arXiv:2603.18282v2 [cs.CV] 20 Mar 2026

CycleCap: Improving VLMs Captioning Performance via Self-Supervised Cycle Consistency Fine-Tuning

Marios Krestenitis[1,2] **Christos Tzelepis**[3] **Konstantinos Ioannidis**[2] **Stefanos Vrochidis**[2]
Ioannis Kompatsiaris[2] **Georgios Tzimiropoulos**[1] **Shaogang Gong**[1] **Ioannis Patras**[1]

[1]Queen Mary, University of London
[2]Centre for Research and Technology Hellas
[3]City St George's, University of London
Corresponding author: m.krestenitis@qmul.ac.uk

Abstract

Visual–Language Models (VLMs) have achieved remarkable progress in image captioning, visual question answering, and visual reasoning. Yet they remain prone to vision–language misalignment, often producing overly generic or hallucinated descriptions. Existing approaches address this via instruction tuning-requiring costly, large-scale annotated datasets or via complex test-time frameworks for caption refinement. In this work, we revisit image–text alignment through the lens of cycle consistency: given an image and a caption generated by an image-to-text model, the backward mapping through a text-to-image model should reconstruct an image that closely matches the original. In our setup, a VLM serves as the image-to-text component, while a pre-trained text-to-image model closes the loop by reconstructing the image from the generated caption. Building on this, we introduce CycleCap, a fine-tuning scheme to improve image captioning using Group Relative Policy Optimization (GRPO) with a reward based on the similarity between the original and reconstructed images, computed on-the-fly. Unlike previous work that uses cycle consistency loss for preference dataset construction, our method leverages cycle consistency directly as a self-supervised training signal. This enables the use of raw images alone, eliminating the need for curated image–text datasets, while steering the VLM to produce more accurate and grounded text descriptions. Applied to four VLMs ranging from 1B to 7B parameters, CycleCap yields consistent improvements across captioning and hallucination benchmarks, surpassing state-of-the-art methods that rely on supervised cycle consistency training.

***K*eywords** Image captioning · Cycle consistency · Self-supervised learning · Visual-language models

1 Introduction

Visual-Language Models (VLMs) are rapidly advancing as effective methods for understanding visual and textual information. Their ability to process both modalities has them demonstrating remarkable performance in several tasks, such as image captioning, visual question answering and visual reasoning [33, 2, 15, 57, 31]. Despite the impressive results, VLMs still face significant challenges in effectively aligning vision and language [59, 8, 37]. Images convey concrete structural and perceptual details that are often difficult to fully express in text, while text captures abstract semantics that may lack direct visual counterparts. Rethinking text-to-image and image-to-text generations as translation tasks, the mapping between images and text is inherently many-to-many rather than one-to-one due to this semantic ambiguity. Such misalignment between modalities can lead to overly generic descriptions or even hallucinated content that fails to accurately ground visual information.

Several approaches have been proposed to mitigate this issue, including instruction tuning strategies and reinforcement learning with human or automated feedback [62, 32, 61, 68]. However, these methods often require large-scale paired datasets, typically obtained through costly human annotation [6, 61]. Another line of work focuses on refining pre-acquired captions at inference time. To this end, multi-stage frameworks have been developed that integrate

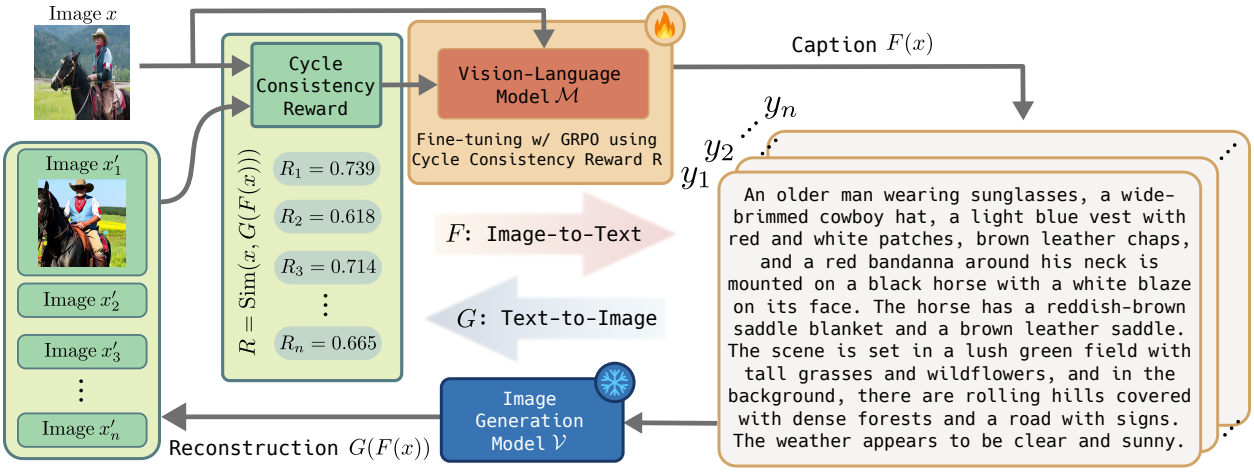

Figure 1: Overview of the CycleCap framework. The Visual-Language Model generates multiple captions $\{y_i\}_{i=1}^n$ for an image x. Each caption is used by a frozen Image Generation Model to reconstruct an image $x'_i = G(y_i)$, whilst the similarity between x and $G(y_i)$ is measured to obtain the cycle consistency reward R_i, $i = 1, \ldots, n$. These rewards guide fine-tuning of the Visual–Language Model via GRPO to encourage captions that better reflect visual content in a self-supervised manner.

VLMs, Large Language Models (LLMs), and supplementary detectors to enhance text descriptions and reduce hallucinations [46, 17, 47]. Despite bypassing additional training, these approaches rely on complex and expensive pipelines, while the extensive computation required at inference time limits their scalability.

In this paper, we revisit image-text alignment through the lens of cycle consistency [67]. As illustrated in Fig. 1, given an image $x \in \mathcal{X}$ and the image-to-text and text-to-image mappings $F\colon \mathcal{X} \to \mathcal{Y}$ and $G\colon \mathcal{Y} \to \mathcal{X}$, respectively, the translation between the two modalities is considered accurate when $G(F(x)) \approx x$, where the accuracy is measured by the similarity between x and $G(F(x))$. Our key idea is to use cycle consistency directly as a self-supervised training signal to improve image-text alignment. In this light, we introduce CycleCap, a cycle-consistent fine-tuning framework for VLMs that adapts Group Relative Policy Optimization (GRPO) [51] and incorporates a reward function that encourages greater similarity between x and $G(F(x))$. This steers the model to generate more accurate and detailed descriptions, improving image-text alignment and visual grounding.

Our work differs from previous approaches that have explored the use of cycle consistency primarily as an evaluation or post-process refinement tool for text descriptions [14, 26, 10, 16]. These methods have shown that higher cycle consistency correlates with more accurate image-to-text generations, but they typically treat it as an outcome of the model rather than a training objective. It also differs from studies that have leveraged cycle consistency loss to construct image-text preference datasets [58, 3], where ensembles of image-to-text and text-to-image models were used to rank image-text pairs for supervised fine-tuning of VLMs with Direct Preference Optimization (DPO) [49].

In contrast, our approach employs cycle consistency directly as a self-supervised learning signal during training. The proposed GRPO reinforcement learning scheme enforces cycle consistency on-the-fly, allowing the model to learn from image input alone and improve image-text alignment without relying on expensive preference datasets or post-processing refinements. We apply our method to fine-tune four novel VLMs of varying sizes, ranging from 1B to 7B parameters, and demonstrate consistent performance improvements across captioning and hallucination benchmarks. Our method surpasses current state-of-the-art approaches that leverage cycle consistency loss for supervised training. Our main contributions are summarized as follows:

- We introduce CycleCap, a novel cycle consistent fine-tuning framework for VLMs that uses the power of GRPO learning to acquire dense and accurate image-to-text outputs. Our approach leverages the idea of cycle consistency, allowing the use of only image data instead of requiring curated and expensive image-text tuning datasets.
- We deploy our method for four VLMs of varying size, ranging from 1B to 7B parameters, and demonstrate consistent performance improvements across captioning and hallucination benchmarks.
- We evaluate our method against state-of-the-art approaches that rely on cycle consistency to construct preference datasets for fine-tuning, and show that it surpasses these methods while operating in a fully self-supervised manner, without requiring costly and time-consuming dataset creation.

- We conduct extensive ablations showing that CycleCap is robust to different image-similarity metrics and text-to-image backbones, and that stronger reconstruction models or perceptual metrics yield higher gains, demonstrating the flexibility and scalability of the framework.

2 Related Work

2.1 Visual-Language Models

Visual–Language models (VLMs) integrate visual and textual understanding within a unified architecture, making them the foundation of multimodal understanding and reasoning. Early approaches, such as CLIP [48] and ALIGN [28], established large-scale image–text alignment through contrastive learning, while models like BLIP-2 [33] bridged pretrained visual encoders and language models using lightweight adapters. Since then, research has shifted toward building instruction-tuned and conversational foundation models that align multimodal behavior with natural language instructions. In this light, several multimodal architectures have been introduced, such as MiniGPT-4 [65], LLaVA [38], Qwen-VL [4], Gemini [53], and GPT4V [42], that extend pretrained LLMs (e.g., LLaMa [54] or Qwen [4]) with visual encoders and multimodal alignment training, enabling more grounded and context-aware responses. Parallel work explores post-training optimization methods—including reinforcement learning with human or automated feedback [13, 30, 7], or more recent DPO [49] and GRPO [51], to improve models' reasoning, alignment with human preferences, and instruction-following capabilities through robust post-training methods.

2.2 Detailed Image Captioning

Despite their scale and versatility, current VLMs still face challenges in achieving fine-grained visual grounding and rich descriptive detail. Producing detailed image descriptions is a challenging task that requires fine spatial reasoning, attribute recognition, and the ability to avoid hallucinated or generic content. A key factor behind this challenge lies in the large-scale image–text datasets [50, 1] used for pretraining, which often contain short or noisy text descriptions that limit the models' ability to learn rich and fine-grained visual–textual correspondences. To overcome this issue, researchers have explored enriching these datasets by re-captioning them with more detailed descriptions, as in ShareGPT4V [11] and IIW [21], or have focused on inference-time methods [46, 47] to improve caption detail at test time. Evaluating detailed captions is also challenging, since traditional reference-based metrics such as BLEU [45] and CIDEr [55] are sensitive to writing style and struggle with long, descriptive sentences. To address this, DetailCaps4870 [41], CompreCap [41], CapsBench [36], and CAPability [40] propose more comprehensive evaluation protocols, often combining LLM-based judging with multilevel alignment metrics. Although these advances improve both caption quality and its assessment, generating accurate and richly detailed captions without additional supervision remains an open challenge.

2.3 Cycle consistency

Cycle consistency, originally introduced by CycleGAN [67] for unpaired image-to-image translation, has been applied across diverse tasks [22, 56, 24, 60]. Recently, it has been extended to VLMs and multimodal systems. Recent works [26, 14, 10] have proposed using cycle consistency as a metric to evaluate the captioning performance of image-to-text models, based on the idea that a faithful text description should enable accurate reconstruction of the original image. Beyond evaluation, Li et al. [34] leveraged cycle consistency for vision–language learning by designing a specialized architecture trained with token-level reconstruction loss to exploit unpaired image-text data. Although the approach was introduced to reduce reliance on large paired datasets, it still required approximately 3M of paired image–text examples for initialization. Moreover, the cycle consistency objective was tightly integrated into the model architecture, making the method dependent on the specialized design rather than serving as a training scheme. On the other hand, CyclePrompt [16] employed a cyclical framework to refine captions offline, generating images from the captions of original images with DALL-E 3 [9] and using GPT-4V [42] to improve those captions based on differences between the original and the reconstructed images.

Building on this idea, Wang et al. [58] introduced the RICO, an iterative caption refinement process that uses FLUX.1-dev [29] to reconstruct images from captions and GPT-4o [44] to refine captions based on discrepancies between the original and reconstructed images. This iterative process is used to construct a preference dataset of caption pairs (initial versus refined) to train the captioning model RICO-Flash. Similarly, Bahng et al. [3] proposed the CyclePref framework, a cycle-based preference construction pipeline. They used an ensemble of 11 image-to-text models (0.5B–40B parameters) to generate caption variations and Stable Diffusion 3 (SD3) [18] for image reconstruction. Generated captions were ranked by comparing the image reconstructions to the original inputs, forming the CyclePrefDB-I2T preference dataset to improve captioning performance.

Despite these advancements, previous works either treat cycle consistency as an evaluation characteristic or use it as a ranking criterion to build costly preference datasets. In contrast, our approach, CycleCap, leverages cycle consistency directly as a self-supervision signal to improve VLM performance and enhance captioning capabilities without requiring annotated image-text pairs or relying on multiple large-scale components or external APIs such as GPT-4.

3 CycleCap Framework

In this section, we present the CycleCap framework for fine-tuning VLMs to improve image captioning quality. As shown in Fig. 1, CycleCap is a simple yet effective approach built on two complementary components:

- An image-to-text VLM model $\mathcal{M}$ denoted as $F\colon \mathcal{X} \to \mathcal{Y}$ that maps images $x \in \mathcal{X}$ to captions $y \in \mathcal{Y}$.
- An image generation model $\mathcal{V}$ that performs the reverse mapping $G\colon \mathcal{Y} \to \mathcal{X}$.

Our objective is to fine-tune $\mathcal{M}$ to generate textual descriptions that are better grounded to the visual content of an image.

Cycle Consistency Reward Given an image $x \in \mathcal{X}$, the VLM model $\mathcal{M}$ produces a textual description $y = F(x) \in \mathcal{Y}$. To evaluate how well the text reflects the image content, we measure how accurately it can reconstruct the original image through $G(y) \in \mathcal{X}$. To this end, we define the cycle consistency reward as

$$R = \text{Sim}(x, G(F(x))), \tag{1}$$

where $\text{Sim}(\cdot,\cdot)$ denotes a similarity metric between the original and reconstructed image. In our experiments, we measure this similarity with DreamSim [20], which assesses the perceptual and semantic correspondence between images. This formulation allows for estimating the quality of the generated text descriptions without requiring reference captions. A caption that enables accurate reconstruction of the input image implicitly captures visual semantics and structure, and thus (1) serves as a self-supervised measure of description quality that we later use to guide model fine-tuning.

Fine-tuning with GRPO We adopt the cycle consistency reward (1) to fine-tune $\mathcal{M}$ with GRPO. The objective is to maximize R, encouraging the model to generate captions that better reflect the visual content. During training, only the parameters of $\mathcal{M}$ are updated, while the text-to-image model $\mathcal{V}$ remains frozen.

As illustrated in Figure 1, for each image $x \in \mathcal{X}$, the VLM generates a group of n candidate captions $y_i \in \mathcal{Y}$, $i = 1, \ldots, n$. Each caption is fed to the image generator $\mathcal{V}$ to generate the reconstructed image $x'_i = G(y_i) \in \mathcal{X}$. Each caption y_i is assigned with a similarity score $R_i = \text{Sim}(x, G(y_i))$. Then, the relative advantage of each caption within the group is calculated as

$$A_i = \frac{R_i - \bar{R}}{s_R}, \tag{2}$$

where $\bar{R}$ and s_R are the mean and standard deviation of the rewards $\{R_1, \ldots, R_n\}$ within the group, respectively. The GRPO loss is then defined as

$$\mathcal{L}_{\text{GRPO}} = -\,\mathbb{E}\left[\frac{1}{n}\sum_{i=1}^{n} \min\left(\rho_i(\theta)A_i,\ \text{clip}(\rho_i(\theta),\, 1-\varepsilon,\, 1+\varepsilon)\, A_i\right)\right] + \beta\, D_{\text{KL}}\left(\pi_\theta \,\|\, \pi_{\text{ref}}\right), \tag{3}$$

where $\rho_i = \frac{\pi_\theta(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)}$ is the likelihood ratio between the updated policy π_θ and the previous policy $\pi_{\theta_{\text{old}}}$, ε is the clipping threshold and β the weight of the KL regularization term used to constrain the updated policy π_θ to stay close to a frozen reference policy π_{ref}. This encourages the model to increase the likelihood of captions that yield higher reconstruction rewards, leading to more accurate text descriptions, while preventing divergence from the original model behavior.

4 Experiments

In this section, we conduct extensive comparisons with widely used VLMs and benchmarks to evaluate the performance gains of our approach.

4.1 Implementation Details

We deploy CycleCap to fine-tune a set of vision–language models of varying sizes, ranging from 1B to 7B parameters – specifically, InternVL3-1B [66], Qwen2-VL-2B [57], Qwen2.5-VL-3B [5], and Qwen2-VL-7B [57]. Training is performed on the COCO 2014 train split [35], which contains approximately 83,000 everyday scene images. All models are fine-tuned for one epoch with a learning rate of 10^{-5}, a global batch size of 64, and the number of GRPO rollouts (caption generations per image) is set to 8. For efficiency, we apply LoRA adaptation [25] with rank 64. During training, each model is prompted to produce a detailed description of the image; the exact prompt template is provided in Fig. A1 in the Appendix. For the backward text-to-image mapping $G\colon \mathcal{Y} \to \mathcal{X}$, we employ Stable Diffusion 3 (SD3) [18] similarly to CyclePref [3] with its default parameters. Moreover, for a fair comparison with RICO [58], we additionally employ FLUX.1-dev [29] text-to-image generator. To mitigate variance in the reward signal introduced by stochastic image generation, we fix the random seed per image sample. All models were trained using 2 A100 GPUs.

4.2 Evaluation Benchmarks

We study the effect of CycleCap fine-tuning on model performance, mainly on the captioning task. To this end, we employ a set of widely used multimodal benchmarks that assess detailed image captions in terms of completeness and correctness. Specifically, we use CompreCap [41], which contains 560 annotated images and measures caption quality from a structured scene-graph perspective by evaluating object-level coverage as well as the accuracy of object attributes and inter-object relations. Moreover, we use CAPability [40], a comprehensive benchmark that evaluates captions across multiple aspects of an image (i.e., object category, number, and color, spatial relations, scene, camera angle, OCR, style and character identification), each containing approx. 1,000 human-annotated images, and checks whether captions correctly and thoroughly describe those elements compared to ground-truth annotations using a GPT-based evaluator. We also use CapsBench [36], which contains 200 images and tests whether generated captions provide sufficient visual grounding by prompting an LLM to answer "yes/no" questions about the image, based solely on the caption. Furthermore, we examine the model performance in terms of hallucinations as a critical aspect that affects caption accuracy. To this end, we use MMHal [52], which contains 96 images paired with challenging image-related questions and scores the responses using an LLM on a 0–6 scale, where 0 denotes fully hallucinated and 6 indicates highly accurate and well-reasoned answers.

For the CAPability benchmark, we use the official prompt to generate image captions. For CompreCap and CapsBench, we extract captions from the evaluated models using the official CompreCap prompt. In the cases of GPT-based evaluators, we employ GPT-4o-mini [43]. For the smaller benchmarks, CapsBench and MMHal, the evaluation is repeated three times to ensure consistency.

4.3 Comparison with Baseline Models

Table 1 summarizes the improvements in captioning performance achieved by the proposed CycleCap fine-tuning compared to the corresponding baseline models. In these experiments, we use the standard SD3 [18] as an image generator. As shown, CycleCap consistently improves caption quality across all benchmarks and model sizes. On CompreCap [41], it yields 2–3% improvements on the Unified Score. Improvements on CAPability [40], a more challenging benchmark that holistically examines the model's captioning abilities across multiple evaluation dimensions, further confirm that CycleCap leads to captions that are more informative, semantically complete, and thoroughly aligned with diverse aspects of the visual content. Similarly, CycleCap leads to improvements of over 2% in most cases, when evaluated on CapsBench [36]. Accordingly, gains in MMHal [52] indicate that the generated descriptions are more informative while containing less hallucinations. Notably, these gains are observed even for the larger model Qwen2-VL-7B, suggesting that CycleCap complements existing large-scale multimodal pretraining rather than merely compensating for smaller model architectures. These results confirm that using the cyclic consistency reward benefits the model's ability to generate detailed and faithful image descriptions.

Furthermore, in Fig. 2a, we show per-image win-rates between fine-tuned and baseline models on CompreCap, indicating how often CycleCap outperforms the baseline under the same metric. We show that CycleCap exceeds the baseline in $> 50\%$ of cases across all metrics, indicating consistent improvements at the benchmark level rather than gains concentrated on a small subset of examples. While the relative win-rates are particularly pronounced for smaller models (1B–3B), especially in object coverage and attribute scores, improvements are consistent across all model scales and benchmarks (see Tab. 1).

Table 1: Comparison of the proposed CycleCap with different baseline models (InternVL3-1B [66], Qwen2-VL-2B [57], Qwen2.5-VL-3B [5], Qwen2-VL-7B [57]) on captioning (CompreCap [41], CAPability [40], CapsBench [36]) and hallucination (MMHal [52]) benchmarks.

Model	CompreCap [41]				CAPability [40] [0, 100]	CapsBench [36] [0, 100]	MMHal [52] [0, 6]
	Object Coverage [0, 100]	Attribute Score [0, 5]	Relation Score [0, 5]	Unified Score [0, 100]			
InternVL3-1B [66]							
Baseline	73.35	2.75	2.83	60.24	69.37	71.93	3.29
CycleCap (Ours)	**77.35**	**2.89**	**2.87**	**62.49**	**70.89**	**73.37**	**3.36**
Qwen2-VL-2B [57]							
Baseline	73.10	2.72	2.76	59.35	69.20	69.70	3.63
CycleCap (Ours)	**76.94**	**2.85**	**2.86**	**62.09**	**70.96**	**72.11**	**3.71**
Qwen2.5-VL-3B [5]							
Baseline	72.30	2.72	2.76	59.21	68.70	69.52	3.78
CycleCap (Ours)	**76.92**	**2.87**	**2.88**	**62.42**	**71.45**	**73.56**	**4.09**
Qwen2-VL-7B [57]							
Baseline	77.86	2.84	2.80	61.73	70.47	74.17	3.85
CycleCap (Ours)	**79.32**	**2.91**	**2.86**	**63.06**	**72.95**	**76.38**	**4.02**

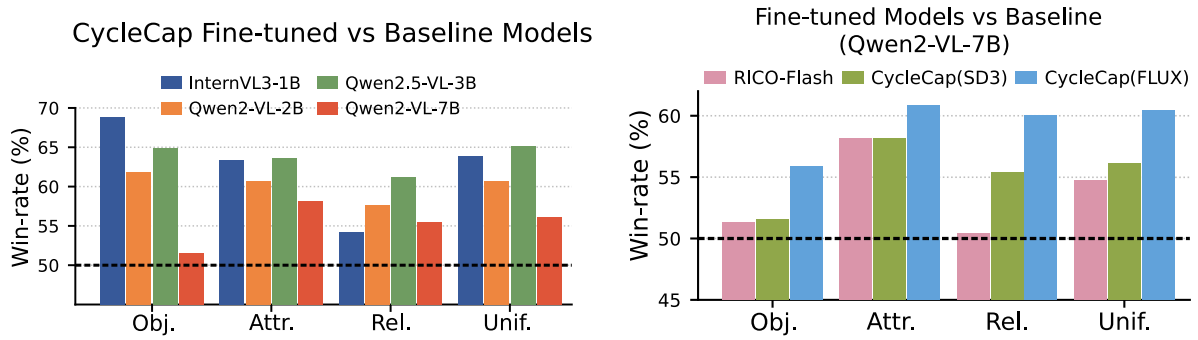

(a) (b)

Figure 2: Win-rates (%) of (a) CycleCap fine-tuned models versus the corresponding baseline and (b) Qwen2-VL-7B [57] fine-tuned with different methods versus the baseline for CompreCap [41] benchmark.

4.4 Comparison with State-of-the-Art (SOTA)

In this section, we compare CycleCap with SOTA methods that leverage cycle consistency to improve captioning performance. More specifically, we compare with RICO-Flash [58], a version of Qwen2-VL-7B [57] fine-tuned on a 10K preference dataset derived through the RICO framework, an iterative caption refinement process that uses GPT-4o [27] feedback by comparing the original with one reconstructed from its caption using FLUX.1-dev [29]. Moreover, we compare with CyclePref [3], which constructs a larger 398K image-text preference dataset (called CyclePrefDB-I2T) using an ensemble of 11 image-to-text models, ranging from 0.5 to 40B parameters. Caption preference is ranked based on the similarity between the input image and an image reconstructed from its caption using Stable Diffusion 3 (SD3) [18]. Then, CyclePrefDB-I2T is used to fine-tune Qwen-VL-7B-Chat [4] for the captioning task with DPO. In our experiments, we follow the setup of [3] and fine-tune Qwen2-VL-7B with DPO on the CyclePrefDB-I2T dataset.

In Tab. 2 we summarize the comparisons with SOTA. The proposed CycleCap consistently achieves higher performance across all benchmarks compared to supervised fine-tuning methods. Notably, CycleCap also improves performance on the CAPability [40] benchmark, a more challenging evaluation where prior methods show limited or even negative gains over the baseline. Overall, CycleCap with SD3 improves the SOTA by a margin comparable to the gap previously observed between the baseline and prior SOTA approaches [58, 3]. It surpasses RICO-Flash [58] and CyclePref [3] in caption completeness and factual grounding, demonstrating that direct cyclic self-supervision provides a more effective training signal than preference-based refinement approaches. On top of that, using FLUX.1-dev for text-to-image generation, similarly to RICO-Flash [58], leads to additional performance gains. This trend is also reflected in Fig. 2b, which reports per-image win-rates against the baseline Qwen2-VL-7B for RICO-Flash and CycleCap (with SD3 and

Table 2: Comparison of CycleCap with state-of-the-art cycle-consistency-based methods on CompreCap [41], CAPability [40], CapsBench [36], and MMHal [52] benchmarks, using Qwen2-VL-7B [57] as baseline. For CompreCap we report here only Unified Score, the analytic table is included in the Appendix.

Method	T2I Model	CompreCap	CAPability	CapsBench	MMHal
Baseline (Qwen2-VL-7B)	—	61.73	70.47	74.17	3.85
CyclePref [3]	SD3	62.03	70.59	74.27	3.95
RICO-Flash [58]	FLUX.1-dev	62.93	68.83	75.30	3.92
CycleCap (Ours)	SD3	63.06	72.95	76.38	4.02
CycleCap (Ours)	FLUX.1-dev	**63.64**	**73.73**	**77.25**	**4.02**

Table 3: Comparison of CycleCap with state-of-the-art methods on the CAPability [40] benchmark, all using Qwen2-VL-7B [57] as baseline. Bold and underlined values indicate the best and second best results, respectively.

Metric	Baseline (Qwen2-VL-7B [57])	CyclePref [3] (SD3)	RICO-Flash [58] (FLUX.1-dev)	CycleCap (Ours) (SD3)	CycleCap (Ours) (FLUX.1-dev)
Obj. Category	73.52	72.61	70.97	<u>74.17</u>	**77.57**
Obj. Number	60.67	60.00	57.85	<u>65.12</u>	**65.74**
Obj. Color	89.67	89.33	88.64	**90.36**	<u>89.94</u>
Spatial Relation	69.11	67.50	68.24	**71.29**	<u>71.08</u>
Scene	77.07	76.36	75.56	<u>78.60</u>	**79.06**
Camera Angle	38.00	39.24	36.90	**40.20**	<u>38.42</u>
OCR	96.99	97.27	<u>97.81</u>	**97.88**	97.61
Style	74.37	76.23	75.35	<u>80.20</u>	**81.50**
Character Ident.	54.83	56.77	48.14	<u>58.73</u>	**62.67**
Average	70.47	70.59	68.83	<u>72.95</u>	**73.73**

FLUX.1-dev respectively) under the CompreCap metrics. CycleCap variants achieve consistently higher win-rates, indicating more consistent improvements relative to the baseline compared to RICO-Flash.

Moreover, in Tab. 3 we compare CycleCap to SOTA across each evaluation dimension of the CAPability [40] benchmark, providing a more fine-grained assessment. CycleCap achieves the best overall performance, with the FLUX.1-dev variant obtaining the highest average score. While prior cycle-based approaches show limited or inconsistent gains over the baseline, CycleCap consistently improves performance across all dimensions. Notably, our method demonstrates substantial improvements particularly in object category and number recognition, scene understanding, style description, and character identification – dimensions that require precise semantic understanding and strong visual grounding. The consistent gains across SD3 and FLUX.1-dev variants indicate that CycleCap enhances caption quality by generating richer descriptions that capture diverse visual attributes and relational aspects of the image.

Overall, Tables 2 and 3 highlight the effectiveness of CycleCap. Notably, our method requires only an image generator, whereas competing SOTA methods require additional substantially heavier external components apart from the image generator itself. Specifically, CyclePref [3] is built upon an ensemble of large and costly image-to-text models (up to 40B parameters), while RICO-Flash [58] depends on iterative refinement using the costly GPT-4o. In both cases, computational cost scales with performance, as improvements to the preference-based supervision are inherently tied either to increasing the number and scale of ensemble models or to performing additional caption refinement iterations through GPT-4o API calls. By contrast, CycleCap replaces static preference-based supervision with a continuous, image-grounded feedback that allows the model to directly optimize for image-text alignment using only raw images rather than imitating preference judgments and without relying on external API calls or multiple heavyweight components. Using CycleCap's GRPO optimization scheme encourages progressive improvements through exploration of multiple caption hypotheses, allowing the model to refine its outputs beyond the limitations of static preference pairs and without scaling external supervision complexity.

CycleCap on top of RICO-Flash We additionally apply CycleCap (with SD3 [18]) on top of RICO-Flash [58], to examine the gains that yield from our method when deployed on an already strengthened model for captioning task through cycle consistency–ranked preference pairs. Results in Tab. 4 show that CycleCap on top of RICO-Flash

Table 4: Performance of CycleCap applied on top of RICO-Flash [58] across captioning and hallucination benchmarks, with Qwen2-VL-7B [57] as the baseline.

Method	CompreCap				CAPability	CapsBench	MMHal
	Obj	Attr	Rel	Uni			
Baseline (Qwen2-VL-7B)	77.86	2.84	2.80	61.73	70.47	74.17	3.85
RICO-Flash [58]	79.09	2.93	2.83	62.93	68.83	75.30	3.92
CycleCap (Ours)	79.32	2.91	2.86	63.06	**73.73**	76.38	**4.02**
RICO-Flash [58] + CycleCap (Ours)	**80.59**	**2.95**	**2.88**	**63.85**	73.49	**77.72**	4.01

Table 5: Effect of different similarity metrics (LPIPS [64], CLIP [48], DreamSim [20]) on CycleCap performance for fine-tuning the baseline Qwen2-VL-2B [57].

Method	CompreCap				CapsBench	MMHal
	Obj	Attr	Rel	Uni		
Baseline (Qwen2-VL-2B)	73.10	2.72	2.76	59.35	69.70	3.63
CycleCap w/ LPIPS	74.21	2.68	2.80	59.70	64.05	3.49
CycleCap w/ CLIP	73.96	2.78	2.83	60.58	71.16	**3.94**
CycleCap w/ DreamSim	**76.94**	**2.85**	**2.86**	**62.09**	**72.11**	3.71

achieves improved performance. Note that CycleCap alone yields gains on the baseline model that surpass those of RICO-Flash. Applying CycleCap on top of RICO-Flash further improves performance, leading to the strongest performance across CompreCap and CapsBench. On CAPability, the combined approach substantially improves over RICO-Flash, which struggled to surpass the baseline, while achieving gains competitive with CycleCap alone. For MMHal, all CycleCap-based variants report similar scores, suggesting that performance on this benchmark is close to a saturation point for this training setup. Overall, the results indicate that our method is complementary to RICO-Flash approach, and that it scales effectively with captioning models boosting further their performance.

4.5 Ablation Studies

To further assess the performance and flexibility of CycleCap, we conduct ablation studies along two key factors that influence fine-tuning effectiveness: (a) the image similarity metric between the input image x and its reconstruction $G(F(x))$ used for the cycle consistency reward, and (b) the text-to-image model $\mathcal{V}$ used for the backward mapping G.

4.5.1 Cycle consistency metrics

In this section, we study the effect of different image similarity metrics on CycleCap. Besides DreamSim [20], we evaluate LPIPS [64] and CLIP-based embeddings [48] to measure similarity between original and reconstructed images. ollowing the setup described in Sect. 4.3, we fine-tune Qwen2-VL-2B [57] for each configuration.

As shown in Tab. 5, the DreamSim-based variant achieves the highest scores across nearly all benchmarks, indicating that perceptual metrics capturing both low- and high-level visual features provide a more effective cyclic supervision signal. LPIPS, which primarily focuses on structural similarity, yields smaller gains – mainly on CompreCap, which emphasizes on caption completeness. CLIP similarity leads to improvements in hallucination robustness but underperforms in overall caption quality compared to DreamSim, which offers the most balanced results. These findings indicate that while CycleCap benefits from different similarity metrics, the choice of image-space metric remains a critical factor influencing its effectiveness, with more robust metrics leading to higher performance.

4.5.2 Text-to-image models

In this section, we study the effect of the employed text-to-image model $\mathcal{V}$ on CycleCap's performance. To this end, we replace SD3 [18] with FLUX.1-schnell [29] for the reconstruction step and fine-tune Qwen2-VL-2B [57] under the same training setup. The results are reported in Tab. 6. Both variants yield clear improvements over the baseline model across all captioning benchmarks, confirming that the cyclic supervision scheme is robust to different reconstruction backbones. The performance of the two cases is comparable, SD3 provides slightly stronger gains on

Table 6: Effect of different text-to-image models (SD3 [18] and FLUX.1-schnell [29]) on the proposed CycleCap performance for fine-tuning baseline Qwen2-VL-2B [57].

Method	CompreCap				CapsBench	MMHal
	Obj	Attr	Rel	Uni		
Baseline (Qwen2-VL-2B)	73.10	2.72	2.76	59.35	69.70	3.63
CycleCap w/ SD3	**76.94**	2.85	2.86	62.09	72.11	**3.71**
CycleCap w/ FLUX.1-schnell	75.96	**2.87**	**2.89**	**62.14**	**72.18**	3.65

Baseline:

The image depicts a scene from a movie or a dramatic event set in an urban environment at night. The central figure is a man dressed in a red suit, standing with his arms outstretched. His expression is intense and somewhat menacing, with a slight smile on his face. The background is filled with a crowd of people, many of whom appear to be cheering or celebrating. The scene is illuminated by streetlights and emergency vehicle lights, creating a dramatic and somewhat chaotic atmosphere. The overall mood of the image is one of tension and excitement.

CycleCap:

The image depicts a dramatic scene set in a city at night. The central figure is a man dressed in a red suit, standing in the middle of a street. He has **long, wavy hair** and is extending his arms wide, as if he is addressing a crowd or **performing a dramatic gesture**. The **man's expression is intense and somewhat menacing**, suggesting he might be a character of significance, possibly a villain or a charismatic leader. The background is filled with a crowd of people, many of whom appear to be cheering or celebrating. The scene is illuminated by various lights, including street lamps and emergency vehicle lights, which add to the dramatic atmosphere. **The buildings in the background are tall and dark, with some lights visible through the windows, indicating a bustling urban environment**. There is a sense of chaos and excitement in the air, with **smoke or fog** adding to the dramatic effect. The overall mood of the image is intense and dramatic, suggesting a moment of high stakes or a pivotal scene in a movie or a live event.

Baseline:

The image depicts an underwater scene featuring a person snorkeling near a vibrant coral reef. The snorkeler is wearing a full-body wetsuit and fins, and is equipped with a snorkel and a mask. The coral reef is lush and colorful, with a variety of coral species in shades of purple, pink, and orange. Several small, brightly colored fish, including orange and red species, are swimming around the coral. The water is clear, allowing for a detailed view of the underwater environment. The sunlight filters through the water, creating a serene and picturesque scene.

CycleCap:

The image depicts an underwater scene teeming with vibrant marine life and coral formations. The water is clear, allowing for a detailed view of the underwater environment. The sunlight filters through the water, creating a serene and tranquil atmosphere. In the foreground, there is a person snorkeling, wearing a wetsuit and fins. **The snorkeler is swimming towards the right side of the image**, exploring the coral reef. The person appears to be **a woman with long hair tied back**, and she is wearing a snorkel mask and fins. The coral reef is lush and diverse, featuring a variety of coral species. **The corals are predominantly purple and pink, with some areas showing a mix of white and brown. The coral formations are dense and intricate, with many branches and tentacles extending from the reef**. Scattered throughout the reef are numerous small fish, primarily in shades of orange and red. These fish are swimming around the coral, adding to the dynamic and lively nature of the scene. The fish appear to be of different species, contributing to the biodiversity of the reef. **In the background, the water extends into the distance, with gentle waves and ripples on the surface.** The light from the sun creates a gradient of light and shadow, enhancing the depth and dimension of the underwater world. Overall, the image captures the beauty and complexity of an underwater coral reef, showcasing the rich marine life and the serene environment that supports it. The snorkeler's presence adds a human element to the scene, highlighting the connection between humans and the natural world.

Figure 3: Qualitative comparison of captions generated by the baseline Qwen2-VL-7B [57] and the proposed CycleCap on samples from the CapsBench [36] dataset. Newly added information is indicated in **bold**. Our method produces more detailed, organized, and accurate descriptions compared to the baseline outputs.

MMHal, while FLUX.1-schnell achieves slightly higher performance on CompreCap and CapsBench. Overall, these results demonstrate the flexibility of CycleCap with different text-to-image models, allowing practitioners to trade off computational efficiency, reconstruction speed, and generation style according to available resources.

4.6 Qualitative Results

In this section, we provide qualitative results of the CycleCap method. Figure 3 presents captions of sample images from the CapsBench dataset, generated with the baseline Qwen2-VL-7B [57] model and the CycleCap fine-tuned version. The CycleCap captions are noticeably denser and more accurate, providing richer descriptions of spatial layout, object and character attributes, and environmental context. They also exhibit improved structure, often organizing the scene into coherent sentences covering the foreground, background, and overall atmosphere. These examples demonstrate that CycleCap improves the model's ability to capture fine-grained visual details and produce more comprehensive captions.

Additionally, Fig. 4 presents qualitative comparisons between CycleCap (using SD3 [18]) and the SOTA methods listed in Tab. 2. For each case, we generate an image caption and then reconstruct an image from that description. For a fair comparison, we use FLUX.1-schnell [29] to generate all reconstructions, since it has not been part of any of the comparative frameworks, and apply the same generation seed across all methods. The illustrations show that CycleCap's captions produce reconstructions that more closely resemble the input image in both detail and structural fidelity.

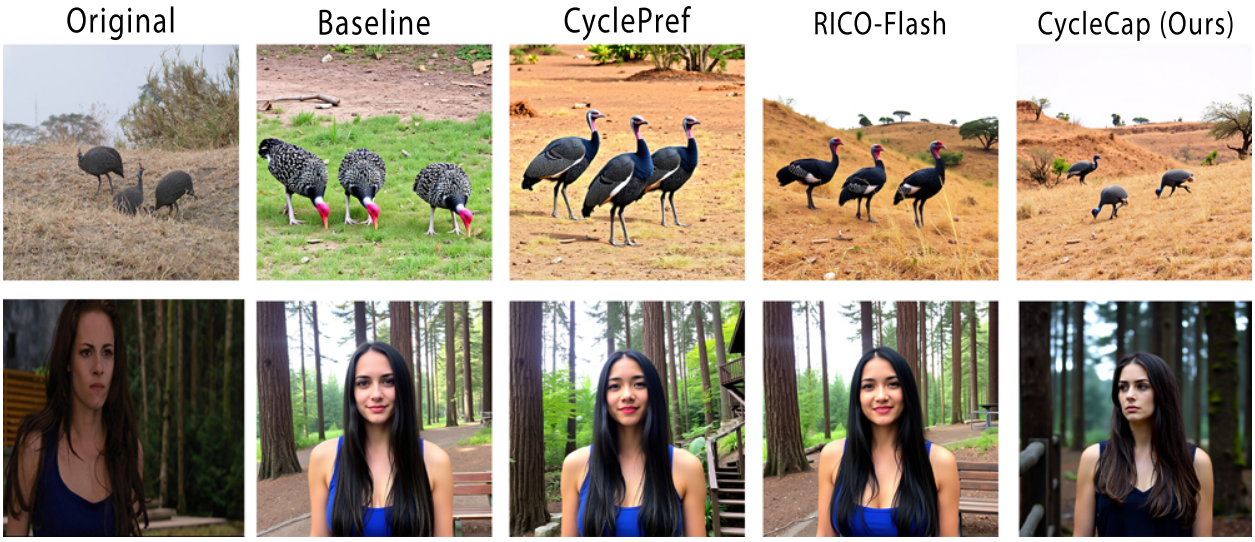

Figure 4: Qualitative comparison of image reconstructions with captions generated by SOTA and our method (CycleCap) deployed for Qwen2-VL-7B. The visualization shows that the model fine-tuned with CycleCap captures more structural details and object attributes, leading to reconstructions closer to the original image. The generated captions are provided in the Appendix.

5 Conclusion

In this work, we introduce CycleCap, a simple yet effective fine-tuning framework for improving the captioning performance of VLMs. Our approach builds on the idea of cycle consistency, where an accurate image-to-text generation should enable a faithful text-to-image reconstruction. CycleCap uses this principle as a direct self-supervised training signal. We build a cycle consistency reward based on the similarity between the input and the reconstructed image, and combined with a GRPO-based optimization scheme it allows the model to explore multiple caption hypotheses and progressively reinforce those that yield higher image similarity. This provides a continuous, image-grounded learning signal using only raw image data, bypassing the need for costly annotated image–text datasets. Across four VLMs ranging from 1B to 7B parameters, CycleCap delivers consistent improvements on detailed captioning and hallucination benchmarks. We further show that CycleCap outperforms state-of-the-art methods that rely on cycle-consistency–derived preference datasets for fine-tuning. Unlike these approaches, which depend on the quality and scale of static preference pairs, CycleCap offers an efficient self-supervised approach that learns directly from images. Ablation studies demonstrate the flexibility of our framework – it benefits from stronger perceptual similarity metrics and higher-fidelity generators, suggesting that the method will naturally improve further as generative models advance. We believe CycleCap opens up new opportunities for self-supervised visual–textual learning and offers a strong foundation for future research in multimodal generation. Currently, our approach is deployed for the image–text–image cycle and relies solely on image-space similarity. In future work, we will explore the extension of the training signal with cross-domain measures that, in addition to image-space similarity, assess image-text alignment across GRPO rollouts.

Appendix

A Additional implementation details

In Tab. A1 below we present the full set of training parameters used to fine-tune the models with the proposed CycleCap framework. Training time ranged from 270 (1B) to 430 (7B) GPU-hours on 2×A100. Additionally, in Fig. A1 we report the prompt used in training to generate image captions.

Table A1: CycleCap training parameters.

Category	Parameters	
Training	Batch size: 64 Scheduler: Linear	Learning rate: 10^{-5} Epochs: 1
LoRA	Rank r: 64 Target modules: all linear projection layers	Dropout: 0.05 Vision tower: frozen
GRPO	KL weight β: 0.04 Generations n: 8	Clip value ε: 0.02
Optimization	Optimizer: AdamW Gradient checkpointing	Precision: bfloat16

```
You are given an image. Your task is to write a detailed and accurate caption that could be used by
a text-to-image model to recreate the image faithfully.
Include information about:
  · The main objects, their colors, shapes, sizes, materials, poses, orientations, and visible parts.
  · Scene layout, spatial relationships, and sense of depth or scale.
  · Lighting conditions, camera angle, and overall environment.
  · Human or animal appearance (pose, clothing, expression, etc.).
  · Background details and textures.
  · If relevant, the artistic or visual style (photo, painting, digital art, cartoon, etc.).
  · Any visible text, symbols, or color palette characteristics.
Respond only with the caption, without additional commentary.
```

Figure A1: The designed prompt of the CycleCap framework used to generate image captions during training.

B Additional experiments

B.1 Evaluation on visual understanding and reasoning benchmarks

Our main objective is to use the CycleCap fine-tuning framework to improve captioning performance. Hence, enhancing performance on broader visual-language understanding, reasoning, or complex VQA tasks lies outside the scope of this work. Nevertheless, for the sake of completeness, we examine how CycleCap affects these abilities by evaluating the models presented in Tab. 1 in the main paper, on a set of widely used benchmarks that include a variety of visual inputs, ranging from natural images to maps and charts, and that measure visual-language capabilities in high-level perception and reasoning tasks. Tab. A2 summarizes the results. We observe that the CycleCap fine-tuned models exhibit comparable performance to their baseline counterparts. This indicates that CycleCap does not harm the generic multimodal capabilities of the models, and can even lead to small improvements, despite being trained solely on the captioning task.

Table A2: Comparison of CycleCap with baseline models on visual-language understanding and reasoning benchmarks. The evaluated models are the same as those reported in Tab. 1 in the main paper.

Model	MME_{sum} [19]	$\text{MMBench}_{\text{test}}$ [39]	MMStar [12]	MMMU_{val} [63]	Hall-Bench [23]
			InternVL3-1B [66]		
Baseline	1873.19	70.51	**52.31**	**40.22**	47.21
CycleCap (Ours)	**1892.90**	**71.41**	52.00	39.33	**49.31**
			Qwen2-VL-2B [57]		
Baseline	1866.25	71.08	43.34	40.44	50.68
CycleCap (Ours)	**1872.50**	**71.97**	**43.70**	**40.89**	**51.10**
			Qwen2.5-VL-3B [5]		
Baseline	2147.22	77.57	**56.01**	45.89	**57.51**
CycleCap (Ours)	**2148.82**	**78.75**	55.04	**46.22**	56.04
			Qwen2-VL-7B [57]		
Baseline	**2294.35**	78.30	57.28	50.44	57.93
CycleCap (Ours)	2293.37	**78.81**	**58.18**	**51.11**	**58.88**

B.2 Analysis of Caption Length

In Tab. A3 we report the average length (in tokens) of the captions generated by each model shown in Tab. 1 in the main paper, on the Comprecap [41] benchmark. The results show that CycleCap fine-tuned models produce caption lengths comparable to baseline models for all cases. This suggests that the improvements in captioning performance are not solely associated with increased verbosity, but rather reflect qualitative improvements in the generated descriptions, such as better structure and stronger semantic alignment with the visual content.

Table A3: Average caption length (in tokens) on the CompreCap [41] benchmark.

	InternVL3-1B [66]		Qwen2-VL-2B [57]		Qwen2.5-VL-3B [5]		Qwen2-VL-7B [57]	
	Baseline	CycleCap	Baseline	CycleCap	Baseline	CycleCap	Baseline	CycleCap
Avg. Length	152.64	179.80	209.39	202.73	176.06	198.73	384.56	385.23

B.3 Effect of GRPO's number of generations

We analyze the sensitivity of CycleCap to the number of generations, n, in GRPO [51], which corresponds to the number of generated captions per image during training to compute the relative advantage in (3). To this end, we fine-tune the Qwen2-VL-7B [57] model with CycleCap for $n = 2, 4$ (in addition to $n = 8$) using SD3 [18], following a similar process to the one described in Section 4.1 of the main paper. Tab. A4 reports CycleCap's performance for different values of n on captioning benchmarks.

Across all configurations, CycleCap consistently improves over the baseline. Notably, even with $n = 2$ or $n = 4$, CycleCap surpasses SOTA approaches in most cases (see Tab. 2 in the main paper), emphasizing the gains that arise from the GRPO-based training objective. Increasing the number of caption generations to $n = 8$ achieves the best overall results across the evaluated benchmarks. In Fig. A2, we illustrate the relationship between CycleCap performance and training cost for different values of n. As shown, increasing n scales approximately linearly with training cost, with larger numbers of generated captions leading to improved performance.

B.4 Controlled comparison with DPO on CyclePrefDB-I2T

In Section 4.4 in the main paper, we compare CycleCap with prior cycle-consistency-based methods, including CyclePref [3], which applies DPO [49] using the CyclePrefDB-I2T [3] preference dataset. Since the models reported in Tab. 2 in the main paper are trained on different datasets, we additionally evaluate both approaches using the same training data for a more controlled comparison.

Specifically, we fine-tune Qwen2-VL-7B [57] with CycleCap using the images from CyclePrefDB-I2T (approximately 7k images). In CycleCap training, a group of $n = 8$ candidate captions is generated for each image during GRPO optimization. Since CyclePrefDB-I2T [3] originally contains about 55 preference pairs per image, we randomly sample

Table A4: Effect of the number of GRPO generations n (captions generated per image) on CycleCap performance on captioning benchmarks CompreCap (Unified Score) [41], CAPability [40], and CapsBench [36].

Method	Generations n	CompreCap	CAPability	CapsBench
Baseline (Qwen2-VL-7B)	—	61.73	70.47	74.17
CycleCap	2	62.57	72.40	74.24
CycleCap	4	62.58	72.17	76.22
CycleCap	8	**63.06**	**72.95**	**76.38**

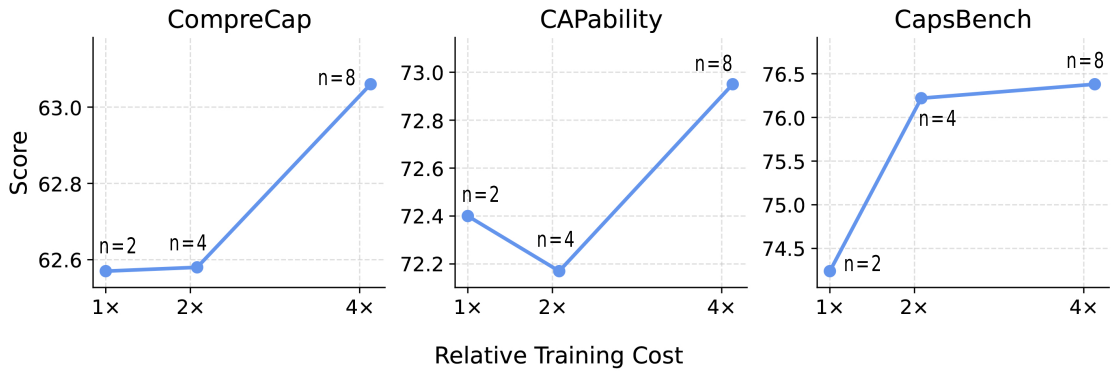

Figure A2: Effect of the number of GRPO generations n on CycleCap performance vs. relative training cost. We report performance on CompreCap (Unified Score) [41], CAPability [40], and CapsBench [36] benchmarks. Relative training cost is normalized with respect to $n = 2$ ($1\times$).

8 preference pairs per image and use this subset for DPO fine-tuning, so that both methods receive a comparable number of supervision signals per image. We fine-tune each model for one epoch using identical optimization settings (Tab. A1) and report the results in Tab. A5.

DPO fine-tuning with the sub-sampled CyclePrefDB-I2T [3] pairs leads to modest improvements over the baseline model. CycleCap, however, consistently achieves stronger gains across the adopted evaluation benchmarks. More specifically, by generating candidate captions during training and optimizing them through cycle consistency rewards, the proposed method provides a stronger learning signal than relying solely on static preference pairs from the dataset.

Table A5: Controlled comparison between DPO and CycleCap using the sub-sampled CyclePrefDB-I2T [3] training set with Qwen2-VL-7B [57] as the base model. Results are reported on CompreCap (Unified Score) [41], CAPability [40], CapsBench [36], and MMHal [52] benchmarks.

Method	CompreCap	CAPability	CapsBench	MMHal
Baseline (Qwen2-VL-7B)	61.73	70.47	74.17	3.85
DPO	61.78	71.11	74.41	3.95
CycleCap (Ours)	**62.33**	**71.89**	**74.47**	**4.01**

B.5 Detailed results of Table 2

Tab. A6 provides the detailed evaluation results corresponding to Tab. 2 in the main paper.

C Additional qualitative comparisons

In Fig. A3 we present additional qualitative comparisons between captions generated by the baseline Qwen2-VL-7B [57] and its CycleCap fine-tuned variant. Moreover, in Fig. A4 and A5 we present the generated captions of each method that correspond to the qualitative comparison of Fig. 4 in the main paper. As shown, CycleCap produces more detailed and accurate captions compared to SOTA methods.

Table A6: Detailed results corresponding to Tab. 2 in the main paper, comparing the proposed CycleCap with SOTA cycle-consistency-based approaches.

Method	T2I Model	CompreCap				CAPability	CapsBench	MMHal
		Obj	Attr	Rel	Uni			
Baseline (Qwen2-VL-7B)	—	77.86	2.84	2.80	61.73	70.47	74.17	3.85
CyclePref [3]	SD3	77.29	2.85	2.85	62.03	70.59	74.27	3.95
RICO-Flash [58]	FLUX.1-dev	79.09	2.93	2.83	62.93	68.83	75.30	3.92
CycleCap (Ours)	SD3	79.32	2.91	2.86	63.06	72.95	76.38	4.02
CycleCap (Ours)	FLUX.1-dev	**79.67**	**2.93**	**2.90**	**63.64**	**73.73**	**77.25**	**4.02**

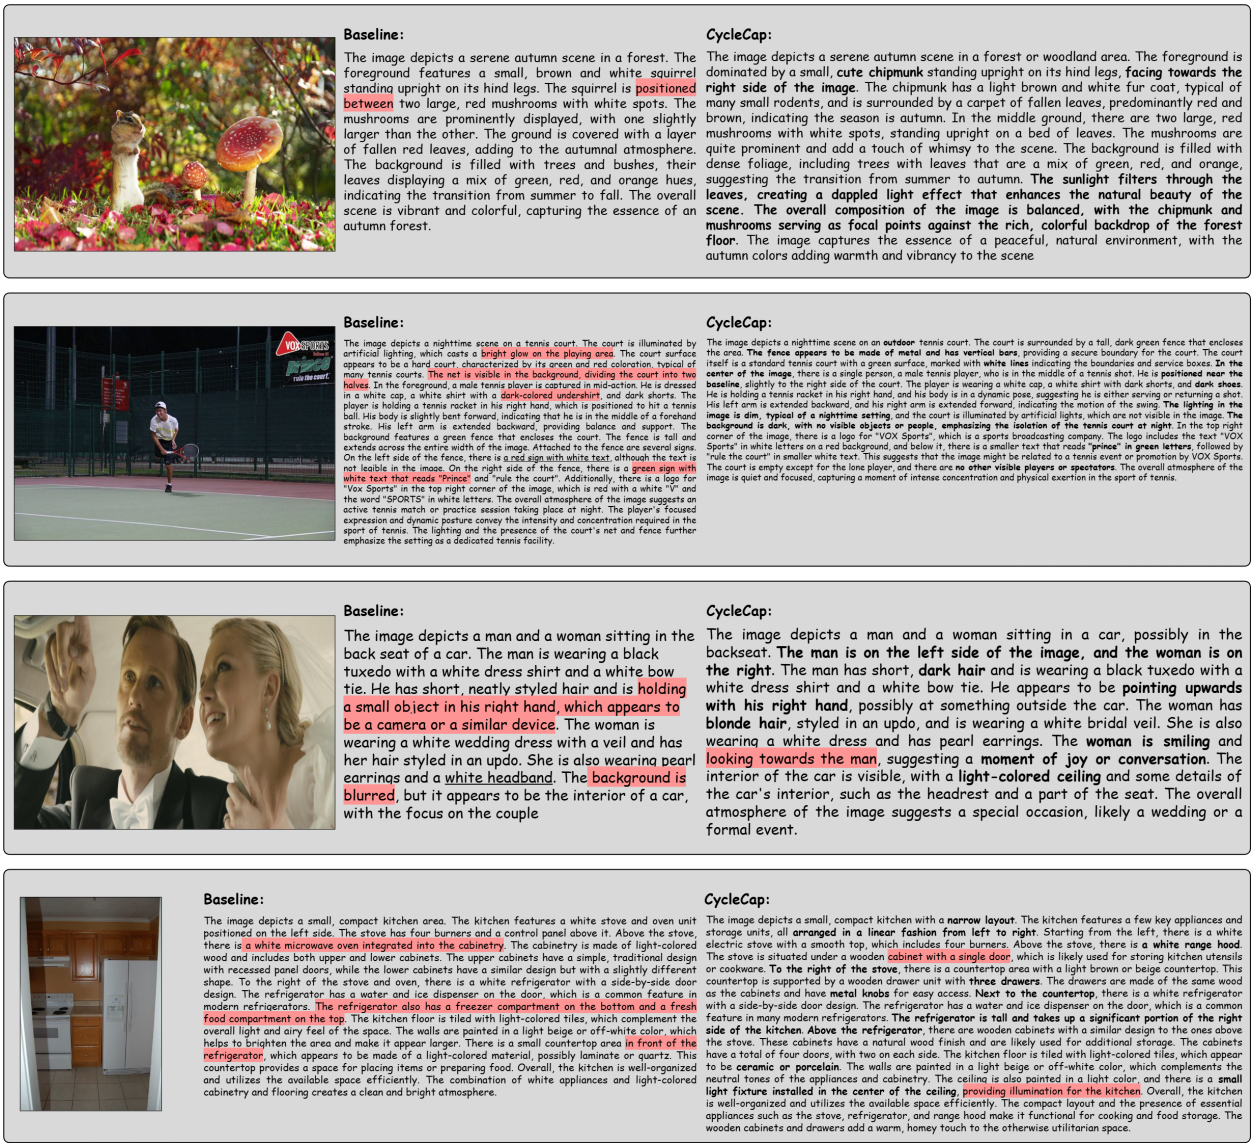

Figure A3: Additional qualitative comparisons between captions generated by the baseline Qwen2-VL-7B [57] model and the CycleCap fine-tuned version, for image samples from the CapsBench [36] and CompreCap [41] benchmarks. Newly added information is shown in **bold**, content present in the baseline but not preserved in the CycleCap caption is underlined, and hallucinations are highlighted in pink. Overall, CycleCap tends to produce more structured and grounded descriptions, while it adds details related to object attributes, spatial layout, and broader scene context.

Baseline: Three guinea fowl are foraging in a grassy area.
CyclePref: The image depicts three guinea fowl standing on a dry, grassy terrain. The background includes some sparse vegetation and a few trees. The guinea fowl have a distinctive appearance with their dark feathers and white speckles.

RICO-Flash: Three guinea fowl birds with dark feathers and white spots, standing on a dry grassy hillside with sparse vegetation and a few trees in the background. The sky is overcast.
CycleCap (Ours): The image depicts a natural, outdoor scene featuring three guinea fowl foraging on a dry, grassy hillside. The background includes sparse vegetation and a few scattered trees, with a hazy sky above. The guinea fowl are characterized by their dark, spotted feathers and distinctive blue heads. The terrain appears to be part of a savanna or grassland ecosystem.

Figure A4: Generated captions corresponding to the qualitative reconstruction comparison shown in the first row of Fig. 4 in the main paper.

Baseline: The image shows a person with long, dark hair wearing a blue tank top, standing in a forested area with tall trees and a wooden bench in the background. The setting appears to be outdoors, possibly in a park or a natural reserve.

CyclePref: The image shows a person with long, dark hair wearing a blue tank top, standing in a forested area with tall trees and green foliage. The background includes a wooden structure and a staircase.
RICO-Flash: A woman with long dark hair, wearing a blue tank top, stands in a forested area with tall trees and a wooden bench in the background. The setting appears to be a mix of natural and man-made elements.
CycleCap (Ours): The image shows a woman with long, dark brown hair standing in a forested area. She is wearing a dark blue sleeveless top and appears to be in a contemplative or serious mood. The background features tall trees with moss-covered trunks and a wooden fence or railing on the left side. The setting seems to be a natural, wooded environment, possibly in a park or forest. The lighting is dim, suggesting it might be early morning or late afternoon.

Figure A5: Generated captions corresponding to the qualitative reconstruction comparison shown in the second row of Fig. 4 in the main paper.

References

[1] Agrawal, H., Desai, K., Wang, Y., Chen, X., Jain, R., Johnson, M., Batra, D., Parikh, D., Lee, S., Anderson, P.: Nocaps: Novel object captioning at scale. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 8948–8957 (2019)

[2] Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. Advances in neural information processing systems **35**, 23716–23736 (2022)

[3] Bahng, H., Chan, C., Durand, F., Isola, P.: Cycle consistency as reward: Learning image-text alignment without human preferences. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22934–22946 (2025)

[4] Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., et al.: Qwen technical report. arXiv preprint arXiv:2309.16609 (2023)

[5] Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2. 5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)

[6] Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al.: Training a helpful and harmless assistant with reinforcement learning from human feedback. arXiv preprint arXiv:2204.05862 (2022)

[7] Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., et al.: Constitutional ai: Harmlessness from ai feedback. arXiv preprint arXiv:2212.08073 (2022)

[8] Bai, Z., Wang, P., Xiao, T., He, T., Han, Z., Zhang, Z., Shou, M.Z.: Hallucination of multimodal large language models: A survey. arXiv preprint arXiv:2404.18930 (2024)

[9] Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al.: Improving image generation with better captions. Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf **2**(3), 8 (2023)

[10] Chan, C., Bahng, H., Durand, F., Isola, P.: On the cycle consistency of image-text mappings (2025)

[11] Chen, L., Li, J., Dong, X., Zhang, P., He, C., Wang, J., Zhao, F., Lin, D.: Sharegpt4v: Improving large multi-modal models with better captions. In: European Conference on Computer Vision. pp. 370–387. Springer (2024)

[12] Chen, L., Li, J., Dong, X., Zhang, P., Zang, Y., Chen, Z., Duan, H., Wang, J., Qiao, Y., Lin, D., et al.: Are we on the right way for evaluating large vision-language models? Advances in Neural Information Processing Systems **37**, 27056–27087 (2024)

[13] Christiano, P.F., Leike, J., Brown, T., Martic, M., Legg, S., Amodei, D.: Deep reinforcement learning from human preferences. Advances in neural information processing systems **30** (2017)

[14] Cui, T., Bai, J., Wang, G.H., Chen, Q.G., Xu, Z., Luo, W., Zhang, K., Shi, Y.: Evaluating image caption via cycle-consistent text-to-image generation. arXiv preprint arXiv:2501.03567 (2025)

[15] Dai, W., Li, J., Li, D., Tiong, A., Zhao, J., Wang, W., Li, B., Fung, P.N., Hoi, S.: Instructblip: Towards general-purpose vision-language models with instruction tuning. Advances in neural information processing systems **36**, 49250–49267 (2023)

[16] Diesendruck, M., Lin, J., Imani, S., Mahalingam, G., Xu, M., Zhao, J.: Learning how to ask: Cycle-consistency refines prompts in multimodal foundation models. arXiv preprint arXiv:2402.08756 (2024)

[17] Dong, H., Li, J., Wu, B., Wang, J., Zhang, Y., Guo, H.: Benchmarking and improving detail image caption. arXiv preprint arXiv:2405.19092 (2024)

[18] Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning (2024)

[19] Fu, C., Chen, P., Shen, Y., Qin, Y., Zhang, M., Lin, X., Yang, J., Zheng, X., Li, K., Sun, X., et al.: Mme: A comprehensive evaluation benchmark for multimodal large language models. arXiv preprint arXiv:2306.13394 (2023)

[20] Fu, S., Tamir, N., Sundaram, S., Chai, L., Zhang, R., Dekel, T., Isola, P.: Dreamsim: Learning new dimensions of human visual similarity using synthetic data. arXiv preprint arXiv:2306.09344 (2023)

[21] Garg, R., Burns, A., Karagol-Ayan, B., Bitton, Y., Montgomery, C., Onoe, Y., Bunner, A., Krishna, R., Baldridge, J.M., Soricut, R.: Imageinwords: Unlocking hyper-detailed image descriptions. In: Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing. pp. 93–127 (2024)

[22] Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 270–279 (2017)

[23] Guan, T., Liu, F., Wu, X., Xian, R., Li, Z., Liu, X., Wang, X., Chen, L., Huang, F., Yacoob, Y., et al.: Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14375–14385 (2024)

[24] Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. In: International conference on machine learning. pp. 1989–1998. Pmlr (2018)

[25] Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. ICLR **1**(2), 3 (2022)

[26] Huang, J.H., Zhu, H., Shen, Y., Rudinac, S., Kanoulas, E.: Image2text2image: A novel framework for label-free evaluation of image-to-text generation with text-to-image diffusion models. In: International Conference on Multimedia Modeling. pp. 413–427. Springer (2025)

[27] Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. arXiv preprint arXiv:2410.21276 (2024)

[28] Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International conference on machine learning. pp. 4904–4916. PMLR (2021)

[29] Labs, B.F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., Kulal, S., Lacey, K., Levi, Y., Li, C., Lorenz, D., Müller, J., Podell, D., Rombach, R., Saini, H., Sauer, A., Smith, L.: Flux.1 kontext: Flow matching for in-context image generation and editing in latent space (2025), `https://arxiv.org/abs/2506.15742`

[30] Lee, H., Phatale, S., Mansoor, H., Lu, K.R., Mesnard, T., Ferret, J., Bishop, C., Hall, E., Carbune, V., Rastogi, A.: Rlaif: Scaling reinforcement learning from human feedback with ai feedback (2023)

[31] Li, B., Zhang, Y., Guo, D., Zhang, R., Li, F., Zhang, H., Zhang, K., Zhang, P., Li, Y., Liu, Z., et al.: Llava-onevision: Easy visual task transfer. arXiv preprint arXiv:2408.03326 (2024)

[32] Li, D., Jiang, B., Huang, L., Beigi, A., Zhao, C., Tan, Z., Bhattacharjee, A., Jiang, Y., Chen, C., Wu, T., et al.: From generation to judgment: Opportunities and challenges of llm-as-a-judge. In: Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing. pp. 2757–2791 (2025)

[33] Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In: International conference on machine learning. pp. 19730–19742. PMLR (2023)

[34] Li, T., Bhardwaj, S., Tian, Y., Zhang, H., Barber, J., Katabi, D., Lajoie, G., Chang, H., Krishnan, D.: Leveraging unpaired data for vision-language generative models via cycle consistency. arXiv preprint arXiv:2310.03734 (2023)

[35] Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)

[36] Liu, B., Akhgari, E., Visheratin, A., Kamko, A., Xu, L., Shrirao, S., Lambert, C., Souza, J., Doshi, S., Li, D.: Playground v3: Improving text-to-image alignment with deep-fusion large language models. arXiv preprint arXiv:2409.10695 (2024)

[37] Liu, H., Xue, W., Chen, Y., Chen, D., Zhao, X., Wang, K., Hou, L., Li, R., Peng, W.: A survey on hallucination in large vision-language models. arXiv preprint arXiv:2402.00253 (2024)

[38] Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. Advances in neural information processing systems **36**, 34892–34916 (2023)

[39] Liu, Y., Duan, H., Zhang, Y., Li, B., Zhang, S., Zhao, W., Yuan, Y., Wang, J., He, C., Liu, Z., et al.: Mmbench: Is your multi-modal model an all-around player? In: European conference on computer vision. pp. 216–233. Springer (2024)

[40] Liu, Z., Xie, C.W., Wen, B., Yu, F., Chen, J., Li, P., Zhang, B., Yang, N., Li, Y., Gao, Z., et al.: Capability: A comprehensive visual caption benchmark for evaluating both correctness and thoroughness. arXiv preprint arXiv:2502.14914 (2025)

[41] Lu, F., Wu, W., Zheng, K., Ma, S., Gong, B., Liu, J., Zhai, W., Cao, Y., Shen, Y., Zha, Z.J.: Benchmarking large vision-language models via directed scene graph for comprehensive image captioning. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 19618–19627 (2025)

[42] OpenAI: Gpt-4v (vision) system card. Tech. Rep. System Card, OpenAI (Sep 2023), https://cdn.openai.com/papers/GPTV_System_Card.pdf

[43] OpenAI: Gpt-4o mini. https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/ (2024), large language model

[44] OpenAI: Hello GPT-4o. https://openai.com/index/hello-gpt-4o/ (May 2024), accessed: 2025-11-12

[45] Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting of the Association for Computational Linguistics. pp. 311–318 (2002)

[46] Peng, R., He, H., Wei, Y., Wen, Y., Hu, D.: Patch matters: Training-free fine-grained image caption enhancement via local perception. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 3963–3973 (2025)

[47] Pi, R., Zhang, J., Zhang, J., Pan, R., Chen, Z., Zhang, T.: Image textualization: An automatic framework for creating accurate and detailed image descriptions. arXiv preprint arXiv:2406.07502 (2024)

[48] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PmLR (2021)

[49] Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. Advances in neural information processing systems **36**, 53728–53741 (2023)

[50] Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. Advances in neural information processing systems **35**, 25278–25294 (2022)

[51] Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al.: Deepseekmath: Pushing the limits of mathematical reasoning in open language models. arXiv preprint arXiv:2402.03300 (2024)

[52] Sun, Z., Shen, S., Cao, S., Liu, H., Li, C., Shen, Y., Gan, C., Gui, L., Wang, Y.X., Yang, Y., et al.: Aligning large multimodal models with factually augmented rlhf. In: Findings of the Association for Computational Linguistics: ACL 2024. pp. 13088–13110 (2024)

[53] Team, G., Anil, R., Borgeaud, S., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., Millican, K., et al.: Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805 (2023)

[54] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023)

[55] Vedantam, R., Lawrence Zitnick, C., Parikh, D.: Cider: Consensus-based image description evaluation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4566–4575 (2015)

[56] Wang, N., Song, Y., Ma, C., Zhou, W., Liu, W., Li, H.: Unsupervised deep tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1308–1317 (2019)

[57] Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., Chen, K., Liu, X., Wang, J., Ge, W., et al.: Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. arXiv preprint arXiv:2409.12191 (2024)

[58] Wang, Y., Cai, Y., Ren, S., Yang, S., Yao, L., Liu, Y., Zhang, Y., Wan, P., Sun, X.: RICO: Improving accuracy and completeness in image recaptioning via visual reconstruction. In: Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing. pp. 21785–21804 (2025), https://aclanthology.org/2025.emnlp-main.1105/

[59] Yarom, M., Bitton, Y., Changpinyo, S., Aharoni, R., Herzig, J., Lang, O., Ofek, E., Szpektor, I.: What you see is what you read? improving text-image alignment evaluation. Advances in Neural Information Processing Systems **36**, 1601–1619 (2023)

[60] Yi, Z., Zhang, H., Tan, P., Gong, M.: Dualgan: Unsupervised dual learning for image-to-image translation. In: Proceedings of the IEEE international conference on computer vision. pp. 2849–2857 (2017)

[61] Yu, T., Yao, Y., Zhang, H., He, T., Han, Y., Cui, G., Hu, J., Liu, Z., Zheng, H.T., Sun, M., et al.: Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13807–13816 (2024)

[62] Yuan, W., Pang, R.Y., Cho, K., Li, X., Sukhbaatar, S., Xu, J., Weston, J.E.: Self-rewarding language models. In: Forty-first International Conference on Machine Learning (2024)

[63] Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., et al.: Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9556–9567 (2024)

[64] Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)

[65] Zhu, D., Chen, J., Shen, X., Li, X., Elhoseiny, M.: Minigpt-4: Enhancing vision-language understanding with advanced large language models. arXiv preprint arXiv:2304.10592 (2023)

[66] Zhu, J., Wang, W., Chen, Z., Liu, Z., Ye, S., Gu, L., Tian, H., Duan, Y., Su, W., Shao, J., et al.: Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. arXiv preprint arXiv:2504.10479 (2025)

[67] Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)

[68] Ziegler, D.M., Stiennon, N., Wu, J., Brown, T.B., Radford, A., Amodei, D., Christiano, P., Irving, G.: Fine-tuning language models from human preferences. arXiv preprint arXiv:1909.08593 (2019)