

DreamLIP: Language-Image Pre-training with Long Captions

Kecheng Zheng^{*,1,2}, Yifei Zhang^{*,3}, Wei Wu⁴, Fan Lu⁴,
Shuailei Ma⁶, Xin Jin⁵, Wei Chen¹, and Yujun Shen²

¹ State Key Lab of CAD&CG, Zhejiang University

² Ant Group ³ Shanghai Jiao Tong University

⁴ University of Science and Technology of China

⁵ Eastern Institute of Technology ⁶ Northeastern University, China

{zkechengzk, xiaomabufei, shenyujun0302}@gmail.com,

qidouxiong619@sjtu.edu.cn, {lufan,wuvy}@mail.ustc.edu.cn, jinxin@eias.ac.cn

Abstract. Language-image pre-training largely relies on how precisely and thoroughly a text describes its paired image. In practice, however, the contents of an image can be so rich that well describing them requires lengthy captions (*e.g.*, with 10 sentences), which are usually missing in existing datasets. Consequently, there are currently no clear evidences on *whether and how language-image pre-training could benefit from long captions*. To figure this out, we first re-caption 30M images with detailed descriptions using a pre-trained Multi-modality Large Language Model (MLLM), and then study the usage of the resulting captions under a contrastive learning framework. We observe that, each sentence within a long caption is very likely to describe the image partially (*e.g.*, an object). Motivated by this, we propose to dynamically sample sub-captions from the text label to construct multiple positive pairs, and introduce a grouping loss to match the embeddings of each sub-caption with its corresponding local image patches in a self-supervised manner. Experimental results on a wide range of downstream tasks demonstrate the consistent superiority of our method, termed **DreamLIP**, over previous alternatives, highlighting its fine-grained representational capacity. It is noteworthy that, on the tasks of image-text retrieval and semantic segmentation, our model trained with 30M image-text pairs achieves on par or even better performance than CLIP trained with 400M pairs. Project page is available at <https://zyf0619sjtu.github.io/dream-lip>.

Keywords: Language-image pre-training· Long caption· Multi-modal learning·

1 Introduction

Language-image pre-training largely relies on how precisely and thoroughly a text describes its paired image. Existing image text pairing datasets typically

* Equal contribution

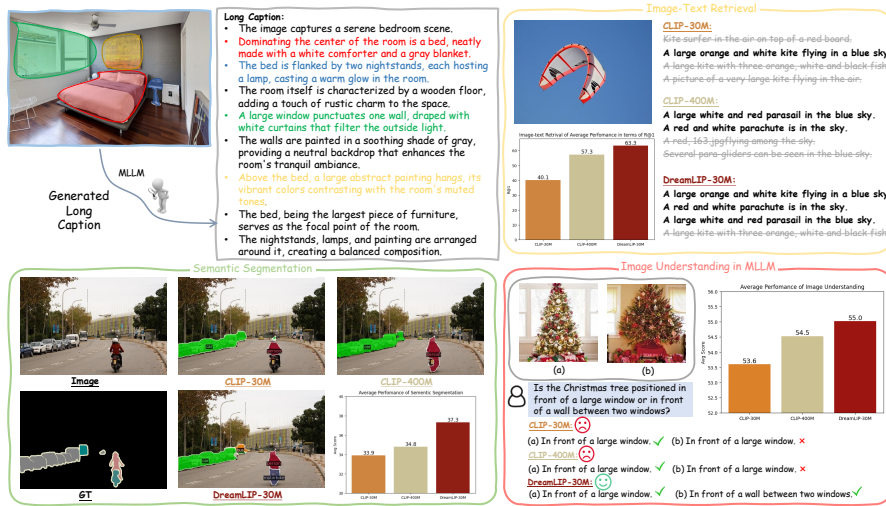


Fig. 1: Richness of an image’s content often necessitates long captions for adequate description, with each sentence likely conveying a fragment of the image’s entirety. Thanks to the long captions, our **DreamLIP** trained with 30M image-text pairs achieves on par or even better performance than CLIP trained with 400M pairs on the tasks of image-text retrieval, semantic segmentation, and image understanding in MLLM.

include short captions, which barely scratches the surface of the intricate real image. In practice, the contents of an image can be so rich that describing them well requires long captions (*e.g.*, with 10 sentences). These sentences are included in the long captions of corresponding image which can typically delineate various local regions of the image. Thus, long captions can unleash the potential of a real-world image and enrich semantic learning of language-image models. This multifaceted relationship between long captions and images presents an untapped reservoir of semantic richness.

Although long captions have enormous potential, there is a shortage of million-scale datasets to evaluate it in vision-language pre-training. One of the most straightforward approaches is: utilizing pre-trained Multi-modality Large Language Models (MLLM) to generate long captions. Actually, recent works [1, 12, 46, 48, 52, 57] also have explored the use of synthetic data to improve vision-language modeling. Specifically, RLEG [57] exploits DALL-E 2 [41] to generate feature embedding online for learning effective vision-language representation. StableRep [46] shows that training on synthetic images generated by Stable Diffusion [43] and their multi-positive contrastive learning method can match the performance of training on real images. LaCLIP [12] uses large language models to rewrite the captions of the real images as data augmentation. On the other hand, some methods [32, 52] utilize the pre-trained visual-language model to generate concise synthetic *short captions* that focus on the image content to address the noise in the original web data. Although existing methods

have attempted to improve the representation ability of multi-modal models using short captions, there are currently no clear evidences on *whether and how language-image pre-training could benefit from long captions*.

To figure this out, we first re-caption 30M images with detailed descriptions using a pre-trained MLLM, and then study the usage of the resulting captions under a contrastive learning framework. As shown in Fig. 1, we observe that each sentence within a long caption is very likely to describe the image partially (*e.g.*, bed). Motivated by this, we propose to dynamically sample sub-captions from the text label to construct multiple positive pairs. This innovative process allows us to forge an image-text dataset with long captions, wherein we employ a multi-positive loss framework to intricately intertwine sub-captions with their corresponding images, crafting a rich tapestry of aligned modalities. Taking it a step further, our approach also delves into learning subcaption-specific groupings of image patches. By applying a subcaption-specific grouping loss, we achieve a fine-grained alignment that meticulously pairs local image patches with their respective textual embeddings across different sub-captions. This nuanced pairing brings forth a better semantic alignment between the two modalities.

Experimental results across a diverse array of downstream tasks consistently demonstrate the superiority of our DreamLIP. These results underscore its exceptional fine-grained representational abilities when compared to previous alternatives. Notably, in the tasks of image-text retrieval, semantic segmentation, and image understanding in MLLM, our model trained on 30 million image-text pairs datasets achieves performance that is comparable to, or even surpasses, that of CLIP, despite the latter being trained on a dataset consisting of 400 million pairs.

2 Related Work

Vision-Language Pre-training. Some recent works like CLIP [40] and ALIGN [19] have shown that contrastive vision-language pre-training can provide rich and general representations for numerous downstream tasks. However, these methods only apply contrastive loss between the entire image and its caption, ignoring the local alignment between text and image. To this end, FILIP [53] and PyramidCLIP [16] modified the original contrastive loss to act between text tokens and image patches. Furthermore, HILCLIP [17] captures the hierarchical nature of high-level and fine-grained semantics conveyed in images and texts through hierarchy-aware attention. UniCLIP [25] instead integrated the contrastive loss of both inter-domain and intra-domain pairs into a single universal space. SoftCLIP [15] relaxed the strict one-to-one constraint with a soft target to enable image-text pairs to have some local similarities and model many-to-many relationships between the two modalities. Instead, in order to get a more robust representation, CLOOB [14] used Hopfield network to regulate the convergence of the learned representations, and FIBER [10] introduced dual modality encoders to obtain better fused multi-modal features. In addition, FLIP [28] and MaskCLIP [9] used image patch masking as an effective method

for vision-language pre-training. Finally, since the image-text pairs crawled from the internet are replete with substantial noise, UniCL [51] proposed to integrate image-label pairs from supervised datasets, while CoCa [55], BLIP [26] proposed to denoise text captions by regenerating them.

Improving Vision-Language Pre-training with Synthetic Data. Synthetic data has been employed to improve models’ performance on many computer vision tasks ranging from semantic segmentation [5, 42], object detection [20] and image classification [56]. Recent works [12, 46, 52, 57] have explored the effect of synthetic data to improve vision-language pre-training. RLEG [57] exploits DALL-E 2 [41] to generate feature embedding online for learning effective vision-language representation. StableRep [46] shows that training on synthetic images generated by Stable Diffusion [43] and their multi-positive contrastive learning method can match the performance of training on real images. Instead, SynthCLIP [18] creates synthetic image-text data using text-to-image generative networks and large language models to eliminate the impact of misalignment, long-tail distribution, and harmful contents in real data. On the other hand, ALIP [52] utilizes the OFA model to generate correct synthetic captions that focus on the image content to address the noise in the original web data. Similarly, LaCLIP [12] uses large language models to rewrite the captions of the real images as data augmentation. The closest to our work is the concurrent new works [24, 32] that enhance visual-language representation learning by utilizing multi-modal large language models to generate captions for each image. Although existing methods have attempted to improve the representation ability of multi-modality models using short captions, they have yet to investigate the employment of *long captions*, which are rich in image details and offer untapped informational potential.

3 Method

3.1 Preliminary

CLIP [40] mainly consists of two components: an image encoder E_I and a text encoder E_T , which can project image and text into the same feature embedding space. Specifically, the images I_1, I_2, \dots, I_N and the corresponding raw short texts T_1, T_2, \dots, T_N are sampled from the training dataset during each training step. The features of image I_i and text T_i are extracted using dedicated encoders and normalization functions: $\mathbf{v}_i = E_I(I_i; \boldsymbol{\theta}), \mathbf{t}_i = E_T(T_i; \boldsymbol{\beta})$. The image and text features are used to compute the InfoNCE [38] loss, where the paired image-text forms the positive pairs, and the unpaired ones are treated as negative samples. In this way, the text-to-vision loss can be obtained by:

$$\mathcal{L}^{t2v} = \sum_{i=1}^N \log \frac{\exp(\cos \mathbf{v}_i, \mathbf{t}_i / \tau)}{\sum_{j=1}^N \exp(\cos \mathbf{v}_j, \mathbf{t}_i / \tau)}, \quad (1)$$

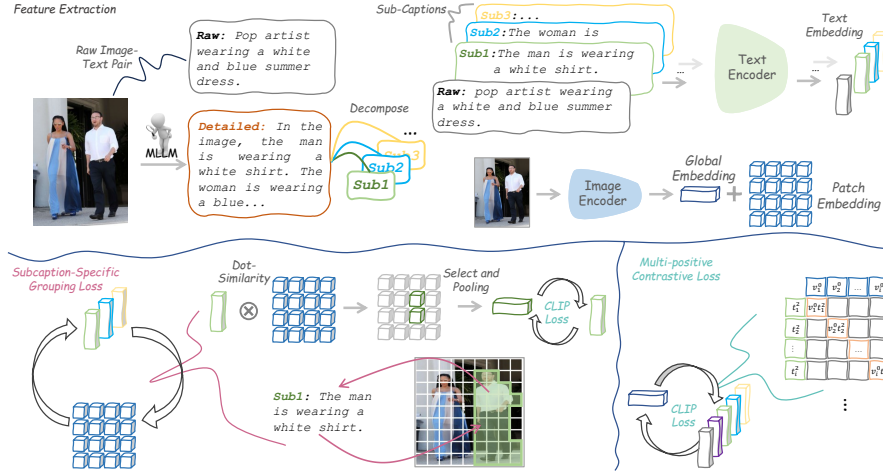


Fig. 2: Illustration of DreamLIP. Firstly, we dynamically sample sub-captions from the text label to construct multiple positive pairs. Then, a global multi-positive contrastive loss is used to align text embeddings of sub-captions and global image embedding. Meanwhile, we introduce a grouping loss to match the embeddings of each sub-caption with its corresponding local image patches in a self-supervised manner.

where cosh , i denotes the cosine similarity and τ is a learnable temperature parameter. Similarly, the vision-to-text loss can be obtained by:

$$L^{v2t} = \sum_{i=1}^N \log \frac{\exp(\text{cosh}t_i, \mathbf{v}_i i / \tau)}{\sum_{j=1}^N \exp(\text{cosh}t_j, \mathbf{v}_i i / \tau)}. \quad (2)$$

The total training loss is $L_{\text{CLIP}} = (L^{t2v} + L^{v2t})/2$.

3.2 Synthetic Long Caption from Image

The utilization of paired image-text datasets is pivotal in fostering models capable of perceiving semantically aligned information across modalities. However, most of these datasets include a non-negligible portion of noisy and mismatched image-text pairs, substantially affecting visual-language representation learning. Thus, existing methods apply some short captions generated from MLLMs for real images to improve representation learning of language-vision pretrained models. The existing approaches, which correlate each image with a short caption, barely scratch the surface of the intricate tapestry woven by real-world data. They ignore exploring how to unleash the potential of long captions that can describe a real image in detail.

Given an image-text paired dataset $D = f(I_i, T_i)g_{i=1}^N$, containing N image-text pairs, we adopt a pretrained MLLM as a captioner f to generate a collection

of long and short captions of images:

$$C = fC_i^l, C_i^s g_{i=1}^N = ff(I_i, q_l), f(I_i, q_s)g_{i=1}^N, \quad (3)$$

where C_i^l and C_i^s denote the long caption and the short caption of image I_i . In addition, q_s and q_l denote the text inputs for generating long and short captions to the MLLM, respectively. We use the simple question template “Describe the image in details” to query the detailed long captions. Meanwhile, “Describe the image in short.” is used as the prompt for generating short captions. The question of detailed long captions has little impact on the diversity of answers, so we can obtain comprehensive captions of each image. Short captions, being concise and less prone to inaccuracies, naturally complement their longer captions by providing a succinct essence of the content.

How long are these long captions? For most raw/short caption labels, the number of captions in the label is one or two, with a token count of approximately 20. An image can include much effective information that needs substantial captions to describe its visual content. Thus, We have counted the number of tokens and sub-captions to see the amount of information in the generated long captions. Although there are some hallucinations in the generated caption, long captions can still bring much effective information that accurately describes the image. We also use different MLLMs (*e.g.*, InstructBLIP and LLaVA-1.5) to generate long captions to help CLIP model training. The experimental results and statistics of different long captions are presented in Sec. 4.6.

3.3 Global Multi-Positive Contrastive Learning

Sometimes complex and multiple information can be conveyed by a single image, which conveys its meaning or essence more effectively than a mere verbal description. Thus, a picture, which is worth a thousand words, should be described in multiple sentences. Inspired by this, we would like to design a strategy for long captions, towards completely utilizing the information of an image. Given original caption T , generated long caption $C^l = [c_1, \dots, c_M]$ and short caption $C^s = [c_s]$, we firstly conduct a sub-caption set that includes different sub-captions. Then we can easily implement a straightforward random sampling process to sample some sub-captions from the sub-caption set:

$$S_{i,j} \sim \text{Uniform}([T, c_s, c_1, \dots, c_M]), \quad (4)$$

where $S_{i,j}$ refers to j -th sub-caption of i -th sample from the sub-caption set. Following LaCLIP [12], the training multi-positive loss over the images becomes:

$$L_{\text{MPCL}}^{t2v} = \sum_{i=1}^N \sum_{j=1}^K \log \frac{\exp(\cos/\mathbf{v}_i, \mathbf{t}_{i,j}/\tau)}{\sum_{n=1}^N \exp(\cos/\mathbf{v}_n, \mathbf{t}_{i,j}/\tau)}, \quad (5)$$

where $\mathbf{t}_{i,j}$ refers to the text embedding of sub-caption $S_{i,j}$, and K denotes the number of sampling sub-captions. Similar to CLIP loss, $L_{\text{MPCL}} = (L_{\text{MPCL}}^{t2v} + L_{\text{MPCL}}^{v2t})/2$. This innovative process allows us to forge an image-text dataset with

long captions, wherein we employ a multi-positive loss framework to intricately intertwine sub-captions with their corresponding images, crafting a rich tapestry of aligned modalities.

3.4 Subcaption-specific Grouping Loss

Multi-positive contrastive learning with long captions can help the model perceive better global representation. To enhance the fine-grained representation of CLIP, existing methods [53] try to utilize word tokens to align the image patches, which may introduce the alignment of words unrelated to visual concepts (*e.g.*, emotion words or conjunctions). Meanwhile, words in a short caption may not take care of whole image. Thus, we would like to delve into learning subcaption-specific groupings of image patches for improving the fine-grained ability of the vision-language pre-training model. Multiple sub-captions from the long caption can typically delineate various local regions of the image.

Given an image and its long caption, we first calculate the cosine similarity between text embedding \mathbf{t}_i of i -th sub-caption and its image patch embedding $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_{HW}]$ to localize the subcaption-specific groupings of image patches. We view the cross-attention weights as a similarity matrix $\mathbf{W} = f\hat{w}_{i,j}g$ and sparse it to let each sub-caption only focus on a few visual tokens, *i.e.*

$$\mathbf{w}_{i,j} = \begin{cases} \hat{w}_{i,j} & \text{if } \hat{w}_{i,j} \geq \sigma \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where σ is the sparsity threshold. After that, these weights are used to select some subcaption-specific grouping visual tokens and pool them together:

$$\hat{\mathbf{v}}_j = \sum_{n=1}^{HW} \frac{\mathbf{w}_{i,j}}{\sum_j \mathbf{w}_{i,j}} \mathbf{v}_n. \quad (7)$$

The subcaption-specific grouping loss can be formulated as follows

$$\mathcal{L}_{\text{Sub}} = \sum_{i=1}^N \sum_{j=1}^{M+2} \log \frac{\exp(\cos(\hat{\mathbf{v}}_{i,j}, \mathbf{t}_{i,j})/\tau)}{\sum_{n=1}^K \exp(\cos(\hat{\mathbf{v}}_{i,n}, \mathbf{t}_{i,j})/\tau)}, \quad (8)$$

By applying this subcaption-specific grouping loss, we achieve a fine-grained alignment that meticulously pairs local image patches with their respective textual embeddings across different sub-captions. This nuanced pairing brings forth an unprecedented level of semantic alignment between the two modalities.

3.5 Overall objective

The overall DreamLIP objective is a weighted sum of the multi-positive contrastive loss and the finegrained alignment contrastive loss:

$$\mathcal{L}_{\text{DreamLIP}} = \lambda_{\text{MPCL}} \mathcal{L}_{\text{MPCL}} + \lambda_S \mathcal{L}_{\text{Sub}}, \quad (9)$$

where the loss weights λ_{MPCL} and λ_S are hyperparameters.

Table 1: Zero-shot image-text retrieval on the test splits of Flickr30k and MSCOCO. All models are pre-trained on YFCC15M. We use ViT-B/32 as image backbone. * denotes that we report results for pre-trained ViT-B/32 from OpenCLIP code base. Long caption generated from ShareGPT4V [4] is used.

Data	Method	Text Retrieval						Image Retrieval					
		Flickr30k			MSCOCO			Flickr30k			MSCOCO		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
CC3M	CLIP [40]	26.6	52.5	63.2	13.4	32.0	43.3	18.3	39.4	49.7	10.1	25.6	35.7
	DreamLIP	57.6	84.4	89.6	33.4	60.7	72.0	42.2	69.0	77.7	23.4	47.2	58.6
CC12M	CLIP [40]	49.3	77.3	85.0	29.3	54.4	65.3	35.5	61.8	71.6	19.0	41.0	52.5
	DreamLIP	78.7	94.6	97.6	53.4	77.1	84.7	61.0	83.9	89.8	36.7	62.3	72.3
YFCC15M	CLIP [40]	34.9	63.9	75.9	20.8	43.9	55.7	23.4	47.2	58.9	13.0	31.7	42.7
	SLIP [36]	47.8	76.5	85.9	27.7	52.6	63.9	32.3	58.7	68.8	18.2	39.2	51.0
	DeCLIP [27]	51.4	80.2	88.9	28.3	53.2	64.5	34.3	60.3	70.7	18.4	39.6	51.4
	UniCLIP [25]	52.3	81.6	89.0	32.0	57.7	69.2	34.8	62.0	72.0	20.2	43.2	54.4
	MCD [21]	57.6	82.6	91.1	32.3	58.7	71.2	36.4	64.8	74.1	20.7	43.5	55.3
	HiCLIP [17]	-	-	-	34.2	60.3	70.9	-	-	-	20.6	43.8	55.3
	HiDeCLIP [17]	-	-	-	38.7	64.4	74.8	-	-	-	23.9	48.2	60.1
	FILIP [53]	-	-	-	33.4	60.1	-	-	-	-	23.0	46.2	-
	ALIP [52]	70.5	91.9	95.7	46.8	72.4	81.8	48.9	75.1	82.9	29.3	54.4	65.4
	DreamLIP	84.9	97.3	98.7	55.7	80.5	88.2	66.0	86.4	91.4	39.8	66.0	75.5
Merged-30M	CLIP [40]	57.8	84.6	91.7	35.0	61.9	73.7	44.0	70.9	79.7	23.5	47.3	59.1
	DreamLIP	87.2	97.5	98.8	58.3	81.6	88.8	66.4	88.3	93.3	41.1	67.0	76.6
LAION-400M*	CLIP [40]	78.7	94.0	97.3	53.7	77.1	85.4	61.8	85.5	90.8	34.8	60.4	71.1

4 Experiments

4.1 Implementation Details and Datasets

Pretraining Datasets. To make a fair comparison with the state-of-the-art contrastive vision-language pretraining approaches, we adopt the CC3M, CC12M and YFCC15M datasets. In addition, we construct a 30M version of pretraining data by including Conceptual Caption 3M (CC3M) [44] and 12M (CC12M) [44]. We mainly conduct ablation studies to validate our model on the CC3M data.

Downstream Datasets. Following CLIP, we select 11 visual recognition datasets under the zero-shot setting, namely ImageNet [8], CIFAR 10 & CIFAR 100 [23], StanfordCars [22], Caltech101 [13], Flowers102 [37], SUN397 [49], DTD [6], FGVAircraft [34], OxfordPets [39], and Food101 [2]. The same zero-shot classification protocol is applied following [40], which uses predefined prompts as text inputs. Although CLIP only evaluates on visual recognition, we also provide comprehensive comparisons on vision-language tasks which are more desired in evaluating multi-modal models, including image-text retrieval on MSCOCO/Flickr30k Caption [30, 54], semantic segmentation (*i.e.*, ADE20K-150 [58], ADE20K-847 [58], VOC-20 [11], PC-59 [35] and PC-459 [35]) as well as vision-language reasoning on ScienceQA-IMG [33], TextVQA [45], POPE [29] and MMVP [47].

Implementation Details. Two variants of Vision Transformer are used as the image encoder in our experiments, *i.e.*, ViT-B/32 and ViT-B/16, while the text encoder is a vanilla Transformer following CLIP as a fair comparison. The embedding size of both image and text features are 512 throughout our paper. To make a fair comparison with CLIP family baselines, we train all models for 32 epochs under the same set of pretraining hyperparameters including learning rate, warmup steps, weight decay, *etc.* The input image size is set to 224 × 224, and the input text sequence length is truncated or padded to 77. Following CLIP, the learnable temperature parameter τ is initialized as 0.07.

4.2 Image-Text Retrieval

After pre-training, the proposed model is evaluated in a zero-shot setting on image-text retrieval tasks, *i.e.*, COCO, and Flickr30K. The pre-trained model is applied to extract embeddings from images and texts, respectively. Similarity scores between image embeddings and text embeddings are used for ranking. We use the R@K to report the recall of top-K retrieval items. As shown in Table 1, quantitative experimental results demonstrate our superiority over state-of-the-art alternatives in terms of all metrics. Thanks to long captions, our model can achieve better performance than CLIP trained by 400M image-text datasets.

4.3 Image Recognition

We have verified the zero-shot classification capability of DreamLIP on 11 common classification benchmarks. Top-1 accuracy is used for evaluation.

For the image classification task, we construct prompts with class label names following the setting in CLIP. Then text embeddings are extracted from these text inputs with class label names. Given an input image, the distances from image embeddings to the text embeddings are computed, and the class label is predicted based on the closest distance. We compare DreamLIP with Zero-shot CLIP and some state-of-the-art methods on the 11 datasets as mentioned above, demonstrated in Table 2. DreamLIP outperforms other SOTA methods on average over 11 datasets, which approves the ability of our pre-trained model to the downstream tasks. It indicates that the long captions are able to enhance the zero-shot performance of CLIP directly transferring to the downstream task.

We also evaluate our model in three multi-classification benchmarks. As shown in Table 5, our method can achieve on par performance with the CLIP trained by 400M datasets. The contents of an image can be so rich that describing them well requires lengthy captions. Thus, an image with lengthy captions can help the model perceive more objects and understand more accurate relations between objects in images.

4.4 Semantic Segmentation

To certify the fine-grained representational capacity of DreamLIP, we evaluate the transferable performance on semantic segmentation tasks following SAN [50].

Table 2: Zero-shot transfer evaluation of different models. Performance on ImageNet and 10 common downstream datasets are reported. We highlight the best performance of each setting in **bold**. **Long caption generated from ShareGPT4V [4] is used.**

Data	Model	Food-101	CIFAR-10	CIFAR-100	SUN397	Cars	Aircraft	DTD	Pets	Caltech-101	Flowers	Average	ImageNet
<i>Model Architecture: ViT-B/32</i>													
CC3M	CLIP [40]	10.2	71.3	32.1	33.8	1.4	1.0	12.0	12.1	50.9	10.8	23.6	17.2
	DreamLIP	16.1	82.0	45.4	41.3	2.5	1.0	13.9	18.8	64.4	14.1	30.0	25.9
CC12M	CLIP [40]	26.5	72.5	38.0	37.1	13.7	2.6	11.4	46.2	74.0	25.7	34.8	32.9
	DreamLIP	48.9	86.4	63.0	55.7	17.9	1.9	23.5	41.9	83.2	25.8	44.8	44.2
YFCC15M	SLIP [36]	33.3	50.7	25.5	34.7	2.8	1.7	14.4	23.5	59.9	49.0	29.6	34.3
	FILIP [53]	43.1	65.5	33.5	50.7	3.3	3.2	24.3	24.1	68.8	52.7	36.9	39.5
	DeCLIP [27]	52.5	66.7	38.7	50.3	3.8	2.1	27.7	33.8	74.7	60.8	41.1	43.2
	HiCLIP [17]	51.2	74.1	46.0	50.6	4.5	3.6	23.1	37.8	67.4	60.9	41.9	40.5
	ALIP [52]	45.4	83.8	51.9	47.8	3.4	2.7	23.2	30.7	74.1	54.8	41.8	40.3
	UniCLIP [25]	48.7	78.6	47.2	50.4	3.4	2.8	23.3	32.5	73.0	8.1	36.8	42.8
	MCD [21]	54.0	80.3	49.6	55.3	4.5	3.0	30.5	40.0	73.2	7.9	39.8	44.7
	DreamLIP	51.7	87.9	60.7	54.8	9.4	7.1	26.8	36.3	79.6	48.6	46.3	46.6
Merged-30M	CLIP [40]	60.7	83.6	54.7	54.6	16.2	6.9	25.9	60.5	81.2	53.5	49.8	49.0
	HiCLIP [17]	63.9	77.6	56.2	60.7	22.2	5.5	38.0	65.6	82.4	62.5	53.5	52.9
	DreamLIP	68.2	91.8	69.2	62.2	20.7	8.0	32.1	62.8	86.1	48.5	55.0	55.7
LAION-400M	CLIP [40]	79.9	91.8	72.0	64.6	77.0	15.8	49.9	84.8	89.3	64.4	62.7	62.0
<i>Model Architecture: ViT-B/16</i>													
CC3M	LaCLIP [12]	14.2	57.1	27.5	35.1	1.6	1.6	16.6	15.6	52.7	14.7	23.7	21.5
	MLLM-A [32]	18.7	58.4	32.4	43.8	3.9	1.5	20.2	32.1	63.5	17.5	29.2	25.0
	CLIP [40]	10.3	54.9	21.8	25.0	0.8	1.4	10.5	12.8	43.3	10.2	19.1	20.3
	DreamLIP	19.4	74.3	44.2	45.9	2.8	1.0	17.0	27.1	63.1	14.7	31.0	31.1
CC12M	LaCLIP [12]	60.7	75.1	43.9	57.0	36.3	5.6	31.0	72.4	83.3	39.9	46.2	48.4
	MLLM-A [32]	60.9	83.0	55.4	59.4	24.1	3.2	30.7	64.8	79.3	36.0	49.7	47.5
	CLIP [40]	25.3	66.5	32.1	39.9	14.7	1.9	13.5	45.0	59.8	15.0	31.4	34.0
	DreamLIP	58.3	87.3	62.6	54.3	29.7	4.9	29.2	60.3	83.1	28.9	49.9	50.2
YFCC15M	CLIP [40]	35.0	67.1	34.8	42.0	5.1	6.3	13.9	20.4	54.5	44.3	32.3	34.1
	DreamLIP	44.2	89.0	62.0	57.1	9.2	6.4	30.5	32.6	79.8	40.2	45.1	48.2
Merged-30M	CLIP [40]	64.5	87.5	60.3	61.1	25.4	6.9	33.7	58.1	84.5	57.3	53.9	55.2
	DreamLIP	75.4	92.3	70.8	63.9	22.7	7.9	33.9	64.1	88.3	51.8	57.1	58.4
LAION-400M	CLIP [40]	85.5	93.0	71.7	66.8	83.5	16.7	52.8	90.1	91.2	63.9	65.5	67.0

Concretely, We replace and fix the pre-trained model in SAN and fine-tune it with COCO-stuff [3]. As shown in Table 3, DreamLIP significantly surpasses CLIP across all scale pretraining datasets (*i.e.*, 3M, 12M, 15M and 30M). Notably, DreamLIP exceeds CLIP by 2.5% on average over 5 widely used semantic segmentation datasets with fewer data (30m *v.s.* 400m). The results indicate DreamLIP can provide reliable fine-grained clues for downstream semantic segmentation tasks with the help of long captions.

Table 3: Transferable performance of semantic segmentation on ADE-847, PC-459, ADE-150, PC-59, and VOC-20. Following SAN [50], we used the full training set of COCO-stuff as the training data and our DreamLIP as pretrained models. We use ViT-B/32 as image backbone. Long caption generated from ShareGPT4V [4] is used.

Data	Method	ADE-847	PC-459	ADE-150	PC-59	VOC-20	avg.
CC3M	CLIP [40]	2.1	5.2	12.3	33.8	65.4	23.8
	DreamLIP	4.1	7.5	17.1	39.9	76.5	29.0
CC12M	CLIP [40]	3.3	6.7	15.7	39.2	79.7	28.9
	DreamLIP	6.1	10.0	23.3	43.6	85.5	33.7
YFCC15M	CLIP [40]	3.2	8.1	14.4	42.0	82.3	30.0
	DreamLIP	6.4	11.1	22.4	48.9	88.2	35.4
Merged-30M	CLIP [40]	5.8	10.2	21.0	45.8	86.9	33.9
	DreamLIP	8.1	12.5	25.3	49.9	90.9	37.3
Laion-400M	CLIP [40]	6.1	12.2	21.3	46.3	88.3	34.8

Table 4: Image understanding performance of DreamLIP in MLLM. We use ViT-B/32 as image backbone. The best results are in **bold** and the second best are underlined. Long caption generated from ShareGPT4V [4] is used.

Data	Method	ScienceQA-IMG	TextVQA	POPE	MMVP	avg.
CC3M	CLIP [40]	65.7	48.0	79.0	14.7	51.8
	DreamLIP	65.1	47.8	79.6	17.3	52.4
CC12M	CLIP [40]	65.3	48.0	79.0	15.3	51.9
	DreamLIP	66.5	47.6	80.3	20.3	53.7
YFCC15M	CLIP [40]	65.7	48.0	79.6	18.3	52.8
	DreamLIP	65.5	48.2	82.0	<u>22.0</u>	54.4
Merged-30M	CLIP [40]	65.5	48.6	81.0	19.3	53.6
	DreamLIP	<u>66.8</u>	<u>48.8</u>	<u>81.8</u>	22.7	55.0
Laion-400M	CLIP [40]	67.3	50.2	80.9	19.7	<u>54.5</u>

4.5 DreamLIP in MLLM

Here we evaluate the image understanding performance of our DreamLIP in MLLM. We follow the training process of LLaVA-1.5 [31] which fixes the visual encoder of CLIP and combines the encoder with LLM, simply replacing LLaVA-1.5’s CLIP encoder with several CLIP encoders shown in Tab. 4 without further tuning. We pick up ScienceQA-IMG [33], TextVQA [45] and POPE [29] from evaluation benchmarks for LLaVA-1.5 which can be directly obtained the results without submitting for website responses. MMVP benchmark [47] which exposes the visual confusion of CLIP in MLLM is also selected. As demonstrated in Tab. 4, our DreamLIP trained on 30M datasets outperforms CLIP utilizing the same data scale and achieves a competitive performance with CLIP employing 400M datasets. It indicates that long captions are beneficial for holistic visual understanding through image-text joint training because they augment CLIP-

Table 5: Zero-shot performance on multi-class recognition tasks. OpenImage and ImageNet-Multi contain multiple objects in an image. hico is a relation recognition dataset. We use ViT-B/32 as image backbone. Long caption generated from ShareGPT4V [4] is used.

Data	Method	ImageNet-Multi	OpenImage		HICO
			Class-Common	Class-Rare	
CC3M	CLIP [40]	12.90	48.82	26.78	10.34
	DreamLIP	21.19	58.86	38.51	14.32
CC12M	CLIP [40]	27.69	58.10	49.23	13.86
	DreamLIP	37.63	71.48	56.94	19.43
YFCC15M	CLIP [40]	26.74	58.27	42.42	12.25
	DreamLIP	37.85	66.50	50.87	18.87
Merged-30M	CLIP [40]	41.87	71.94	59.01	19.83
	DreamLIP	48.19	74.32	61.58	25.36
Laion-400M	CLIP [40]	53.74	73.05	67.44	22.47

like models to mine the rich visual content hidden in images, and finally enhance the image understanding ability in MLLM.

4.6 Ablation Studies

Effectiveness of Each Component. To further explore the effectiveness of short captions, long captions and subcaption-specific grouping loss, we perform ablation experiments based on the zero-shot image-text retrieval task, zero-shot image classification and semantic segmentation.

As shown in Tab. 6, we first introduce short captions into the baseline (*i.e.*, CLIP), which achieves similar performance with direct using long captions during training. It demonstrates that long captions and short captions can help learn richer information from image than raw captions. Directly training with the long captions may not unleash their potential. Thus, we design a uniform sampling strategy for long captions. In this way, we can observe that the performance is better than directing to use long captions, indicating that model can be benefit from the multi-positive pairs including many sub-captions. Then, when combining the short captions and sampling sub-captions together, the performance is further improved. This is because the short captions are concise in describing the whole image, while long captions are more details but has some hallucinations, which can complement each other. Finally, with the help of grouping loss, our DreamLIP can achieve the best performance in terms of all metrics. We design a simple yet effective strategy to use the long captions in language-image pre-training.

Long Captions from Different MLLMs. Given the significance of synthetic captions in this study, we investigate the impact of captions generated by different MLLMs on downstream tasks. The experiment results are presented

Table 6: Ablation study of different designs. ‘S.C’ refers to short captions generated by MLLM. ‘SGL’ refers to subcaption-specific grouping losses. We use ViT-B/16 as image backbone.

Long Captions		S.C	SGL	Text Retrieval		Image Retrieval		Classification	Segmentation
Direct	Sampling			Flickr30k	MSCOCO	Flickr30k	MSCOCO	ImageNet	VOC-20
				R@1	R@1	R@1	R@1	Acc.(%)	mIOU
				32.6	14.8	21.4	11.5	20.3	64.4
		✓		55.1	32.7	43.3	23.0	25.6	77.7
✓				56.6	30.2	40.9	21.4	24.4	75.7
	✓			63.0	35.7	49.0	25.6	30.0	81.8
	✓	✓		68.3	40.8	53.4	29.4	30.1	82.9
	✓	✓	✓	69.5	42.8	54.4	30.4	31.1	84.5




Fig. 3: Statistics of long captions generated by MLLMs (*i.e.*, InstructBLIP, LLaVA-1.5 and ShareGPT4V).

in Table 7. It is worth noting that using long captions generated by InstructBLIP [7] achieves better performance than using short captions. Additionally, we provide some stronger MLLMs to generate long captions, which can bring a more significant improvement. As shown in Figure 3, detailed captions, characterized by a greater number of tokens and subcaptions, can capture the contents of an image more comprehensively than typically possible with short captions. Long captions generated by ShareGPT4V [4] achieve the best performance, and have longer tokens and subcaptions, which demonstrates the effectiveness of its long captions.

Number of sub-captions. We evaluate the performance for different numbers K of sampled sub-captions from the sub-caption set. As shown in Table 8, we observe that as the number of sub-captions increase, the performance gradually improves in terms of zero-shot classification, image-text retrieval, and semantic segmentation. However, when the number K of sub-captions reached approximately 8, the performance showed little to no further improvement. This phenomenon can be attributed to the fact that the number of tokens and sub-captions derived from long captions in the synthetic caption dataset reached its

Table 7: Ablation study of long captions from different MLLMs. We use ViT-B/16 as image backbone.

MLLM	Text Retrieval		Image Retrieval		Classification	Segmentation
	Flickr30k	MSCOCO	Flickr30k	MSCOCO	ImageNet	VOC-20
	R@1	R@1	R@1	R@1	Acc.(%)	mIOU
CLIP <i>w/o</i> MLLM	32.6	14.8	21.4	11.5	20.3	64.4
(1) InstructBLIP	58.7	34.4	45.2	24.9	27.8	79.2
(2) LLaVA-1.5	66.8	42.4	53.3	29.9	29.0	81.8
(3) ShareGPT4V	69.5	42.8	54.4	30.4	31.1	84.5
(1)+(2)+(3) *	74.4	46.4	62.4	34.9	34.6	88.2

peak, as shown in Figure 3. Increasing the number of sub-captions may result in redundant samples, which do not provide additional information to enhance model training.

Table 8: Ablation study of sampling number of sub-captions from long captions. We use ViT-B/16 as image backbone.

K	Text Retrieval		Image Retrieval		Classification	Segmentation
	Flickr30k	MSCOCO	Flickr30k	MSCOCO	ImageNet	VOC-20
	R@1	R@1	R@1	R@1	Acc.(%)	mIOU
CLIP [40]	32.6	14.8	21.4	11.5	20.3	64.4
3	65.4	37.4	49.5	26.9	29.4	82.0
4	68.0	38.2	51.6	28.6	30.8	82.1
5	68.5	39.0	53.6	29.2	30.9	79.6
6	69.1	40.6	53.2	29.5	31.3	81.1
7	70.0	41.4	53.1	29.6	31.1	84.0
8	70.9	41.5	53.0	29.8	31.0	84.5
9	70.5	41.9	54.0	30.5	31.1	83.0
10	69.5	42.8	54.4	30.4	31.1	84.5

5 Conclusion

We re-caption 30M images with detailed descriptions using a pre-trained MLLM and explore the usage of these long captions under a contrastive learning framework. Specifically, we propose to dynamically sample sub-captions from the text label to construct multiple positive pairs, and introduce a grouping loss to match the embeddings of each sub-caption with its corresponding local image patches in a self-supervised manner. Experimental results on a wide range of downstream tasks demonstrate the consistent superiority of our method, termed DreamLIP, over previous alternatives, highlighting its fine-grained representational capacity. This work represents a promising direction for enhancing multi-modality foundation models, and we anticipate it will inspire further research.

Acknowledgements

This work is supported in part by NSFC 62302246 and ZJNSFC under Grant LQ23F010008, and supported by High Performance Computing Center at Eastern Institute of Technology, Ningbo, and Ningbo Institute of Digital Twin. Wei Chen is supported by the National Science Foundation of China (62132017), Zhejiang Provincial Natural Science Foundation of China (LD24F020011).

References

1. 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)
2. Bossard, L., Guillaumin, M., Van Gool, L.: Food-101—mining discriminative components with random forests. In: *Eur. Conf. Comput. Vis.* (2014)
3. Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1209–1218 (2018)
4. 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. *arXiv preprint arXiv:2311.12793* (2023)
5. Chen, Y., Li, W., Chen, X., Gool, L.V.: Learning semantic segmentation from synthetic data: A geometrically guided input-output adaptation approach. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1841–1850 (2019)
6. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2014)
7. Dai, W., Li, J., Li, D., Tiong, A.M.H., Zhao, J., Wang, W., Li, B., Fung, P.N., Hoi, S.: InstructBLIP: Towards general-purpose vision-language models with instruction tuning. *Adv. Neural Inform. Process. Syst.* **36** (2024)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2009)
9. Dong, X., Bao, J., Zheng, Y., Zhang, T., Chen, D., Yang, H., Zeng, M., Zhang, W., Yuan, L., Chen, D., et al.: Maskclip: Masked self-distillation advances contrastive language-image pretraining. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 10995–11005 (2023)
10. Dou, Z.Y., Kamath, A., Gan, Z., Zhang, P., Wang, J., Li, L., Liu, Z., Liu, C., LeCun, Y., Peng, N., et al.: Coarse-to-fine vision-language pre-training with fusion in the backbone. *Adv. Neural Inform. Process. Syst.* **35**, 32942–32956 (2022)
11. Everingham, M., Winn, J.: The pascal visual object classes challenge 2012 (voc2012) development kit. *Pattern Anal. Stat. Model. Comput. Learn., Tech. Rep* **2007**(1-45), 5 (2012)
12. Fan, L., Krishnan, D., Isola, P., Katabi, D., Tian, Y.: Improving clip training with language rewrites. *Adv. Neural Inform. Process. Syst.* **36** (2024)
13. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2004)
14. Fürst, A., Rumetshofer, E., Lehner, J., Tran, V.T., Tang, F., Ramsauer, H., Kreil, D., Kopp, M., Klambauer, G., Bitto, A., et al.: Cloob: Modern hopfield networks with infoloob outperform clip. *Adv. Neural Inform. Process. Syst.* **35**, 20450–20468 (2022)

15. Gao, Y., Liu, J., Xu, Z., Wu, T., Liu, W., Yang, J., Li, K., Sun, X.: Softclip: Softer cross-modal alignment makes clip stronger. arXiv preprint arXiv:2303.17561 (2023)
16. Gao, Y., Liu, J., Xu, Z., Zhang, J., Li, K., Ji, R., Shen, C.: Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *Adv. Neural Inform. Process. Syst.* **35**, 35959–35970 (2022)
17. Geng, S., Yuan, J., Tian, Y., Chen, Y., Zhang, Y.: HiCLIP: Contrastive language-image pretraining with hierarchy-aware attention. In: *Int. Conf. Learn. Represent.* (2023)
18. Hammoud, H.A.A.K., Itani, H., Pizzati, F., Torr, P., Bibi, A., Ghanem, B.: Synthclip: Are we ready for a fully synthetic clip training? arXiv preprint arXiv:2402.01832 (2024)
19. 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: *Int. Conf. Mach. Learn.* (2021)
20. Johnson-Roberson, M., Barto, C., Mehta, R., Sridhar, S.N., Rosaen, K., Vasudevan, R.: Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks? arXiv preprint arXiv:1610.01983 (2016)
21. Kim, B., Jo, Y., Kim, J., Kim, S.: Misalign, contrast then distill: Rethinking misalignments in language-image pre-training. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2563–2572 (2023)
22. Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for fine-grained categorization. In: *ICCVW* (2013)
23. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
24. Lai, Z., Zhang, H., Wu, W., Bai, H., Timofeev, A., Du, X., Gan, Z., Shan, J., Chuah, C.N., Yang, Y., et al.: From scarcity to efficiency: Improving clip training via visual-enriched captions. arXiv preprint arXiv:2310.07699 (2023)
25. Lee, J., Kim, J., Shon, H., Kim, B., Kim, S.H., Lee, H., Kim, J.: Unclip: Unified framework for contrastive language-image pre-training. arXiv:2209.13430 (2022)
26. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: *Int. Conf. Mach. Learn.* (2022)
27. Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., Yu, F., Yan, J.: Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. arXiv:2110.05208 (2021)
28. Li, Y., Fan, H., Hu, R., Feichtenhofer, C., He, K.: Scaling language-image pre-training via masking. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 23390–23400 (2023)
29. Li, Y., Du, Y., Zhou, K., Wang, J., Zhao, W.X., Wen, J.R.: Evaluating object hallucination in large vision-language models. arXiv preprint arXiv:2305.10355 (2023)
30. 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: *Eur. Conf. Comput. Vis.* pp. 740–755. Springer (2014)
31. Liu, H., Li, C., Li, Y., Lee, Y.J.: Improved baselines with visual instruction tuning. arXiv preprint arXiv:2310.03744 (2023)
32. Liu, Y., Wang, K., Shao, W., Luo, P., Qiao, Y., Shou, M.Z., Zhang, K., You, Y.: Mllms-augmented visual-language representation learning. arXiv preprint arXiv:2311.18765 (2023)

33. Lu, P., Mishra, S., Xia, T., Qiu, L., Chang, K.W., Zhu, S.C., Tafjord, O., Clark, P., Kalyan, A.: Learn to explain: Multimodal reasoning via thought chains for science question answering. *Adv. Neural Inform. Process. Syst.* **35**, 2507–2521 (2022)
34. Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. *arXiv:1306.5151* (2013)
35. Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: The role of context for object detection and semantic segmentation in the wild. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 891–898 (2014)
36. Mu, N., Kirillov, A., Wagner, D., Xie, S.: Slip: Self-supervision meets language-image pre-training. In: *Eur. Conf. Comput. Vis.* (2022)
37. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *Sixth Indian Conference on Computer Vision, Graphics & Image Processing* (2008)
38. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv:1807.03748* (2018)
39. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: *Int. Conf. Comput. Vis.* (2012)
40. 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: *Int. Conf. Mach. Learn.* (2021)
41. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
42. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: *Eur. Conf. Comput. Vis.* pp. 102–118. Springer (2016)
43. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 10684–10695 (June 2022)
44. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: *Association for Computational Linguistics* (2018)
45. Singh, A., Natarajan, V., Shah, M., Jiang, Y., Chen, X., Batra, D., Parikh, D., Rohrbach, M.: Towards vqa models that can read. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8317–8326 (2019)
46. Tian, Y., Fan, L., Isola, P., Chang, H., Krishnan, D.: Stablerep: Synthetic images from text-to-image models make strong visual representation learners. *Adv. Neural Inform. Process. Syst.* **36** (2024)
47. Tong, S., Liu, Z., Zhai, Y., Ma, Y., LeCun, Y., Xie, S.: Eyes wide shut? exploring the visual shortcomings of multimodal llms. *arXiv preprint arXiv:2401.06209* (2024)
48. Wu, S., Fei, H., Zhang, H., Chua, T.S.: Imagine that! abstract-to-intricate text-to-image synthesis with scene graph hallucination diffusion. *Adv. Neural Inform. Process. Syst.* **36** (2023)
49. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: Large-scale scene recognition from abbey to zoo. In: *Int. Conf. Comput. Vis.* (2010)
50. Xu, M., Zhang, Z., Wei, F., Hu, H., Bai, X.: Side adapter network for open-vocabulary semantic segmentation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2945–2954 (2023)
51. Yang, J., Li, C., Zhang, P., Xiao, B., Liu, C., Yuan, L., Gao, J.: Unified contrastive learning in image-text-label space. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 19163–19173 (2022)

52. Yang, K., Deng, J., An, X., Li, J., Feng, Z., Guo, J., Yang, J., Liu, T.: Alip: Adaptive language-image pre-training with synthetic caption. In: *Int. Conf. Comput. Vis.* pp. 2922–2931 (2023)
53. Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., Liang, X., Li, Z., Jiang, X., Xu, C.: Filip: fine-grained interactive language-image pre-training. *arXiv:2111.07783* (2021)
54. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics* pp. 67–78 (2014)
55. Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917* (2022)
56. Yuan, J., Zhang, J., Sun, S., Torr, P., Zhao, B.: Real-fake: Effective training data synthesis through distribution matching. *arXiv preprint arXiv:2310.10402* (2023)
57. Zhao, L., Zheng, K., Zheng, Y., Zhao, D., Zhou, J.: Rleg: Vision-language representation learning with diffusion-based embedding generation. *Int. Conf. Mach. Learn.* (2023)
58. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 633–641 (2017)