

SC-Tune: Unleashing Self-Consistent Referential Comprehension in Large Vision Language Models

Tongtian Yue^{1,3*} Jie Cheng^{2,3*} Longteng Guo^{1,3*} Xingyuan Dai^{2,3}
 Zijia Zhao^{1,3} Xingjian He^{1,3} Gang Xiong^{2,3} Yisheng Lv^{2,3} Jing Liu^{1,3†}
¹Laboratory of Cognition and Decision Intelligence for Complex Systems, CASIA
²State Key Laboratory of Multimodal Artificial Intelligence Systems, CASIA
³School of Artificial Intelligence, University of Chinese Academy of Sciences

Abstract

Recent trends in Large Vision Language Models (LVLMs) research have been increasingly focusing on advancing beyond general image understanding towards more nuanced, object-level referential comprehension. In this paper, we present and delve into the self-consistency capability of LVLMs, a crucial aspect that reflects the models' ability to both generate informative captions for specific objects and subsequently utilize these captions to accurately re-identify the objects in a closed-loop process. This capability significantly mirrors the precision and reliability of fine-grained visual-language understanding. Our findings reveal that the self-consistency level of existing LVLMs falls short of expectations, posing limitations on their practical applicability and potential. To address this gap, we introduce a novel fine-tuning paradigm named **Self-Consistency Tuning (SC-Tune)**. It features the synergistic learning of a cyclic describer-locator system. This paradigm is not only data-efficient but also exhibits generalizability across multiple LVLMs. Through extensive experiments, we demonstrate that SC-Tune significantly elevates performance across a spectrum of object-level vision-language benchmarks and maintains competitive or improved performance on image-level vision-language benchmarks. Both our model and code will be publicly available at <https://github.com/ivattyue/SC-Tune>.

1. Introduction

Recently, Large Vision Language Models have witness remarkable progress [13, 25, 59]. By introducing learnable parameters to map visual features to the semantic space of Large Language Models (LLMs) [7, 12, 38, 45], LVLMs demonstrate strong capabilities for visual content percep-

*Equal Contribution.

†Corresponding author.




Figure 1. Upper: two examples demonstrating the shortfall in self-consistent referential comprehension of LVLMs, which are attributed to limited grounding (the first example) and limited captioning (the second example) capabilities, respectively. Lower: We use Pr@0.5 as the self-consistency evaluation metric. We consider it to be correct when the IoU between prediction bbox and the ground truth is greater than 0.5. The self-consistency levels of different LVLMs show pronounced performance gap between in-domain (RefCOCO) and out-of-domain (Object365 and OpenImages) datasets.

tion and expression. A line of works [3, 9, 10, 57] go beyond basic image-level comprehension, and have emerged to cultivate object-level referential comprehension capability with LVLMs, namely *understanding* or *identifying* a specific object within an image. Referential comprehension capability bases on two symmetry object-level tasks, *i.e.*, referring expression generation (REG) and referring expression comprehension (REC). REG is to describe an object

in an image with discriminative referring expression, while REC is the reverse task of REG, aiming at localizing a particular object given an expression. Recent LVLMs handle spatial coordinate inputs and outputs in natural language, thereby unifying both REG and REC tasks in sequence modeling framework, which are optimized with multi-task instruction tuning on a large-scale collection of fine-grained region-text data [22, 32, 36, 55].

A LVLM model that generates a referring expression for an object (REG) and then accurately locates back the object based on that expression (REC) demonstrates a deeper understanding of the content and context of the image. Such bbox-caption-bbox cyclical consistency is a crucial capability of LVLMs performing object-level understanding tasks, which we refer to as *self-consistency* on referential comprehension (see Figure 1). It ensures that both the generation and comprehension aspects are aligned and accurate, leading to more robust models. If a model generates an incorrect referring expression or fails to accurately localize an object, the inconsistency becomes apparent when attempting to reverse the task. In real-world scenarios where accuracy in both generating and comprehending referential expressions is crucial (such as in autonomous navigation, visually impaired assistance, embodied agents, or interactive AI systems), self-consistency ensures reliability and usability of the model.

Surprisingly, through preliminary experiments, we found that the *self-consistency capability of current object-level LVLMs (e.g. MiniGPT-v2 [9] and Qwen-VL [3]) dramatically drops on out-of-domain images*. As illustrated in Figure 1, we randomly select 4k bounding boxes from both the in-domain (seen during training) dataset RefCOCO [55] and the out-of-domain (unseen during training) datasets, *i.e.*, OpenImages [21] and Object365 [40], respectively. Starting from these bboxes, we sequentially perform REG and REC and examine whether the model can locate the correct region based on its generated region caption. As is shown in the figure, there exists pronounced performance gap of self-consistency between in-domain and out-of-domain images, implying poor generalization of their referential comprehension ability.

In this paper, we foster the self-consistent referential comprehension capability of object-level LVLMs by proposing a model fine-tuning paradigm named *Self-Consistency Tuning (SC-Tune)*. We go beyond conventional multi-task learning of both REG and REC tasks, which neglects their consistent correlation. Self-consistency tuning centers on the synergistic learning of a cyclic dual-component system: a describer and a locator, which are two roles of the same pre-trained LVLM. In this system, the describer generates contextually enriched captions from bboxes. Subsequently, the locator operates on those generated captions, utilizing them as directives for

precise object localization within images. The describer is trained through a bbox-caption-bbox self-consistency reward cycle, under Proximal Policy Optimization (PPO) [39] reinforcement learning paradigm. This mechanism ensures the generation of captions that not only describe but also discriminatively guide the localization process. Conversely, the locator is supervisory trained to parse nuanced linguistic cues from the evolving synthetic captions of the describer, enabling precise and robust object localization even in out-of-domain images.

The two components are updated under an iterative training cycle: freezing one component when training the other one. This alternating cycle is complemented by the synchronization of their parameters post each training cycle, fostering a harmonious growth of the overall system in fine-grained understanding. By refining the interplay between caption generation and object localization, we enhance the model’s self-consistency in both in-domain and out-of-domain images. Such Self-consistency tuning thereby increases the performance and robustness of LVLMs in fine-grained referential comprehension.

Our main contributions can be summarized as follows:

- We propose self-consistency as a crucial metric for model reliability in fine-grained referential comprehension and systematically examine this capability on existing LVLMs.
- We propose self-consistency tuning to effectively foster the self-consistent referential comprehension capability of LVLMs, which is data efficient and generalizable across multiple LVLMs.
- By incorporating our proposed self-consistency tuning with state-of-the-art LVLMs, we observe notable enhancements in zero-shot performance across multiple object-level vision-language benchmarks, while maintaining competitive or even improved performance on image-level vision-language benchmarks.

2. Related Work

Referential Comprehension of LVLMs. In everyday human interactions, referencing specific objects or areas within a visual context is a frequent occurrence. Therefore, it is significant to augment LVLMs with robust referential comprehension capabilities. The mainstream paradigms for integrating location information into the understanding of LVLMs can be summarized as the followings, namely textual coordinate representation and regional feature extraction. The representative work of the former can be traced back to Pix2Seq [11]. It leverages discrete coordinate tokens to encode spatial information, thereby unifying the referential comprehension training into the sequence modeling task. Notable works include OFA [48], Unified-io [29], Shikra [10] and Kosmos-2 [36]. For the latter paradigm, represented by PVIT [8] and GPT4RoI [57], it utilizes Re-

gion of Interest (ROI) [17] to extract object-level features, then align these features to the textual modality. However, these works ignore the native self-consistency between reference generation and grounding when training referential comprehension capabilities. The disruption in the bidirectional consistency between visual and textual modalities indicates unsatisfactory alignment quality.

Reinforcement Learning in LLMs. Recent advancements in Large Language Models (LLMs) have been significantly influenced by Reinforcement Learning from Human Feedback (RLHF). This method has proven effective in aligning LLMs with human-centric values and preferences [4, 34]. RLHF hinges on training a reward model (RM) aligned with human preferences, subsequently refining LLMs based on the reward signals provided by the RM. Notably, LLMs are often updated using PPO algorithm [33, 41], with the addition of a Kullback–Leibler (KL) penalty to control deviations from the initial model [26, 34]. Extending beyond text, RLHF has been adapted for image generation. By introducing a RM, the alignment quality of prompts and images is improved [24], making it more consistent with human preferences [52]. To save the resources on collecting high-quality human preferences, [23] proposes to use AI instead of human in labeling preferences, a method termed RL from AI Feedback (RLAIF), which has demonstrated comparable efficacy. Similarly, in this work, we employ the locator to provide feedback as the reward signal for the describer. By treating the caption generation as a sequential decision-making problem, we leverage reinforcement learning to fine-tune the describer towards self-consistency.

3. Self-Consistency Tuning

In this section, we introduce our training framework aimed at eliminating the sense of isolation between tasks, enhancing the model’s self-consistent referring comprehension capability. This approach cyclically fine-tunes a dual-component system towards self-consistency. We briefly review the architectural design of current LVLMs in section 3.1. Subsequently, in section 3.2, we provide a detailed introduction to our SC-Tune pipeline, as illustrated in Figure 3. Then we elaborate the training process and loss function for each component respectively in section 3.3 and 3.4.

3.1. LVLM Architecture

The structure of current LVLMs can be primarily summarized into three parts: the visual encoder, LLM and the bridging component that connects the two modalities. Initially, input images are processed through the visual encoder for feature extraction. These visual features are then mapped onto the semantic space of LLM

through the bridging component, where they are concatenated with the input text and fed into LLM for content generation. LVLMs acquire referential comprehension capabilities through region-text data, which is sourced from both small-scale human annotations [22, 32, 55] and extensive weakly supervised data [36]. LVLMs proportionally resize the top-left and bottom-right coordinates to integers within a fixed range (e.g., [0,1000] for Qwen-VL and [0,100] for MiniGPT-v2). Subsequently these discrete and finite coordinates are represented in textual form. In REG and REC tasks, these textually represented coordinates serve as either input or expected output for LVLMs. Both tasks utilize a language modeling loss for optimization. However, the REC and REG capabilities trained in this manner are sub-optimal due to a lack of symmetry.

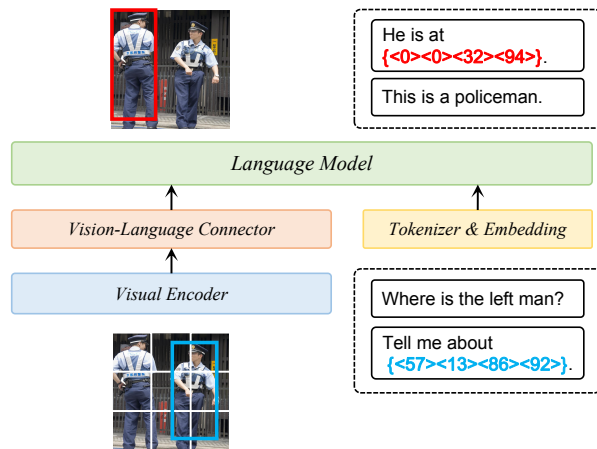


Figure 2. Architecture of object-level LVLMs. They mainly comprises three parts: the visual backbone for feature extraction, the vision-language connector for semantic alignment and a LLM. It has preliminary REC and REG capabilities by generating and understanding coordinates represented by text.

3.2. Overview of SC-Tune

Our framework regards two roles, namely “describer” and “locator”, of one pre-trained LVLM as a dual-component system. The describer focuses on generating detailed captions based on given bboxes, while the locator aims to reconstruct these bboxes from the self-generated captions. The central goal of our fine-tuning process is to find a set of parameters that has both enhanced capabilities.

To achieve that goal, we draw inspiration from the “target network” utilized in Deep-Q Networks (DQN) [31]. In DQN, one network is frozen as the “target network” to provide training signals for the other network. As the other one inches closer to its goal, the frozen target network is updated by the parameters copied from the other network to facilitate continuous improvement. Similarly in our framework, we treat each component as the target network for the

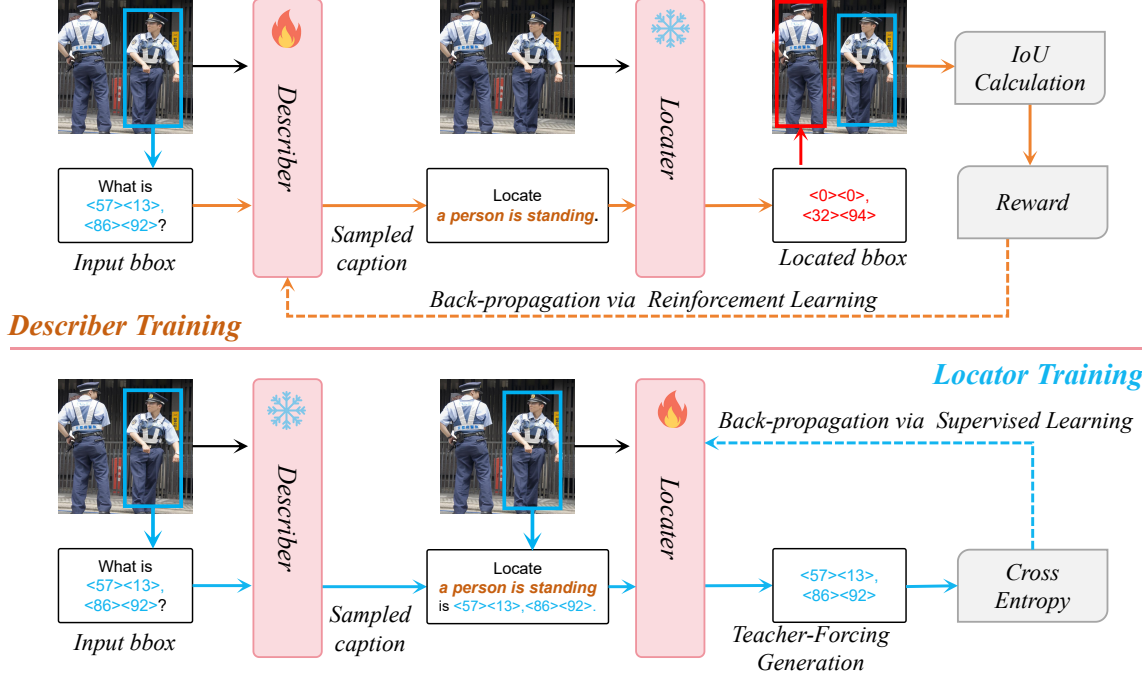


Figure 3. An illustration of our *SC-Tune* framework. It mainly comprises describer training cycle and locator training cycle. Each cycle employs a bbox-caption-bbox pipeline with respective loss function designed for fostering the self-consistent referential comprehension capability of object-level LVLMs. Training alternates between these two cycles, with parameter synchronization post each training cycle.

other to generate training signals. Moreover, we iteratively train the two components and synchronize their parameters when switching the focus of training. This method not only inherits the advantage of the target network to make training more stable, but also promotes the mutual improvement of the two capabilities within one model.

Specifically, the parameters of the describer and locator are the same initially. During the training stage of describer, we freeze the locator and only tune the describer for a pre-defined number of steps, H . Following this, we synchronize the parameter of describer to the locator. Conversely, in the training stage of locator, we freeze the describer and only tune the locator for H steps. After this stage, we synchronize the parameter of locator to the describer. In each training stage, we employ a bbox-caption-bbox pipeline (show in Figure 3), leveraging the task symmetry and consistent correlation for self-consistency fine-tuning. This cyclical training process continues until it effectively integrates the improvements of both capabilities into a cohesive model.

3.3. Describer Training

This subsection details the training process for the describer, which is responsible for generating detailed captions \hat{C} from given bounding boxes B . Concurrently, the locator is frozen to localize bboxes \hat{B} for these self-generated captions \hat{C} . In this reconstruction process, our aim is to maximize the similarity between the input and output

bboxes, *i.e.*, $\text{IoU}(B, \hat{B})$, for self-consistency, which implicitly requires these captions \hat{C} to be informative and discriminative.

However, the operation of sampling from vocabulary during generation is non-differentiable, resulting that the describer cannot be updated directly through back-propagating reconstruction loss. To address this challenge, we employ reinforcement learning (RL), which does not require derivation of the action sampling process.

PPO Loss. In the setting of RL, both the describer and locator are regarded as “agents” with policy π parameterized by θ and ϕ , respectively. Initially, $\theta = \phi$. At timestep t , the previously generated tokens $\{\hat{c}_1, \dots, \hat{c}_{t-1}\}$, the input image I , and the bbox B are regarded as the “state” s_t . The generating caption token \hat{c}_t is served as the “action” taken with the policy π_θ , *i.e.*, $\hat{c}_t \sim \pi_\theta(\cdot|s_t)$. The objective is to maximize the cumulative rewards $\sum_{t=1}^T R(\hat{c}_t, s_t)$, where T is the terminal token generation step. To achieve the objective, we use PPO algorithm to optimize the policy in the trust region for stable training. The RL loss function is formulated as follows:

$$\mathcal{L}^{RL}(\theta) = -\mathbb{E}_t [\min(r_t(\theta)A_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)] \quad (1)$$

where $r_t(\theta)$ is the importance sampling ratio: $r_t(\theta) = \frac{\pi_\theta(\hat{c}_t|s_t)}{\pi_{\theta_{\text{old}}}(\hat{c}_t|s_t)}$, and θ_{old} is the policy parameters before update.

ϵ is a hyperparameter and A_t is the advantage function. clip function is to limit values to a range.

Reward. Given the self-generated caption $\hat{C} = \{\hat{c}_1, \dots, \hat{c}_T\}$ and input image I , the frozen locator π_ϕ generates a predicted bbox \hat{B} , *i.e.*, $\hat{B} = f(\pi_\phi(\hat{C}, I))$, where f represents regular expression search from generated sentence to parse bbox. Due to the goal of reconstructing the original bbox B , it is intuitive to set the $\text{IoU}(B, \hat{B})$ as the reward function. Notably, it is a sentence-level reward for describer. However, this reward function does not constrain the semantic information of captions, probably causing forgetfulness in linguistic rule. Therefore, we add a token-level KL penalty from the initial model to mitigate over-optimization of the reconstruction reward. The complete reward function is formulated as follows:

$$R(\hat{c}_t, s_t) = \mathbb{1}_{t=T}(t) \cdot \text{IoU}(B, \hat{B}) - \beta \log \frac{\pi_\theta(\hat{c}_t | s_t)}{\pi_{\theta_{\text{ref}}}(\hat{c}_t | s_t)} \quad (2)$$

where $\mathbb{1}$ is the indicator function, θ_{ref} is the initial parameters, and β is the KL coefficient.

Value-Free Advantage. The advantage function is used to evaluate the long-term benefits of a chosen action compared to other actions. In standard PPO algorithm, it needs another value function to compute the advantage, which requires additional network and computing resources. Alternatively, we apply the form of advantage function in ‘‘REINFORCE with baseline’’ algorithm [43] formulated in Eq. (3), where the baseline is employed to reduce variance theoretically.

$$A_t(\hat{c}_t, s_t) = \sum_{i=t}^T R(\hat{c}_i, s_i) - R(c_T^*, s_T^*) \quad (3)$$

where the superscript $*$ represents the baseline. We choose greedily decoded captions C^* generated from initial model θ_{ref} as the baseline. Note that there is no need of gradients for action \hat{c}_t , reward $R(\hat{c}_t, s_t)$, advantage $A(\hat{c}_t, s_t)$, and old sampling probability $\pi_{\theta_{\text{old}}}(\hat{c}_t | s_t)$, only the latest sampling probability $\pi_\theta(\hat{c}_t | s_t)$ needs to calculate gradient in the training loss (1).

3.4. Locator Training

In this subsection, we detail the training process for the locator. Unlike the describer that is in the middle of the bbox-caption-bbox pipeline, the gradient of the loss function can be propagated back to the locator directly. Therefore, we adopt self-supervised learning on locator for simplicity.

Specifically, during the training of locator π_ϕ , we freeze the describer π_θ to generate pseudo captions \hat{C} from given

bboxes B and image I . Then we fine-tune locator with the data (B, I, \hat{C}) using loss derived from maximum likelihood estimation (MLE). Assuming the tokens of tokenized instructions building upon (B, I, \hat{C}) are denoted as $\{w_1, \dots, w_T\}$, the training loss of locator is formulated as follows:

$$\mathcal{L}^{\text{MLE}}(\phi) = -\mathbb{E}_t [\log P(w_t | w_1, \dots, w_{t-1})] \quad (4)$$

Our framework iteratively fine-tunes the describer and locator for H steps using PPO loss (1) and MLE loss (4), respectively. Post each training phase, we synchronize their parameters to facilitate synergistic and continuous improvement, as detailed in section 3.2.

4. Experiments

4.1. Experimental Settings

Baseline LVLMS. We verify the effectiveness and universality of our proposed training framework on two prevalent fine-grained LVLMS, *i.e.* Qwen-VL [3] and MiniGPT-v2 [9]. Qwen-VL adopts Qwen-7B [2] as its foundation component, and uses the Vision Transformer (ViT) [14] as the visual encoder, with pre-trained weights from Openclip’s ViT-bigG [19]. For MiniGPT-v2, it takes the visual tokens from EVA-Clip [15] and leverage LLaMA2-Chat (7B) [46] as LLM. Uniformly, for both models, we inherit the checkpoints from stage 3 and execute *SC-Fune* on this basis to further refine the referential comprehension capabilities.

Implementation Details. We adopt AdamW [28] as the optimizer. Parameter synchronization is performed every 200 steps and the training stage is then switched. We train both baseline models for 1 epoch which contains 6 stage switches. Both models are trained with a batch size of 128 and weight decay of 0.1. The learning rate is set to 5e-7 and 1e-6 for describer and locator training stage, respectively. We set the KL coefficient and PPO epoch to 0.01 and 2 for reinforcement learning. The input image is resized to 448 × 448 without any additional data-augmentation. The model is trained on 8 NVIDIA A100 GPUs, lasting roughly 15 hours for Qwen-VL and 10 hours for MiniGPT-v2.

Training Dataset. Our *SC-Tune* is data-efficient. We leverage about 166K images and corresponding bbox locations sampled from Object365 [40] Dataset. Training does not require any text annotations.

Evaluation Datasets and Metrics. We evaluate the baseline models on various object-level and image-level benchmarks to validate the referential comprehension enhanced by our framework. For REC and REG, we consider RefCOCO+/g [32, 55], ReferItGame [20] and Flickr30K Entities [37]. For referential question answering, we evaluate

Table 1. Self-consistency evaluation on various benchmarks. We report self-consistency level using accuracy, where a sample is considered as right when IoU between prediction and ground-truth is higher than 0.5.

Method	Object365	OpenImages	RefCOCO
Qwen-VL	76.9	52.9	87.9
+ <i>SC-Tune</i> (Object365)	94.1	68.8	93.8
+ <i>SC-Tune</i> (OpenImages)	89.6	73.6	92.0
MiniGPT-v2	65.4	38.2	87.2
+ <i>SC-Tune</i> (Object365)	76.6	48.5	90.4

on Visual-7W [60] and PointQA-Local [30]. Furthermore, we conduct assessments on the image-level benchmarks to substantiate the enhanced fine-grained alignment quality brought by our *SC-Tune*. For Image Caption, we employ Nocaps [1] and Flickr30K [37]. Regarding the question-answering tasks, we adopt VQAv2 [16] and GQA [18] as the evaluation benchmarks.

For all benchmarks, we follow the prompt templates used in the original instruction tuning of the baseline models. Greedy search is used for decoding. In REC task, the bounding box predicted by the model is considered as correct for reporting accuracy if its intersection over union (IoU) between prediction and ground-truth is higher than 0.5. In REG and Image Caption tasks, following previous works, we report results using METEOR [5] and CIDEr [47] metrics. In QA tasks, we leverage top-1 accuracy to measure the matching degree between the response and ground-truth.

4.2. Self-Consistency Evaluation Results

We first investigate the consistency enhancement brought by *SC-Tune*, as detailed in Table 1. By default, we perform subsequent performance evaluations using checkpoints obtained from training with Object365. We additionally present the results of Qwen-VL trained on OpenImages as supplement. More detailed experiments can be referred to in the Appendix. For both baseline models, the application of *SC-Tune* on a small scale Object365 dataset results in a notable improvement in self-consistency. The enhancement is applicable to both in-domain data RefCOCO, and out-of-domain data OpenImages. *SC-Tune* demonstrates robustness to data distribution, which is reflected in the similar average increases in self-consistency level: 13 points for Object365 and 12.5 points for OpenImages in Qwen-VL.

4.3. Object-Level Evaluation Results

Reference Expression Comprehension Results Table 2 demonstrates the improvement of referring grounding ability of baseline models under the out-of-domain setting before and after *SC-Tune*. For Qwen-VL, compared to the baseline, using *SC-Tune* achieves an accuracy improvement

of 7.8% and 13.34% on the test split of two prevalent datasets, *i.e.* ReferItGame and Flickr30K, respectively. It is worth noting that under the zero-shot setting, the performance of Qwen-VL equipped with our *SC-Tune* can even be comparable with the fully supervised model represented by CLIP-VG [51]. Similar patterns are also shown on another baseline MiniGPT-v2, with performance improvements of 7.6% and 4.6% respectively. These results provide direct evidence that *SC-Tune* can refine a more general referential comprehension capability.

Table 2. REC results on ReferItGame and Flickr30k Entities Datasets. We report the accuracy metric for all methods.

Method	Zero-Shot	ReferIt	Flickr30K
<i>Specialists</i>			
SeqTR [58]	✗	69.66	81.23
VLTVG [54]	✗	71.98	79.84
CLIP-VG [51]	✗	70.89	81.99
<i>Generalists</i>			
Qwen-VL	✓	61.48	64.15
Qwen-VL + <i>SC-Tune</i>	✓	69.28	77.49
MiniGPT-v2	✓	36.05	55.39
MiniGPT-v2 + <i>SC-Tune</i>	✓	43.68	60.03

Although the main focus of our training framework is not to achieve further improvement on a specific benchmark by adapting to its distribution, from the perspective of completeness, we report the results on the in-domain REC dataset, *e.g.* RefCOCO [55], in Table 3. Introducing *SC-Tune* has a slight performance drop compared to the baselines on the three evaluation subsets of RefCOCO. Intuitively, as tuning progresses, the caption style generated by the model gradually deviates from the in-domain style. The adaptation to this deviation will inevitably lead to a decrease of in-domain capability. Therefore, we argue that this fluctuation is acceptable as a trade-off with the gain in generalization ability. Furthermore, following RLHF pipeline which mixes the pretraining gradients into the original gradients to fix the performance regressions on benchmarks [35], we implement this regularization in our training as an attempt. Specifically, during the locator training stage, we augmented the synthesized captions with additional supervised data from RefCOCO. As illustrated in the Table 3, this strategy enables further enhancement of in-domain performance. However, as previously mentioned, this is not the primary focus of our research.

Reference Expression Generation Results. Table 4 evaluates our *SC-Tune* on REG benchmarks under the out-of-domain setting. Symmetrically, we adopt ReferItGame and Flickr30K for evaluation. Baseline model shows strong grounded captioning capabilities after *SC-Tune*, performing better than other zero-shot counterparts. For Qwen-VL, compared with the baseline, our *SC-Tune* outperforms it by

Table 3. REC results on RefCOCO Dataset under in-domain setting. \mathcal{L}_{sup} represents the incorporation of in-domain RefCOCO data in the locator training stage. The performance metrics for the two baseline models are reproduced using their official open-source code and checkpoints.

Method	RefCOCO		
	val	test A	test B
<i>Specialists</i>			
G-DINO-L [27]	90.56	93.19	88.24
UNINEXT-H [53]	92.64	94.33	91.46
ONE-PEACE [49]	92.58	94.18	89.26
<i>Generalists</i>			
VisionLLM-H [50]	-	86.70	-
OFA-L [48]	79.96	83.67	76.39
Shikra-7B [10]	87.01	90.61	80.24
Qwen-VL	88.88	92.27	84.30
Qwen-VL + <i>SC-Tune</i>	88.04	90.77	84.62
Qwen-VL + <i>SC-Tune</i> + \mathcal{L}_{sup}	89.61	93.43	85.72
MiniGPT-v2	84.84	89.49	81.08
MiniGPT-v2 + <i>SC-Tune</i>	83.59	88.50	80.22
MiniGPT-v2 + <i>SC-Tune</i> + \mathcal{L}_{sup}	85.29	90.77	82.33

over 20 CIDEr points across all tasks of ReferItGame and Flickr30K. Similarly, significant improvements can also be observed on MiniGPT-v2, *i.e.* an increase of 9 points in average CIDEr.

We report the performance on the RefCOCO dataset in Table 5 as a complementary evaluation of the in-domain REG capability. Facilitating more informative and unique descriptions by *SC-Tune*, we found that the in-domain captioning capabilities of baseline models can be further improved. For MiniGPT-v2, this benefit makes it a generalist model that can even compete with specialist models represented by PFOS [42] in terms of CIDEr scores, *e.g.*, 87.7 for PFOS vs 82.5 for MiniGPT-v2 in TestA split and 132.9 for PFOS vs 131.1 for MiniGPT-v2 in TestB split. The improvement in Qwen-VL is also worthy of recognition.

Table 5. REG results on RefCOCO Dataset under in-domain setting.

Method	RefCOCO			
	TestA		TestB	
	Meteor	CIDEr	Meteor	CIDEr
<i>Specialists</i>				
SLR [56]	26.8	69.7	32.9	132.3
EU [44]	31.1	83.7	33.0	133.3
PFOS [42]	30.3	87.7	34.1	132.9
<i>Generalists</i>				
Qwen-VL	23.8	88.0	24.9	110.0
Qwen-VL + <i>SC-Tune</i>	26.3	104.0	28.7	129.1
MiniGPT-v2	16.9	62.1	21.4	113.7
MiniGPT-v2 + <i>SC-Tune</i>	23.2	82.5	26.2	131.1

Referential Question Answering Results. Apart from the two abilities of referential comprehension, *i.e.* REC and

REG, we also evaluate the ability of referential question answering. For Visual-7W [60], it features a which setting, requiring the model to select one matching box from four options based on the given reference. For Local-QA [30], the models are asked to answer questions based on the given bbox. We report the performance on both benchmarks in Table 6. Through tuning Qwen-VL, the accuracy on Visual-7W and Local-QA increases by 11.3% and 8.1% respectively. The patterns shown on MiniGPT-v2 are consistent, which fully proves the universality of *SC-Tune*.

Table 6. Referential QA Results on Visual-7W and Local-QA. The evaluation is under zero-shot setting. We report the accuracy metric for both benchmarks.

Method	Visual-7W	Local
Qwen-VL	34.77	60.75
Qwen-VL + <i>SC-Tune</i>	46.09	72.13
MiniGPT-v2	27.53	49.63
MiniGPT-v2 + <i>SC-Tune</i>	35.29	51.15

4.4. Image-Level Evaluation Results

Intuitively, when the quality of referential comprehension is optimized, it enhances the perceptual and understanding capabilities at the image level. Consequently, we conducted extensive evaluation on multiple image-level benchmarks, with the results presented in Table 7. The results indicate that *SC-Tune* demonstrates improvements across multiple coarse-grained benchmarks. Given that *SC-Tune* requires only 166K images and does not necessitate text annotations, this enhancement is satisfactory.

Table 7. Image-level evaluation results on image caption and QA task. For QA tasks, accuracy metrics are employed, while for image caption, we use CIDEr score. For VQA and GQA, following previous works, we report the metrics on the val split and testdev split, respectively.

Method	Nocaps	Flick30K	VQA	GQA
<i>Specialists SOTAs</i>				
	127.0	84.5	86.1	72.1
<i>Generalists</i>				
BLIP-2	103.9	71.6	65.0	32.3
InstructBLIP	121.9	82.8	-	49.5
Shikra	-	73.9	77.4	-
Qwen-VL	120.2	81.0	78.2	57.5
Qwen-VL + <i>SC-Tune</i>	121.4	86.0	79.0	58.4
MiniGPT-v2	93.5	77.1	72.6	59.1
MiniGPT-v2 + <i>SC-Tune</i>	94.6	78.4	73.4	59.9

4.5. Ablation Studies

Synergistic Effect of Iterative Training. In this section, we conduct an ablation study on the strategy of iterative training with parameter exchange, to demonstrate the synergistic effect of our *SC-Tune*. The results are presented in Table 8 based on Qwen-VL. Taking the second row as an example, we freeze the locator while tuning the describer,

Table 4. REG results on ReferItGame and Flickr30k Entities Datasets under out-of-domain setting. Following previous work, ReferItGame test split is divided into test A and test B based on whether the sample is human.

Method	ReferIt				Flickr30K			
	test A		test B		val		test	
	Meteor	CIDEr	Meteor	CIDEr	Meteor	CIDEr	Meteor	CIDEr
<i>Specialists</i>								
SLR [56]	3.5	10.9	2.8	8.5	3.0	14.1	2.7	13.9
EU [44]	2.0	9.5	1.9	7.7	2.7	14.8	2.4	15.0
DisCLIP [6]	9.7	8.8	9.0	6.3	9.5	6.7	9.6	6.8
<i>Generalists</i>								
Qwen-VL	7.5	19.2	9.3	35.0	16.8	53.6	17.6	57.6
Qwen-VL + <i>SC-Tune</i>	11.6	43.2	13.1	51.7	18.5	73.7	18.4	80.9
MiniGPT-v2	8.9	33.3	8.2	47.5	15.8	78.7	16.3	84.0
MiniGPT-v2 + <i>SC-Tune</i>	11.1	39.5	11.2	56.8	16.3	87.1	16.9	93.5

training it for an equivalent number of steps as in the iterative approach. This process resulted in a specialized model with enhanced captioning capabilities, and vice versa. The results, however, indicate that even specialized model for captioning does not perform as well as the model with *SC-Tune* on REG tasks. It is equally applicable to the grounding specialized model. These results demonstrate the superiority of consistently improve both capabilities over the isolated optimization of either.

Table 8. Ablation study on synergistic effect of iterative training. We report the performance in REC, REG of ReferItGame and Local-QA based on Qwen-VL. For REC and QA tasks, accuracy metrics are employed, while for REG, we use CIDEr score.

Tuning Component		REC@ReferIt	REG@ReferIt		QA@Local
Locator	Describer	test	test A	test B	test
✗	✗	61.48	19.2	35.0	60.75
✗	✓	52.92	38.2	45.8	70.12
✓	✗	64.82	19.8	35.1	69.54
✓	✓	69.28	43.2	51.7	72.13

Training Steps for Each Cycle. In this section, we analyze the training steps for each cycle. The analysis pertaining to the training steps is presented in Table 9. The results indicate that selecting appropriate training steps to balance fluctuations in various capabilities is of significance.

Table 9. Ablation study on training steps based on Qwen-VL.

Steps	REC@ReferIt	REG@ReferIt		QA@Local
	test	test A	test B	test
50	68.41	34.3	46.6	71.15
200	69.28	43.2	51.7	72.13
1000	67.45	39.3	48.3	71.12

Training Data Source. In this section, we conduct an ablation study on the source of training data based on Qwen-VL. It demonstrates the universality of *SC-Tune* across various data distributions. Specifically, we selected another out-of-domain object detection dataset OpenImages [21], and

an in-domain dataset Visual Genome [22]. In terms of data filter and training strategies, we keep consistent with which we conduct in Object365. The experiment results in Table 10 indicate that *SC-Tune* achieves a significant performance improvement over the baseline model across all three data distributions.

Table 10. Ablation study on data source based on Qwen-VL.

Data Source	REC@ReferIt	REG@ReferIt		QA@Local
	test	val	test	test
	61.48	19.2	35.0	60.75
Object365	69.28	43.2	51.7	72.13
OpenImages	66.07	35.5	47.6	72.50
Visual Genome	63.28	26.3	41.2	68.56

5. Conclusion

In this work, we reveal a notable shortfall in the self-consistency levels of current LVLMs. Responding to this gap, we propose the Self-Consistency Tuning (*SC-Tune*), an object-level fine-tuning paradigm designed to improve self-consistent referential comprehension capability. Central to *SC-Tune* is a cyclic training loop of a dual-component system (*i.e.* describer and locator), which promotes a harmonious growth of the overall system in fine-grained understanding. *SC-Tune* not only is data-efficient but also demonstrates robust generalizability across multiple LVLMs. Our comprehensive experiments show that *SC-Tune* significantly improves performance across various object-level vision-language benchmarks, while simultaneously sustaining or augmenting performance in image-level vision-language benchmarks. We plan to release our model and code to the public, expecting to foster future research in this direction.

Acknowledgements We thank all the insightful reviewers for the helpful suggestions. This work was supported by the National Science and Technology Major Project (No.2022ZD0118801), National Natural Science Foundation of China (U21B2043, 62206279).

References

- [1] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019. 6
- [2] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 5
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 2023. 1, 2, 5
- [4] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022. 3
- [5] Satyanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 6
- [6] Lior Bracha, Eitan Shaar, Aviv Shamsian, Ethan Fetaya, and Gal Chechik. Disclip: Open-vocabulary referring expression generation. *arXiv preprint arXiv:2305.19108*, 2023. 8
- [7] 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. 1
- [8] Chi Chen, Ruoyu Qin, Fuwen Luo, Xiaoyue Mi, Peng Li, Maosong Sun, and Yang Liu. Position-enhanced visual instruction tuning for multimodal large language models. *arXiv preprint arXiv:2308.13437*, 2023. 2
- [9] Jun Chen, Deyao Zhu¹, Xiaoqian Shen¹, Xiang Li, Zechun Liu², Pengchuan Zhang, Raghuraman Krishnamoorthi², Vikas Chandra², Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: Large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. 1, 2, 5
- [10] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023. 1, 2, 7
- [11] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*, 2021. 2
- [12] 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. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2023. 1
- [13] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. 1
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [15] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023. 5
- [16] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. 6
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 3
- [18] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 6
- [19] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, 2021. If you use this software, please cite it as below. 5
- [20] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014. 5
- [21] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2017. 2, 8

- [22] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. [2](#), [3](#), [8](#)
- [23] Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023. [3](#)
- [24] Kimin Lee, Hao Liu, Moonkyung Ryu, Olivia Watkins, Yuqing Du, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, and Shixiang Shane Gu. Aligning text-to-image models using human feedback. *arXiv preprint arXiv:2302.12192*, 2023. [3](#)
- [25] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. [1](#)
- [26] Ruibo Liu, Chenyan Jia, Ge Zhang, Ziyu Zhuang, Tony Liu, and Soroush Vosoughi. Second thoughts are best: Learning to re-align with human values from text edits. *Advances in Neural Information Processing Systems*, 35:181–196, 2022. [3](#)
- [27] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. [7](#)
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [5](#)
- [29] Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Motlaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. *arXiv preprint arXiv:2206.08916*, 2022. [2](#)
- [30] Arjun Mani, Nobline Yoo, Will Hinthorn, and Olga Russakovsky. Point and ask: Incorporating pointing into visual question answering. *arXiv preprint arXiv:2011.13681*, 2020. [6](#), [7](#)
- [31] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015. [3](#)
- [32] Varun K Nagaraja, Vlad I Morariu, and Larry S Davis. Modeling context between objects for referring expression understanding. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 792–807. Springer, 2016. [2](#), [3](#), [5](#)
- [33] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*, 2021. [3](#)
- [34] 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](#)
- [35] 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. [6](#)
- [36] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *ArXiv*, abs/2306, 2023. [2](#), [3](#)
- [37] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pages 2641–2649, 2015. [5](#), [6](#)
- [38] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. [1](#)
- [39] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. [2](#)
- [40] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. [2](#), [5](#)
- [41] Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021, 2020. [3](#)
- [42] Mengyang Sun, Wei Suo, Peng Wang, Yanning Zhang, and Qi Wu. A proposal-free one-stage framework for referring expression comprehension and generation via dense cross-attention. *IEEE Transactions on Multimedia*, 25:2446–2458, 2023. [7](#)
- [43] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018. [5](#)
- [44] Mikihiro Tanaka, Takayuki Itamochi, Kenichi Narioka, Ikuro Sato, Yoshitaka Ushiku, and Tatsuya Harada. Generating easy-to-understand referring expressions for target identifications. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5794–5803, 2019. [7](#), [8](#)
- [45] 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. [1](#)
- [46] 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. [5](#)

- [47] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015. 6
- [48] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, pages 23318–23340. PMLR, 2022. 2, 7
- [49] Peng Wang, Shijie Wang, Junyang Lin, Shuai Bai, Xiaohuan Zhou, Jingren Zhou, Xinggang Wang, and Chang Zhou. One-peace: Exploring one general representation model toward unlimited modalities. *arXiv preprint arXiv:2305.11172*, 2023. 7
- [50] 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. *arXiv preprint arXiv:2305.11175*, 2023. 7
- [51] Linhui Xiao, Xiaoshan Yang, Fang Peng, Ming Yan, Yaowei Wang, and Changsheng Xu. Clip-vg: Self-paced curriculum adapting of clip for visual grounding. *IEEE Transactions on Multimedia*, 2023. 6
- [52] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *arXiv preprint arXiv:2304.05977*, 2023. 3
- [53] Bin Yan, Yi Jiang, Jiannan Wu, Dong Wang, Ping Luo, Zehuan Yuan, and Huchuan Lu. Universal instance perception as object discovery and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15325–15336, 2023. 7
- [54] Li Yang, Yan Xu, Chunfeng Yuan, Wei Liu, Bing Li, and Weiming Hu. Improving visual grounding with visual-linguistic verification and iterative reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9499–9508, 2022. 6
- [55] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 69–85. Springer, 2016. 2, 3, 5, 6
- [56] Licheng Yu, Hao Tan, Mohit Bansal, and Tamara L. Berg. A joint speaker-listener-reinforcer model for referring expressions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3521–3529, 2017. 7, 8
- [57] Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601*, 2023. 1, 2
- [58] Chaoyang Zhu, Yiyi Zhou, Yunhang Shen, Gen Luo, Xingjia Pan, Mingbao Lin, Chao Chen, Liujuan Cao, Xiaoshuai Sun, and Rongrong Ji. Seqtr: A simple yet universal network for visual grounding. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 598–615. Springer, 2022. 6
- [59] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 1
- [60] Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4995–5004, 2016. 6, 7