropped visual input back to the multimodal model. By conditioning the final answer on both the query and the selected focus area, this approach not only improves interpretability but also helps leverage the multi-turn in-context learning capabilities of the underlying LMM.

However, despite the advantages brought by methods like Visual CoT, these approaches also impose extremely high requirements on training data. Existing methods rely on supervised learning to teach models to produce intermediate reasoning steps. Specifically, for each intention or query, the training process requires corresponding bounding box annotations to guide the model in identifying the correct focus area. This dependence on exhaustive annotation is impractical, as the same image can correspond to vastly different regions depending on the request. As a result, the annotation complexity grows combinatorially with the diversity of queries, making it impossible to cover all potential cases in a scalable manner.

In this work, we propose a novel learning framework VisRL that optimizes the entire reasoning process based solely on rewards and feedback from the task itself, rather than dense annotations (shown in Figure 1). Specifically, we treat the success or failure of the task as a reward signal and apply reinforcement learning (RL) to train the model. Our approach requires no bounding box annotations, thereby addressing the scalability challenge posed by annotation requirements. Furthermore, this learning paradigm aligns more closely with how humans acquire perceptual skills — humans do not learn to focus on specific regions through meticulously annotated training data for each task, but rather through trial-and-error interaction with the environment, gradually developing the ability to adaptively zoom into relevant regions. By leveraging this reward-driven learning strategy, our framework enables intention-driven visual per-

ception in a more flexible, scalable, and human-like manner.

VisRL adopts an iterative DPO framework to complete the RL process for visual reasoning. This framework consists of multiple cycles of data generation and model optimization. During the data generation phase, we introduce a diversity controller to ensure that the generated bounding boxes cover a wide variety of potential focus regions. Additionally, we apply a filtering mechanism to select questions with appropriate difficulty levels for the current model and to identify the most effective preference pairs associated with each question. During the model optimization phase, we apply a step-level DPO algorithm, ensuring that the model learns to optimize every step of the visual reasoning process.

Our contributions can be summarized as follows.

- We present VisRL, the first framework to apply RL to the problem of intention-driven visual perception. By addressing the data annotation bottleneck, VisRL establishes a learning process that is much closer to human-like visual understanding.
- We design a tailored data generation pipeline for VisRL, incorporating both a diversity controller to enhance visual exploration and further use a step-level DPO algorithm to fully exploit the collected data.
- Extensive experiments across multiple benchmarks demonstrate that VisRL consistently outperforms strong baselines. Moreover, our results show that the effectiveness of VisRL generalizes well across different multimodal models, highlighting its broad applicability.

2. Related Work

Multi-modal Large Language Models With the advancement of large language models (LLMs) [1, 2, 4, 6, 7, 17, 19, 24, 42, 43, 52, 54, 55, 65, 71], multi-modal large language models which integrate vision and language modalities have also experienced rapid development. This progress

enables AI systems to better perceive and understand the real-world interplay between visual and textual information. Notable methods like LLaVA [34] aligns image tokens with pre-trained LLMs by training a projector, while other approaches utilize a Q-Former [29, 30] to learn image embeddings via learnable queries after extracting image features. These LMMs provide strong base models for VISRL, as they have the capability to process both visual and language data simultaneously, enabling the completion of the reasoning process.

Intention-Driven Visual Models Recently, several methods have attempted to enhance models’ intention-driven visual perception capabilities. VisCoT [48] employs a multi-turn interpretable processing mechanism with bounding boxes that dynamically focus on visual inputs. Similarly, SpatialCoT [37] achieves spatial grounding through spatial coordinates, while SegLLM [58] leverages mask-labeled data to enable reasoning about complex user segmentation intentions. V* (SEAL) [61] provides an LLM-guided search mechanism for efficient visual querying. Besides, both MLLM-TPO [64] and VisionLLM v2 [60] achieve intention-driven perception by training decoders for specific downstream tasks. Additionally, MVoT [28] introduces a text-image-text reasoning paradigm by training on interleaved data. These methods generally follow the supervised learning (SL) paradigm, therefore heavily relying on dense-labeled data (e.g., bounding boxes, spatial coordinates, masks, multi-round reasoning conversations), which constrains their ability to scale further. Instead, VISRL uses rewards as learning supervision and gets rid of the requirement on dense annotations. On the other hand, several methods employ tool-usage for enhancement. Specifically, AURORA [5] leverages specialized detection models, while Plug-and-Play [8] adopts a multi-agent framework. Besides, ViRReq [51] leverages the knowledge base to decompose visual recognition. However, these approaches rely on external models or knowledge, but not focusing on enhancing the intrinsic capabilities of the models themselves, while VISRL tries to completing the task via learning to reasoning by the model itself.

Multimodal Models with Chain-of-Thoughts Chain-of-Thought (CoT) reasoning plays an important role for LMMs. The methods can be broadly categorized into two types. (1) Text-thought methods [10, 15, 23, 53, 63] elicit textual CoT reasoning of LLMs in visual reasoning tasks by introducing text thinking tokens inspired by [21], guiding towards the final response. Differently, VISRL emphasizes that visual information should be involved in the reasoning process to fully leverage the strengths of multimodal models, rather than relying solely on language tokens for reasoning. (2) Multi-modal-thought methods involves multimodal information in the reasoning process. The Mind’s Eye [62] elicits

spatial reasoning of LLMs by visualizing their reasoning traces. Additionally, [37, 48, 58] first generate visual marks (e.g. bounding boxes, spatial coordinates, masks), and subsequently perform CoT reasoning based on these fine-grained visual marks. These works still uses SL for optimization while VISRL first explores RL in this direction. Besides, recent repositories [38, 49] also tried to enhance the bounding boxes (bboxes) generation through RL, but they uses ground truth bboxes to give rewards, which is still limited by dense annotation data.

3. Methodology

Preliminary. Reinforcement Learning [14] stands out as a highly effective method for significantly bolstering the robustness, factual accuracy, and safety of large language models [40]. The method consists of two key training stages, namely the reward model training and the policy model training. To avoid this complex training pipeline, [45] proposed Direct Preference Optimization (DPO). DPO streamlines the process by directly leveraging pair-wise preference data to optimize the policy model with an equivalent optimization objective. Specifically, given an input prompt x , and a preference data pair (y_{win}, y_{lose}) , DPO aims to maximize the probability of the preferred output y_{win} and minimize that of the undesirable output y_{lose} . The optimization objective is formulated as:

$$\mathcal{L}_{DPO}(\theta) = -\mathbb{E}_{(x, y_{win}, y_{lose}) \sim D}(P(\theta)), \quad (1)$$

where D is the pair-wise preference dataset, and P_θ is:

$$P(\theta) = \log \sigma \left(\beta \log \frac{\pi_\theta(y_{win} | x)}{\pi_{ref}(y_{win} | x)} - \beta \log \frac{\pi_\theta(y_{lose} | x)}{\pi_{ref}(y_{lose} | x)} \right). \quad (2)$$

σ is the sigmoid function, $\pi_\theta(\cdot | x)$ is the policy model to be optimized, $\pi_{ref}(\cdot | x)$ is the reference model kept unchanged during training, and the hyperparameter β serves to regulate the proximity of the policy model to the reference model.

Method overview. As shown in Figure 2, VISRL adopts DPO to optimize the entire visual reasoning process due to its simplicity. Our method leverages the final task success or failure as the outcome reward, and the grades of intermediate steps as the process reward, guiding the model to gradually refine its reasoning process through reinforcement learning. This reasoning process is divided into two steps. In the first step, the model generates a bounding box B representing the focused area based on the given query or question Q and the original image I . In the second step, the cropped region I^s corresponding to the bounding box, together with the original image and the question, is fed into the multimodal model to produce the final response R .

In Section 3.1, we will describe the data generation process required to support this two-stage reasoning pipeline.

In Section 3.2, we will introduce the optimization strategy, explaining how stepwise DPO is applied to enable the model to acquire intention-driven visual perception ability.

Notably, this data generation and optimization process will be iterated multiple times. The improved model can collect better data, and the better data will further refine the model. The initial model before RL training is denoted with \mathcal{M}_{RL}^0 , and the model which is updated with k iterations is denoted with \mathcal{M}_{RL}^k .

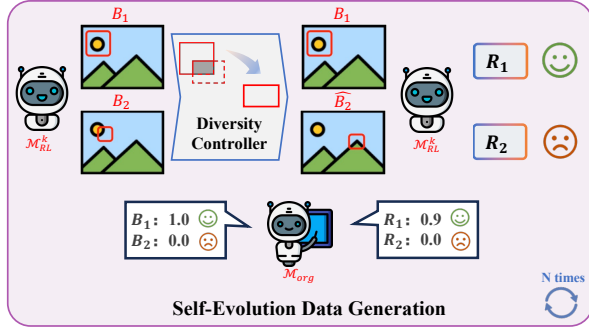


Figure 3. The schematic illustration of our data generation pipeline. Here \mathcal{M}_{RL}^k denotes the model updated with k iterations of data generation and optimization, and \mathcal{M}_{org} is the original model. VISRL uses \mathcal{M}_{RL}^k to generate samples, and use \mathcal{M}_{org} to provide rewards. Hence, different versions of a single model are used in this self-evolution data generation process, and no bounding box annotations and external models are introduced in this process.

3.1. Data Generation

Before RL training, we use SFT as a warm-up stage. This stage changes the original model \mathcal{M}_{org} into \mathcal{M}_{SFT} , so \mathcal{M}_{SFT} is \mathcal{M}_{RL}^0 . This stage requires bounding box annotations, so we directly use samples from VisCoT dataset [48]. Our intent here is just making the model have the basic ability to generate bounding boxes in specific format. Therefore, the used data amount is relatively small here (30k in ours and 438k in VisCoT). The second part of our source datasets consists of question-answer pairs from nine source datasets that span five distinct domains, with the majority of them being Visual Question Answering (VQA) and Image Captioning datasets. These are specifically used for constructing preference data for RL training, therefore requiring no bounding box annotations. We use an additional 180k data here. In Alg. 1, we explain how we construct preference data based on our source datasets in details.

Self-evolution. The critical parts for enabling effective RL learning is generating CoT data, and providing reward signals. Unlike previous manual annotation methods, which is infeasible for providing widely enough coverage for all possible cases, or approaches that rely on more powerful external models, which is more like distillation rather than allowing the model to learn on its own, we adopt a strategy that does not depend on external capabilities—instead, leveraging the model’s own self-evolution. To achieve this, we sample from

Algorithm 1: Self-Evolution Data Generation

Input: One sample from the source dataset including question Q , image I , and response R . The original model \mathcal{M}_{org} to provide criticism. The model updated with k RL iterations \mathcal{M}_{RL}^k to generate bounding boxes (b-boxes) and responses.

// Hyperparams in experiments: $N = 5$, $\mathcal{T} = 0.2$, $\mathcal{T}_{max}^b = \mathcal{T}_{max}^r = 0.8$, $\mathcal{T}_{min}^b = \mathcal{T}_{min}^r = 0.2$

- 1 Candidate preference pool $P \leftarrow \emptyset$.
- 2 **for** $i = 1$ to N **do**
- 3 $B_{i,1} = \mathcal{M}_{RL}^k(Q, I)$, $B_{i,2} = \mathcal{M}_{RL}^k(Q, I)$;
- 4 **if** $\text{IoU}(B_{i,1}, B_{i,2}) \geq \mathcal{T}$ **then**
- 5 $B_{i,2} = \mathcal{R}_S(\mathcal{U}_{\mathcal{I}}(B_{i,1}))$ // Refer to Eqn.(3).
 Increase b-boxes diversity.
- 6 // Crop sub-images and generate responses
 $I_{i,1}^s \leftarrow \text{crop}(I, B_{i,1})$, $I_{i,2}^s \leftarrow \text{crop}(I, B_{i,2})$;
- 7 $R_{i,1} = \mathcal{M}_{RL}^k(Q, I, I_{i,1}^s)$, $R_{i,2} = \mathcal{M}_{RL}^k(Q, I, I_{i,2}^s)$;
- 8 // Score b-boxes and responses
 $s_{i,1}^b = \mathcal{M}_{org}(Q, I_{i,1}^s, R)$, $s_{i,2}^b = \mathcal{M}_{org}(Q, I_{i,2}^s, R)$;
- 9 $s_{i,1}^r = \mathcal{M}_{org}(Q, R_{i,1}, R)$, $s_{i,2}^r = \mathcal{M}_{org}(Q, R_{i,2}, R)$;
- 10 $p_{i,1} = (Q, I, B_{i,1}, I_{i,1}^s, R_{i,1}, s_{i,1}^b, s_{i,1}^r)$,
 $p_{i,2} = (Q, I, B_{i,2}, I_{i,2}^s, R_{i,2}, s_{i,2}^b, s_{i,2}^r)$
- 11 $P \leftarrow P \cup \{p_{i,1}, p_{i,2}\}$;
- 12 $P_{win} = \{p_i \mid s_i^b \geq \mathcal{T}_{max}^b \text{ and } s_i^r \geq \mathcal{T}_{max}^r\}$
- 13 $P_{lose} = \{p_i \mid s_i^b < \mathcal{T}_{min}^b \text{ and } s_i^r < \mathcal{T}_{min}^r\}$
- 14 $i_b = \arg \max_{p_i \in P_{win}} s_i^b$, $i_r = \arg \max_{p_i \in P_{win}} s_i^r$
- 15 **if** $i_b = i_r$ **then**
- 16 $p_{win} \leftarrow p_{i_b/i_r}$
- 17 $i_b = \arg \min_{p_i \in P_{lose}} s_i^b$, $i_r = \arg \min_{p_i \in P_{lose}} s_i^r$
- 18 **if** $i_b = i_r$ **then**
- 19 $p_{lose} \leftarrow p_{i_b/i_r}$
- 20 **return** (p_{win}, p_{lose}) when they both have values // filter the sample that is too hard/easy for current model

\mathcal{M}_{SFT} multiple times to generate CoT data with as much diversity as possible, while use \mathcal{M}_{org} to provide criticism. This approach not only showcases the intrinsic capabilities of the model, but also facilitates a more effective adjustment of the predicted probability distribution toward a stable state through self-generated data.

Sample generation and rewarding. For each input question and image (Q, I) , we first sample from \mathcal{M}_{SFT} twice to obtain the bounding boxes B_1 and B_2 , respectively. Then, we introduce a diversity controller to ensure diversity of bounding boxes. Specifically, we update B_2 according to the Intersection over Union (IoU) and the rejecting threshold value \mathcal{T} between B_1 and B_2 , which is formulated as:

$$\hat{B}_2 = \begin{cases} B_2, & \text{IoU}(B_1, B_2) < \mathcal{T} \\ \mathcal{R}_S(\mathcal{U}_{\mathcal{I}}(B_1)), & \text{IoU}(B_1, B_2) \geq \mathcal{T} \end{cases}, \quad (3)$$

where $\mathcal{U}_{\mathcal{I}}(B_1)$ is a set of bounding boxes that are outside of B_1 but within image I . The operator $\mathcal{R}_S(\cdot)$ represents a random selection from the set $\mathcal{U}_{\mathcal{I}}(B_1)$, and has an area differing by no more than S compared to B_1 .

Based on the bounding boxes B_1 and \hat{B}_2 , we crop the sub-image I_1^s and I_2^s from I , respectively. Then, we input (Q, I, I_1^s) and (Q, I, I_2^s) to the model to obtain final

responses R_1 and R_2 , separately. At this point, we have completed the sampling of two distinct CoT reasoning paths.

Then, we should evaluate whether the paths are good or not. VISRL uses the original model before SFT stage \mathcal{M}_{org} as a critic to score the pairs B_1, B_2 and R_1, R_2 . Scores for the bounding boxes and the final responses are denoted with s_1^b, s_2^b and s_1^r, s_2^r , respectively. For more details including prompt designs, please refer to Sec. 3 in the Supp. Mat..

Data filtering. For each input question and image, repeating the above data generation process N times, we will obtain a set of $2N$ candidates $P = \{p_1, p_2, \dots, p_{2N}\}$, where each p contains original question, image, one bounding box, cropped sub-image, the final response and two scores, i.e., $p_i = (Q, I, B_i, I_i^s, R_i, s_i^b, s_i^r)$.

Moreover, to ensure the validity of preference data pairs in the candidate set P , we apply filtering by setting win and lose thresholds, \mathcal{T}_{max}^b and \mathcal{T}_{min}^b for bounding boxes, \mathcal{T}_{max}^r and \mathcal{T}_{min}^r for responses:

$$\begin{aligned} P_{win} &= \{p_i \mid s_i^b \geq \mathcal{T}_{max}^b \text{ and } s_i^r \geq \mathcal{T}_{max}^r\} \\ P_{lose} &= \{p_i \mid s_i^b < \mathcal{T}_{min}^b \text{ and } s_i^r < \mathcal{T}_{min}^r\}. \end{aligned} \quad (4)$$

For the current question and image (Q, I) , if the win set and the loss set is not empty, this question and image and its corresponding generated data are preserved. The intuition here is that we apply a filter to the questions in the dataset, selecting those with a difficulty level suitable for the current model. Questions that are too difficult will prevent the model from generating meaningful answers, while questions that are too easy will result in the model producing correct answers across the board. Both of these types of questions will be filtered out during the current training round. However, as the model’s capabilities are updated and strengthened, less data are filtered in the next iteration, as indicated by Data Num. of VisRL-Full vs. VisRL-Full-Iter1 in Tab. 6.

Then, for each preserved question, we select the most representative path from each set to obtain the win-lose preference pairs, denoted as $(p_{win}, p_{lose}) = (\mathcal{C}_{max}(P_{win}), \mathcal{C}_{min}(P_{lose}))$. In the case of \mathcal{C}_{max} , it is required that both s^b and s^r of p_i are the maximum over the other elements in the set P_{win} . Conversely, \mathcal{C}_{min} requires that both scores are the minimum in the set P_{lose} . If the condition of being simultaneously maximum/minimum on both s^b and s^r is not met, the data point will also be discarded. Finally, we obtain the preference dataset $D_P = \{(p_{win_1}, p_{lose_1}), (p_{win_2}, p_{lose_2}) \dots\}$

3.2. Reinforced Visual Reasoning

Similar with [27], VISRL uses a step-level DPO method. It is divided into two stages. Stage 1 involves optimizing the bounding box, while stage 2 focuses on the joint optimization of the bounding box and the final response. For stage 1, given

the input question-image $x = (Q, I)$, the objective is:

$$\mathcal{L}_{stage1}(\theta) = -\mathbb{E}_{(x, B_{win}, B_{lose}) \sim D_P} (P_B(\theta)), \quad (5)$$

where each pair-wise preference paths in D_P consists of bounding boxes B_{win} and B_{lose} , the formulation of bounding box preference probability is:

$$P_B(\theta) = \log \sigma \left(\beta_1 \log \frac{\pi_\theta(B_{win} | x)}{\pi_{ref}(B_{win} | x)} - \beta_1 \log \frac{\pi_\theta(B_{lose} | x)}{\pi_{ref}(B_{lose} | x)} \right). \quad (6)$$

After stage1, the policy model updated from π_θ to $\pi_{\hat{\theta}}$, while the reference model is updated to $\pi_{\hat{r}_{ef}}$. Then, for stage2, we further consider the cropped-image from B_{win} to make CoT inference, that is $\hat{x} = (Q, I, I_{win}^s)$, then the formulation of response preference probability is:

$$P_R(\hat{\theta}) = \log \sigma \left(\beta_2 \log \frac{\pi_{\hat{\theta}}(R_{win} | \hat{x})}{\pi_{\hat{r}_{ef}}(R_{win} | \hat{x})} - \beta_2 \log \frac{\pi_{\hat{\theta}}(R_{lose} | \hat{x})}{\pi_{\hat{r}_{ef}}(R_{lose} | \hat{x})} \right), \quad (7)$$

where R_{win} and R_{lose} are the final preference responses in D_P . Based on Eq. 6 and Eq. 7, the objective for jointly optimizing the bounding boxes and the responses in stage 2 can be formulated as:

$$\mathcal{L}_{stage2}(\hat{\theta}) = -(\lambda_B \mathcal{L}_B(\hat{\theta}) + \lambda_R \mathcal{L}_R(\hat{\theta})), \quad (8)$$

where:

$$\mathcal{L}_B(\hat{\theta}) = \mathbb{E}_{(x, B_{win}, B_{lose}) \sim D_P} (P_B(\hat{\theta})), \quad (9)$$

$$\mathcal{L}_R(\hat{\theta}) = \mathbb{E}_{(\hat{x}, R_{win}, R_{lose}) \sim D_P} (P_R(\hat{\theta})). \quad (10)$$

4. Experiments

4.1. Comparisons with Baselines

We evaluate our method with several state-of-the-art methods on an array of different categories as follows. More details are in Supp. Mat.. As shown in Tab. 1, we categorize methods into three groups based on the types of LLM and vision encoder, then evaluate them on comprehensive benchmarks (MME, MMBench) as well as hallucination benchmarks (POPE). In all cases, our method achieves either the best or second-best performance, demonstrating the robust and well-rounded improvement over the baseline. In contrast, other methods exhibit performance drops on specific benchmarks (e.g., SEAL on MMBench, VisCoT on MME). We attribute this phenomenon to the limitations inherent in their training approaches – data-driven SFT struggles to generalize effectively, while tool-usage methods suffer from intrinsic shortcomings on certain datasets. Notably, our approach achieves highly comprehensive and promising results while using only 30k dense-labeled (w. bounding boxes) samples – significantly fewer than other methods, and without relying on any external capabilities. In particular, under Vicuna-7B

Table 1. The evaluation of different baselines on MME [18], MMBench [36], and POPE [31] datasets. Datasets marked with [D] are dense-labeled datasets (e.g., CoT data). In different methods, [B] denotes the base model, [D] represents data-driven SFT methods, and [T] refers to tool-usage methods (e.g., agents). The **best** is highlighted and the second-best is underlined. Remark: the data number considered here includes only the data used to enhance specific model capabilities and pretraining data, excluding general instruction-tuning dataset.

Method	LLM	Vision Encoder	MME	MMBench	POPE	Dataset Num.
LLaVA [B] [34]	Vicuna-7B [13]	CLIP-ViT-L-14-224 [44]	1051.2	34.4	76.5	558K
SEAL [D] [61]	Vicuna-7B	CLIP-ViT-L-14-224	1128.9	33.1	82.4	558K + 387K [D]
LLaVA + P2G [T] [8]	Vicuna-7B	CLIP-ViT-L-14-224	<u>1223.0</u>	—	—	558K + 427K [D]
LLaVA + VisRL	Vicuna-7B	CLIP-ViT-L-14-224	1183.8	<u>37.5</u>	78.2	558K + 30K [D]+180K
LLaVA + VisRL- Iter1	Vicuna-7B	CLIP-ViT-L-14-224	1238.3	38.6	80.4	180K
LLaVA-1.5 [B] [33]	Vicuna-7B	CLIP-ViT-L-14-336	1510.7	64.3	85.8	558K
VisCoT [D] [48]	Vicuna-7B	CLIP-ViT-L-14-336	1453.6	67.9	86.0	558K + 376K [D]
LLaVA-1.5 + VisRL	Vicuna-7B	CLIP-ViT-L-14-336	<u>1526.3</u>	<u>70.1</u>	<u>87.5</u>	558K + 30K [D]+180K
LLaVA-1.5 + VisRL- Iter1	Vicuna-7B	CLIP-ViT-L-14-336	1560.0	71.7	88.8	180K
LLaVA-NeXT [B] [35]	Vicuna-7B-1.5 [72]	CLIP-ViT-L-14-336	1611.1	72.3	—	558K
VisionLLM v2 [D] [60]	Vicuna-7B-1.5	CLIP-ViT-L-14-336	1512.5	77.1	87.5	892K
Insight-V-LLaVA [T] [15]	Vicuna-7B-1.5	CLIP-ViT-L-14-336	1583.9	81.7	—	558K + 215K [D]
LLaVA-NeXT + VisRL	Vicuna-7B-1.5	CLIP-ViT-L-14-336	<u>1619.2</u>	78.8	<u>88.4</u>	558K + 30K [D]+180K
LLaVA-NeXT + VisRL- Iter1	Vicuna-7B-1.5	CLIP-ViT-L-14-336	1637.0	<u>80.0</u>	89.3	180K

Table 2. Performance on the different benchmarks. Base model is evaluated without CoT while others leverage CoT during inference. The **best** results from different LMMs are highlighted.

Training Phase	Base	SFT	SFT+RL1	SFT+RL1+RL2
LLaVA-1.5-7B [33]	0.444	0.538	0.568	0.605
LLaVA-NeXT-7B [35]	0.459	0.556	0.593	0.628
Llama-3.2-V-11B [39]	0.635	0.724	0.749	0.786
MiniCPM-o-2.6-8B [66]	0.428	0.519	0.544	0.591
PaliGemma2-10B [50]	0.377	0.479	0.507	0.562
Yi-VL-6B [67]	0.389	0.492	0.523	0.573
Qwen2.5-VL-7B [3]	0.640	0.730	0.763	0.796

and CLIP-ViT-L-14-336, our method outperforms VisCoT across all benchmarks – the representative data-driven SFT approach. Specifically, our method outperforms VisCoT by **5.00%** (1526.3 vs. 1453.6) on MME and by **1.74%** (87.5 vs. 86.0) on the hallucination benchmark POPE. Moreover, after 1 iteration (Base Model + VisRL- Iter1), VisRL improves performance by **1%** to **4%** across all benchmarks.

4.2. Results on Visual CoT Benchmark

In this section, we comprehensively investigate different training phases in VisRL across various base LMMs. We use Visual CoT benchmark here, which primarily focuses on scenarios where the LMM needs to concentrate on specific regions within a complete image.

Settings. For SFT, the objective is to regularize the model with the capability of outputting bounding box, while RL is to enhance the model’s visual perception capabilities via rewards. To train the LMM with outputting bounding box in SFT phase, we add a CoT prompt (“Please first provide the bounding box coordinate of the region, then refer to the corresponding sub-image to answer the question better.”) to the question, asking the model to identify the most informative region of the image. As shown in Tab. 3, we found that the model has already been capable of outputting bounding boxes under SFT with the dataset of 30k. Then, we sequentially proceed with further training using stage1 (RL1) and

stage2 (RL2) as described in Sec. 3.2.

Results. As shown in Tab. 2, there remains substantial room for improvement after SFT. Tab. 1 in Supp. Mat. details the results, indicates that there is a decline on some datasets even with the use of visual CoT (e.g. InfogVQA), which further corroborates the validity of our revised SFT strategy. Besides, we found that the difference in results obtained from SFT with 30k or 438k is not significant (i.e., VisCoT (0.574) vs. LLaVA-1.5-7B-Ours-SFT (0.538)). These suggests that: on the one hand, the improvement in model capability is more attributed to the introduction of CoT rather than the SFT memory. On the other hand, the model’s capability has already achieved saturated in a data-driven SFT manner and fails to further generalize, which is also the reason why Visual CoT fails in some OOD scenarios. However, our RL method can achieve comprehensive enhancement, with the promising improvement of up to **49.07%** (PaliGemma2-10B: 0.377 vs. 0.562), and the minimum improvement of **23.78%** (Llama-3.2-Vision-11B: 0.635 vs. 0.786). Meanwhile, we does not rely on a large amount of dense-labeled data (i.e. with bounding box annotation), but still has learned more essential visual perception.

Detection Ability. Our approach is grounded in CoT for visual reasoning, thereby placing significant emphasis on the accuracy of intermediate bounding boxes. To substantiate the enhancement in bounding box precision following the implementation of our RL method, we present the detection performance in Tab. 4. Specifically, we compute the IoU between the predicted CoT bounding boxes and the corresponding GT bounding boxes, deeming a prediction correct if the IoU value surpasses 0.5. It is evident that, when using the same base model – LLaVA-1.5-7B, the performance after our SFT with 30k data is somewhat inferior to that of Visual CoT which leverages 438k data. However, during the RL phase, we get rid of bounding box annotation data and exclusively utilized 180k simple question-answer

Table 3. Ratio of successful bounding box outputs of different SFT data number in terms of Qwen2.5-VL-7B. We evaluate on Visual CoT benchmark with 8281 data.

SFT Data Num.	10k	20k	30k	50k	100k
Ratio	28.14%	95.48%	99.87%	99.87%	99.87%

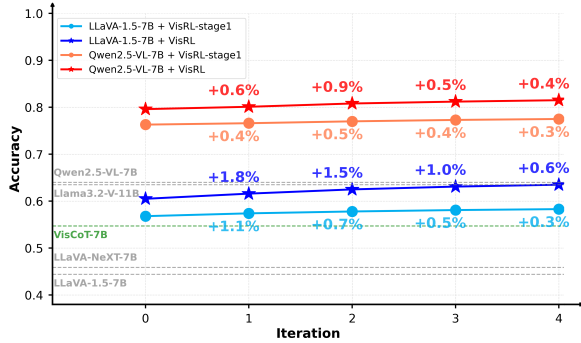


Figure 4. Performance of our VisRL over multiple iterations, attributing to the intertwined improvement of data quality and model capability during the iterative process. The accuracy is calculated as the average value over the 11 datasets listed in Tab. 2.

pairs (further processed via our data generation pipeline in Sec. 3.1 to construct preference data), thereby achieving an accuracy improvement of **15.61%** (0.437 vs. 0.378) over Visual CoT. Moreover, even with only RL1, we still attained an approximate **8.47%** (0.410 vs. 0.378) improvement. Notably, on the DocVQA dataset (which was not included in the RL training phase), RL1 achieved a remarkable **39.71%** (0.190 vs. 0.136) improvement, while RL2 accomplished an impressive **73.68%** (0.231 vs. 0.136) improvement. We attribute these substantial gains to the robust generalization capability of our RL method.

4.3. Performance across Multiple RL Iterations

Fig. 4 illustrates the performance variation curves of our method over multiple iterations on LLaVA-1.5-7B and Qwen2.5-VL-7B. It can be observed that after each iteration (i.e., data is regenerated and VisRL training is reconducted), performance improves significantly, regardless of whether only RL1 is applied or the full process of RL1+RL2 is used. The performance gains range from a minimum of **0.3%** to a maximum of **1.8%**. Since Qwen2.5-VL-7B is closer to the performance upper bound, its growth is relatively slower. Notably, after 4 iterations, LLaVA-1.5-7B even surpasses Llama-3.2-Vision-11B, as indicated by the blue curve. This further validates the promising potential of our method and lays the experimental foundation for future optimizations in online algorithms.

4.4. Visualization

This section presents the qualitative performance of our VisRL in Fig. 5, highlighting the accuracy of our method in identifying critical regions within images and then aid

in CoT reasoning. Compared to VisCoT (red), our VisRL (green) demonstrates greater performance in both localizing the regions of interest and generating the final response, as evidenced by the ground truth (GT) comparison (blue).

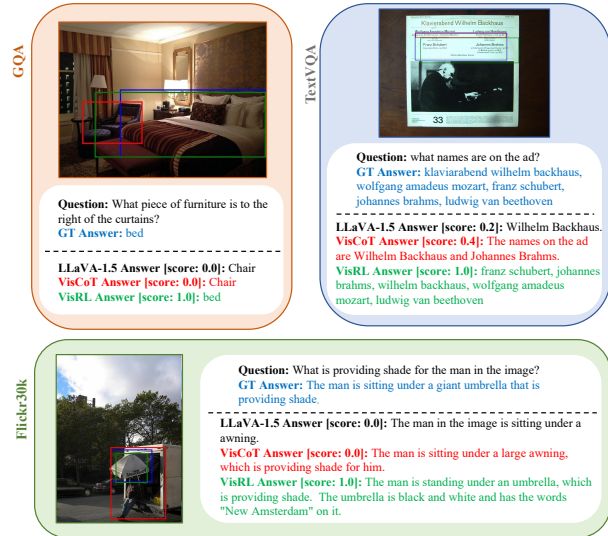


Figure 5. Visualization of LLaVA-1.5 vs. VisCoT vs. VisRL (based on LLaVA-1.5). GT bounding boxes are shown in blue, VisCoT-generated bounding boxes are shown in red, while Ours-generated bounding boxes are in green. The scores are evaluated by the GPT-4o. Our method consistently delivers the best results across various benchmarks. More visualizations are in the Supp. Mat..

4.5. Ablation Study

Different stages. In Tab. 2, we have conducted the ablation study on the usage of different training stages (SFT, SFT+RL1, SFT+RL1+RL2), demonstrating a consistent performance improvement from SFT to RL1 and then to RL2. Furthermore, in Tab. 5, we present the results of training with RL2-Only after SFT with bounding boxes. As observed, applying RL2-Only training already yields a promising improvement compared to SFT-then-CoT (0.759 vs. 0.730). However, it still results in significant performance drops on certain benchmarks compared to the original model – Qwen2.5-VL-7B (e.g., DocVQA, DUDE, InfogVQA). This suggests that directly optimizing both bounding-box detection and response generation in a joint manner poses certain challenges. Therefore, introducing RL1 first to optimize bounding-box detection separately is necessary.

Different training strategy. We attempted to directly learn the final response in Tab. 5, i.e., without the CoT setting. In this experiment, we used the same 30k dataset D_P . Specifically, SFT was trained solely on the chosen response (w/o bounding box); DPO [45] was trained on the win/lose preference response; Kahneman-Tversky Optimisation (KTO) [16] assigned "true" and "false" labels to "win" and "lose" response, respectively; and Proximal Policy Optimization (PPO) [46] sampled 15k preference data to train the reward

Table 4. Detection performance (Top-1 Accuracy@0.5) on the various benchmark, where both "Ours" and "Visual-CoT" similarly utilize LLaVA-1.5-7B as the base model. The amount of dense-labeled data with bounding box annotations is indicated in []. The **best** is highlighted.

Method	Doc/Text					Chart	General VQA		Relation Reasoning			Fine-grained	
	DocVQA	TextCaps	TextVQA	DUDE	SROIE	InfogVQA	Flickr30k	GQA	Open images	VSR	CUB	Avg	
VisCoT [438k] [48]	0.136	0.413	0.468	0.050	0.157	0.072	0.496	0.420	0.576	0.696	0.670	0.378	
VisRL-SFT [30k]	0.133	0.406	0.460	0.048	0.146	0.069	0.495	0.421	0.575	0.691	0.673	0.374	
VisRL-SFT+RL1	0.190	0.433	0.471	0.081	0.216	0.133	0.518	0.450	0.581	0.733	0.701	0.410	
VisRL-SFT+RL1+RL2	0.231	0.464	0.504	0.124	0.252	0.172	0.519	0.487	0.590	0.745	0.715	0.437	

Table 5. Different training strategies for directly fitting the final response in terms of Qwen2.5-VL-7B. Besides, we also ablated on RL2-Only. The red background indicates a decline in performance compared to the original model, while the **best** is highlighted.

Method	Doc/Text					Chart	General VQA		Relation Reasoning			Fine-grained	
	DocVQA	TextCaps	TextVQA	DUDE	SROIE	InfogVQA	Flickr30k	GQA	Open images	VSR	CUB	Avg	
Ori	0.836	0.760	0.847	0.606	0.789	0.685	0.601	0.467	0.289	0.581	0.583	0.640	
SFT	0.778	0.751	0.856	0.566	0.713	0.631	0.627	0.524	0.345	0.626	0.736	0.650	
DPO [45]	0.791	0.793	0.837	0.599	0.776	0.678	0.670	0.470	0.298	0.594	0.707	0.656	
KTO [16]	0.797	0.790	0.833	0.597	0.780	0.694	0.671	0.469	0.299	0.599	0.707	0.658	
PPO [46]	0.794	0.791	0.837	0.597	0.787	0.700	0.672	0.472	0.302	0.606	0.713	0.661	
VisRL-RL2-Only	0.816	0.769	0.897	0.602	0.800	0.681	0.637	0.652	0.786	0.827	0.882	0.759	
VisRL-Full	0.874	0.819	0.897	0.640	0.829	0.753	0.675	0.700	0.814	0.864	0.892	0.796	

model, while the remaining 15k data were used to train PPO. The results indicate that both previous RL and SFT methods exhibit limited effectiveness in directly fitting the final response, with some benchmarks even showing performance degradation (e.g. DocVQA, DUDE, etc.). In contrast, our training approach shown in Tab. 2, which incorporates the CoT reasoning in RL method, consistently improves model performance across all benchmarks.

Data Generation Pipeline. Our data generation process follows a self-evolution paradigm and adopts an actor-critic framework, where the actor is the SFTed model and the critics refer to the pre-SFT model. In this study, we conduct ablation experiments on several modules within the pipeline. As shown in Tab. 6, our objective is to maximize the proportion of positive instances in the constructed preference dataset, specifically WP-LN (win-positive, lose-negative), while minimizing the proportion of negative instances, WN-LP (win-negative, lose-positive). Experimental results indicate that under our pipeline, the proportion of positive instances is **64.64%**. After the first and second iterations – where the model undergoes full VisRL training and then participates in data construction again – this proportion increases by **3.18%** and **2.30%**, respectively. Beside, the number of valid data increases from 30k to 33k to **35k**. This is because, through iterative training, the model’s capability improves to overcome bottlenecks where obtaining the correct answer was previously outside of its ability. Specifically, some questions are inherently too difficult for model previously, meaning that no matter how many times the model responds, the answer remains incorrect, making it impossible to construct win/lose preference.

When replacing the critics with GPT-4o, the proportion of positive instances remains largely unchanged (64.64% vs. 65.31%), but this comes at the cost of increased token consumption and reduced response speed. Conversely, when substituting the critics with the SFTed model itself, the

Table 6. Ablation of our data generation pipeline in terms of Qwen2.5-VL-7B. We evaluate the data quality by comparing the IoU (Top-1 Accuracy@0.5) between the annotated preference data’s bounding boxes and the GT bounding boxes. "W" and "L" represent the win and loss in the data annotation, respectively, while "P" and "N" indicate positive or negative. Ideally, we expect WP-LN to be as large as possible, as highlighted in green.

	WP-LP	WP-LN	WN-LP	WN-LN	Data Num.
w GPT-4o-2024-11-20	0.00%	65.31%	1.32%	33.37%	47k
w SFTed Model	0.00%	54.68%	0.00%	45.32%	3k
w/o Bounding Box Critics	5.42%	31.02%	10.04%	53.51%	86k
w/o Diversity Controller	4.53%	52.02%	4.68%	38.77%	19k
VisRL-Full	0.43%	64.64%	1.64%	33.29%	30k
VisRL-Full-Iter1	0.45%	67.82%	0.82%	30.91%	33k
VisRL-Full-Iter2	0.47%	70.12%	0.00%	29.41%	35k

data volume decreases to **1/10** of the original (3k vs. 30k), accompanied by a **9.96%** reduction in the proportion of positive instances. Furthermore, omitting the evaluation of the generated bounding boxes results in a **33.62%** drop in the proportion of positive instances, while the absence of the diversity controller leads to a **12.62%** reduction. Considering these factors, we adopt the approach described in Sec. 3.1.

5. Conclusion

In this paper, we propose VisRL, a framework for learning intention-driven visual perception abilities from task feedback. This approach enables the model to undergo RL through self-evolution when trained on simple data with only final responses. Specifically, we introduce a novel pipeline for generating CoT preference data based on the model’s own actor-critic process, eliminating the need for external models or human annotations. Using these self-constructed data, we further optimize visual perception in two RL stages: (1) independently optimizing generated bounding box; (2) jointly optimizing both generated bounding box and response. Extensive experiments on various benchmarks and LMMs demonstrate the effectiveness of the proposed framework, establishing a solid foundation for future exploration.

References

- [1] Ebtesam Almazrouei, Ruxandra Cojocaru, Michele Baldo, Quentin Malartic, Hamza Alobeidli, Daniele Mazzotta, Guilherme Penedo, Giulia Campesan, Mugariya Farooq, Maitha Alhammedi, et al. Alghafa evaluation benchmark for arabic language models. In *Proceedings of ArabicNLP 2023*, pages 244–275, 2023. 2
- [2] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 2
- [3] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025. 6
- [4] Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiusi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024. 2
- [5] Mahtab Bigverdi, Zelun Luo, Cheng-Yu Hsieh, Ethan Shen, Dongping Chen, Linda G Shapiro, and Ranjay Krishna. Perception tokens enhance visual reasoning in multimodal language models. *arXiv preprint arXiv:2412.03548*, 2024. 3
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 2
- [7] Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*, 2024. 2
- [8] Jiaying Chen, Yuxuan Liu, Dehu Li, Xiang An, Weimo Deng, Ziyong Feng, Yongle Zhao, and Yin Xie. Plug-and-play grounding of reasoning in multimodal large language models. *arXiv preprint arXiv:2403.19322*, 2024. 3, 6
- [9] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. In *European Conference on Computer Vision*, pages 370–387. Springer, 2024. 1
- [10] Liang Chen, Lei Li, Haozhe Zhao, Yifan Song, and Vinci. R1-v: Reinforcing super generalization ability in vision-language models with less than \$3. <https://github.com/Deep-Agent/R1-V>, 2025. Accessed: 2025-02-02. 3
- [11] Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Zhenyu Tang, Li Yuan, et al. Sharegpt4video: Improving video understanding and generation with better captions. *Advances in Neural Information Processing Systems*, 37:19472–19495, 2025. 1
- [12] Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhi-jian Liu, Song Han, and Jiaya Jia. Longlora: Efficient fine-tuning of long-context large language models. *arXiv preprint arXiv:2309.12307*, 2023. 1
- [13] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, march 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna>, 3(5), 2023. 6
- [14] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017. 3
- [15] Yuhao Dong, Zuyan Liu, Hai-Long Sun, Jingkan Yang, Winston Hu, Yongming Rao, and Ziwei Liu. Insight-v: Exploring long-chain visual reasoning with multimodal large language models. *arXiv preprint arXiv:2411.14432*, 2024. 3, 6
- [16] Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*, 2024. 7, 8
- [17] Luciano Floridi and Massimo Chiriatti. Gpt-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30: 681–694, 2020. 2
- [18] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models, 2024. 6
- [19] Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, et al. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793*, 2024. 2
- [20] Anisha Gunjal, Jihan Yin, and Erhan Bas. Detecting and preventing hallucinations in large vision language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 18135–18143, 2024. 1
- [21] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 3
- [22] Yanming Guo, Yu Liu, Ard Oerlemans, Songyang Lao, Song Wu, and Michael S Lew. Deep learning for visual understanding: A review. *Neurocomputing*, 187:27–48, 2016. 1
- [23] Jiaming Ji, Jiayi Zhou, Hantao Lou, Boyuan Chen, Donghai Hong, Xuyao Wang, Wenqi Chen, Kailie Wang, Rui Pan, Jiahao Li, Mohan Wang, Josef Dai, Tianyi Qiu, Hua Xu, Dong Li, Weipeng Chen, Jun Song, Bo Zheng, and Yaodong Yang. Align anything: Training all-modality models to follow instructions with language feedback. 2024. 3
- [24] Fengqing Jiang. Identifying and mitigating vulnerabilities in llm-integrated applications. Master’s thesis, University of Washington, 2024. 2
- [25] Yang Jiao, Shaoxiang Chen, Zequn Jie, Jingjing Chen, Lin Ma, and Yu-Gang Jiang. Lumen: Unleashing versatile vision-centric capabilities of large multimodal models. *Advances in Neural Information Processing Systems*, 37:81461–81488, 2025. 1

- [26] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9579–9589, 2024. 1
- [27] Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. Step-dpo: Step-wise preference optimization for long-chain reasoning of llms. arXiv preprint arXiv:2406.18629, 2024. 5
- [28] Chengzu Li, Wenshan Wu, Huanyu Zhang, Yan Xia, Shaoguang Mao, Li Dong, Ivan Vulić, and Furu Wei. Imagine while reasoning in space: Multimodal visualization-of-thought. arXiv preprint arXiv:2501.07542, 2025. 3
- [29] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In International conference on machine learning, pages 12888–12900. PMLR, 2022. 3
- [30] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In International conference on machine learning, pages 19730–19742. PMLR, 2023. 3
- [31] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. arXiv preprint arXiv:2305.10355, 2023. 6
- [32] Yanwei Li, Chengyao Wang, and Jiaya Jia. Llama-vid: An image is worth 2 tokens in large language models. In European Conference on Computer Vision, pages 323–340. Springer, 2024. 1
- [33] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. 6
- [34] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. 1, 3, 6
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 6
- [36] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In European conference on computer vision, pages 216–233. Springer, 2024. 6
- [37] Yuecheng Liu, Dafeng Chi, Shiguang Wu, Zhanguang Zhang, Yaochen Hu, Lingfeng Zhang, Yingxue Zhang, Shuang Wu, Tongtong Cao, Guowei Huang, et al. Spatialcot: Advancing spatial reasoning through coordinate alignment and chain-of-thought for embodied task planning. arXiv preprint arXiv:2501.10074, 2025. 3
- [38] Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual-rft: Visual reinforcement fine-tuning. arXiv preprint arXiv:2503.01785, 2025. 3
- [39] Llama Meta. 3.2: Revolutionizing edge ai and vision with open, customizable models, 2024. URL: <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices>. 6
- [40] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. Advances in neural information processing systems, 35:27730–27744, 2022. 3
- [41] Thomas J Palmeri and Isabel Gauthier. Visual object understanding. Nature Reviews Neuroscience, 5(4):291–303, 2004. 1
- [42] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2
- [43] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. OpenAI blog, 1(8):9, 2019. 2
- [44] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In International conference on machine learning, pages 8748–8763. PmlR, 2021. 6
- [45] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. Advances in Neural Information Processing Systems, 36, 2024. 3, 7, 8
- [46] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017. 7, 8
- [47] Hao Shao, Yuxuan Hu, Letian Wang, Guanglu Song, Steven L Waslander, Yu Liu, and Hongsheng Li. Lmdrive: Closed-loop end-to-end driving with large language models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15120–15130, 2024. 1
- [48] Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. Visual cot: Unleashing chain-of-thought reasoning in multi-modal language models, 2024. 1, 3, 4, 6, 8
- [49] Haozhan Shen, Zilun Zhang, Qianqian Zhang, Ruochen Xu, and Tiancheng Zhao. Vlm-r1: A stable and generalizable r1-style large vision-language model. <https://github.com/om-ai-lab/VLM-R1>, 2025. Accessed: 2025-02-15. 3
- [50] Andreas Steiner, André Susano Pinto, Michael Tschanen, Daniel Keysers, Xiao Wang, Yonatan Bitton, Alexey Gritsenko, Matthias Minderer, Anthony Sherbondy, Shangbang Long, et al. Paligemma 2: A family of versatile vlms for transfer. arXiv preprint arXiv:2412.03555, 2024. 6
- [51] Chufeng Tang, Lingxi Xie, Xiaopeng Zhang, Xiaolin Hu, and Qi Tian. Visual recognition by request. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 15265–15274, 2023. 1, 3
- [52] InternLM Team. Internlm: A multilingual language model with progressively enhanced capabilities, 2023. 2
- [53] Omkar Thawakar, Dinura Dissanayake, Ketan More, Ritesh Thawkar, Ahmed Heakl, Noor Ahsan, Yuhao Li, Mohammed

- Zumri, Jean Lahoud, Rao Muhammad Anwer, Hisham Cholakkal, Ivan Laptev, Mubarak Shah, Fahad Shahbaz Khan, and Salman Khan. Llamav-01: Rethinking step-by-step visual reasoning in llms, 2025. 3
- [54] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. [arXiv preprint arXiv:2302.13971](#), 2023. 2
- [55] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. [arXiv preprint arXiv:2307.09288](#), 2023. 2
- [56] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. [arXiv preprint arXiv:2409.12191](#), 2024. 1
- [57] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. [Advances in Neural Information Processing Systems](#), 36:61501–61513, 2023. 1
- [58] XuDong Wang, Shaolun Zhang, Shufan Li, Konstantinos Kallidromitis, Kehan Li, Yusuke Kato, Kazuki Kozuka, and Trevor Darrell. Segllm: Multi-round reasoning segmentation. [arXiv preprint arXiv:2410.18923](#), 2024. 3
- [59] Jeremy M Wolfe and Todd S Horowitz. Five factors that guide attention in visual search. [Nature human behaviour](#), 1(3):0058, 2017. 1
- [60] Jiannan Wu, Muyan Zhong, Sen Xing, Zeqiang Lai, Zhaoyang Liu, Zhe Chen, Wenhai Wang, Xizhou Zhu, Lewei Lu, Tong Lu, et al. Visionllm v2: An end-to-end generalist multimodal large language model for hundreds of vision-language tasks. [Advances in Neural Information Processing Systems](#), 37:69925–69975, 2025. 3, 6
- [61] Penghao Wu and Saining Xie. V?: Guided visual search as a core mechanism in multimodal llms. In [Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition](#), pages 13084–13094, 2024. 3, 6
- [62] Wenshan Wu, Shaoguang Mao, Yadong Zhang, Yan Xia, Li Dong, Lei Cui, and Furu Wei. Mind’s eye of llms: visualization-of-thought elicits spatial reasoning in large language models. [Advances in Neural Information Processing Systems](#), 37:90277–90317, 2025. 3
- [63] Xueqing Wu, Yuheng Ding, Bingxuan Li, Pan Lu, Da Yin, Kai-Wei Chang, and Nanyun Peng. Visco: Benchmarking fine-grained critique and correction towards self-improvement in visual reasoning. [arXiv preprint arXiv:2412.02172](#), 2024. 3
- [64] Ziang Yan, Zhilin Li, Yinan He, Chenting Wang, Kunchang Li, Xinhao Li, Xiangyu Zeng, Zilei Wang, Yali Wang, Yu Qiao, et al. Task preference optimization: Improving multimodal large language models with vision task alignment. [arXiv preprint arXiv:2412.19326](#), 2024. 3
- [65] Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. Baichuan 2: Open large-scale language models. [arXiv preprint arXiv:2309.10305](#), 2023. 2
- [66] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. [arXiv preprint arXiv:2408.01800](#), 2024. 6
- [67] Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Guoyin Wang, Heng Li, Jiangcheng Zhu, Jianqun Chen, et al. Yi: Open foundation models by 01. ai. [arXiv preprint arXiv:2403.04652](#), 2024. 6
- [68] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. [arXiv preprint arXiv:2306.02858](#), 2023. 1
- [69] Jiacheng Zhang, Yang Jiao, Shaoxiang Chen, Jingjing Chen, and Yu-Gang Jiang. Eventhallusion: Diagnosing event hallucinations in video llms. [arXiv preprint arXiv:2409.16597](#), 2024. 1
- [70] Shilong Zhang, Peize Sun, Shoufa Chen, Minn Xiao, Wenqi Shao, Wenwei Zhang, Yu Liu, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on regionof-interest, 2024. In [URL https://openreview.net/forum](#). 1
- [71] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. [arXiv preprint arXiv:2205.01068](#), 2022. 2
- [72] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanhao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. [Advances in Neural Information Processing Systems](#), 36:46595–46623, 2023. 6