



Alpha-CLIP: A CLIP Model Focusing on Wherever You Want

Zeyi Sun^{*1,4}, Ye Fang^{*2,4}, Tong Wu³, Pan Zhang⁴, Yuhang Zang⁴,
 Shu Kong⁵, Yuanjun Xiong⁶, Dahua Lin^{3,4}, Jiaqi Wang^{†4}

¹Shanghai Jiao Tong University ²Fudan University ³The Chinese University of Hong Kong
⁴Shanghai AI Laboratory ⁵University of Macau ⁶MThreads, Inc.

szy2023@sjtu.edu.cn, {fangye, zhangpan, zangyuhang, wangjiaqi}@pjlab.org.cn

<https://aleafy.github.io/alpha-clip>

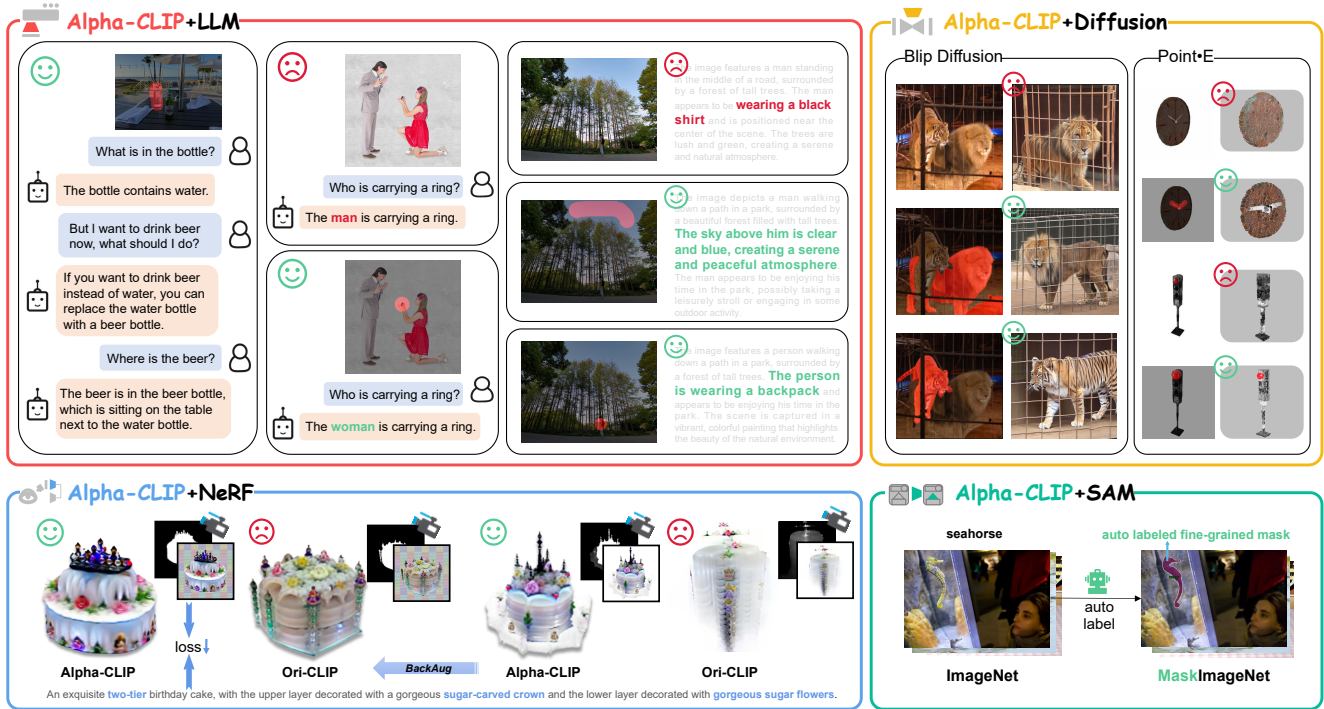


Figure 1. Usage of our proposed Alpha-CLIP. Our Alpha-CLIP can seamlessly replace the original CLIP in a wide range of tasks to allow the whole system to focus on any specified region given by points, strokes or masks. Cases marked with ☹️ are generated with the original CLIP. Cases marked with 😊 are generated with our Alpha-CLIP. All cases shown here are made simply by replacing the original CLIP of the system with a plug-in Alpha-CLIP without further tuning.

Abstract

Contrastive Language-Image Pre-training (CLIP) plays an essential role in extracting valuable content information from images across diverse tasks. It aligns textual and visual modalities to comprehend the entire image, including all the details, even those irrelevant to specific tasks. However, for a finer understanding and controlled editing of images, it becomes crucial to focus on specific regions of interest, which can be indicated as points, masks, or boxes

by humans or perception models. To fulfill the requirements, we introduce Alpha-CLIP, an enhanced version of CLIP with an auxiliary alpha channel to suggest attentive regions and fine-tuned with constructed millions of RGBA region-text pairs. Alpha-CLIP not only preserves the visual recognition ability of CLIP but also enables precise control over the emphasis of image contents. It demonstrates effectiveness in various tasks, including but not limited to open-world recognition, multimodal large language models, and conditional 2D / 3D generation. It has a strong potential to serve as a versatile tool for image-related tasks.

* Equal contribution. † Corresponding authors.

Domains	Components	Tasks	Methods	Advantages over the original CLIP
Image Recognition	Alpha-CLIP	Zero-shot Classification Zero-shot REC	-	Superior classification accuracy Excellent region-text comprehension ability
	Alpha-CLIP + SAM	Data Engine for OVD	Detic [76]	Higher OVD mAP
MLLM	Alpha-CLIP + LLM	VQA, Captioning	BLIP-2 [28], LLaVA-1.5 [33]	Region-focused captioning / VQA Eliminating hallucinations Reducing model bias
2D Generation	Alpha-CLIP + Diffusion	Image Variation	BLIP-Diffusion [27]	Controllable generation Enabling subject-driven generation in complex images
3D Generation	Alpha-CLIP + Diffusion	Generalized Image-to-3D	Point-E [39]	Rectifying absent parts
	Alpha-CLIP + NeRF	Optimized Image-to-3D	PureCLIPNeRF [25]	Improved 3D optimization results

Table 1. Downstream tasks of Alpha-CLIP and their advantages over the original CLIP

1. Introduction

Recent advances in Contrastive Language-Image Pre-training (CLIP) [19, 43] and its diverse variants [10, 30, 55] have established a robust framework for extracting semantically coherent features from both images and text. These features aim to capture all the semantic details within images, exhibiting potent representation capabilities and exceptional generalizability, making them versatile in a variety of downstream tasks, such as open-world recognition [7, 13, 62, 64, 65], Multimodal Large Language Models (MLLMs) [4, 18, 26, 28, 33, 34, 41, 56, 71], and 2D / 3D generation [20, 25, 27, 38, 39, 44, 68].

While CLIP captures the content of the entire image, it is also crucial to focus on the regions of interest to enable a finer understanding [16, 21, 24, 42, 53, 77] and controllable content generation [25, 38, 51, 60]. These regions can be specified by points, masks, or boxes via human interaction or perception models (e.g., SAM [22], GLIP [29] and proposal networks [70]).

To fulfill the demands of downstream tasks, researchers have attempted to acquire region-focused CLIP features using two primary strategies. The first method is to exclude non-relevant areas by cropping the regions of interest into distinct patches [7, 54, 73, 74] or applying masking to the irrelevant parts of images [31], features [31, 60], and attention masks [65, 75]. However, this approach disrupts (in cropping) and omits (in masking) contextual information, which is crucial for precise image understanding and reasoning. The second method is to highlight the regions of interest by circles [52] or mask contour [66] on the images fed to CLIP. Although user-friendly, it changes the original content of the images, which will result in undesirable recognition and generation results (cf. Fig. 2).

To achieve region focus without hurting original image, we propose Alpha-CLIP, which improves CLIP [43] by incorporating regions of interest through an additional alpha channel input. Along with the RGB channels, the introduced alpha channel enables the Alpha-CLIP to focus on designated areas while maintaining an awareness of the contextual information. While initialized with the CLIP [43] model, the training of Alpha-CLIP still requires a large set

of region-text paired data. By harnessing the Segment Anything Model (SAM) [22] and multimodal large models for image captioning, such as BLIP-2 [28], we develop an effective pipeline to generate millions of region-text pairs that are readily convertible to RGBA-text data. After training with a mixture of region-text pairs and image-text pairs, Alpha-CLIP can focus on the specific regions while maintaining the visual recognition accuracy of CLIP [43].

Alpha-CLIP can enhance CLIP across a wide array of downstream tasks, applying a plug-and-play methodology that permeates diverse domains, spanning from perception to generation in 2D and 3D applications, as shown in Fig. 1 and Tab. 1. Specifically, **1) Image Recognition:** Alpha-CLIP not only maintains the visual recognition ability of the original CLIP but also boosts the capability of region-based recognition. Specifically, when provided with ground-truth region to focus on, Alpha-CLIP achieves 4.1% improvement in top-1 accuracy on zero-shot ImageNet classification task. This superior region-based recognition ability helps downstream tasks like Referring Expression Comprehension(REC) [54] or serves as data engine for Open Vocabulary Detection(OVD) [76]. **2) Serving as vision backbone for MLLM:** In conjunction with a large language model, Alpha-CLIP becomes capable of facilitating region level captioning and VQA within a MLLM framework. This integration significantly mitigates the occurrences of hallucinations (e.g., black shirt) and diminishes model bias (e.g., man carrying a ring). **3) 2D generation:** When integrated with a diffusion model, Alpha-CLIP enhances the controllability of BLIP-Diffusion [27] in image variation tasks. In addition, it enables the extraction of subjects from complex images for subject-driven generation, surmounting an obstacle encountered when deploying BLIP-Diffusion with the original CLIP, which only supports single subjects in simplistic images. **4) 3D generation:** In addition to the capabilities in 2D generation, Alpha-CLIP exhibits proficiency in 3D generation as well. It can be effectively deployed in conjunction with a diffusion model, such as Point-E [39], to enhance the quality of 3D object generation. Additionally, it can be utilized with NeRF [37], exemplified by PureCLIPNeRF [25], to optimize the creation of superior 3D objects.

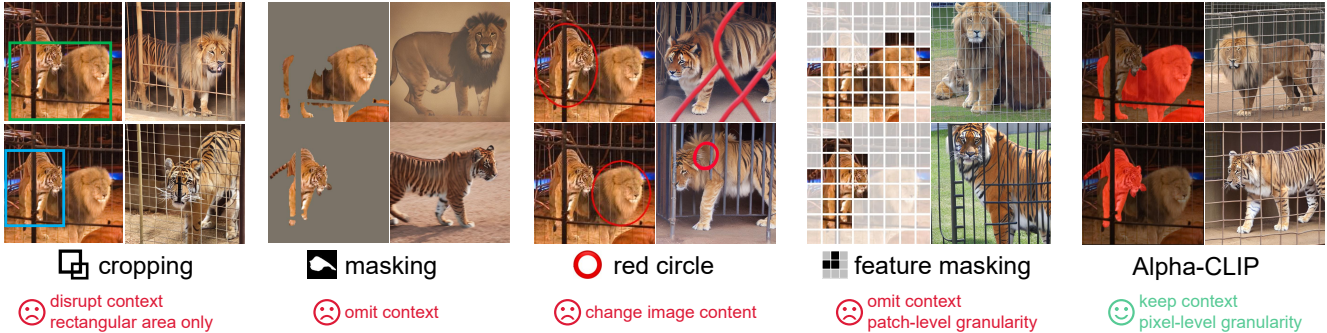


Figure 2. **Alpha-CLIP vs. other methods of region-focusing** for image generation using BLIP-Diffusion [27]. The fine-grained region focusing ability of Alpha-CLIP produces better results than these methods that adopt the original CLIP.

In summary, we propose Alpha-CLIP, which equips the original CLIP model with the capability of region awareness. Through fine-tuning on millions of RGBA region-text pairs, Alpha-CLIP demonstrates significant advantages over the original CLIP across various tasks, including but not limited to image recognition [43, 54, 76], multimodal large language models [28, 34], 2D generation [27, 44] and 3D generation [25, 39].

2. Related Work

Empowering CLIP with region awareness. To enable CLIP [43] to disentangle regions from the whole image for more targeted processing and understanding, various methods have been explored in the field of *segmentation*. Among them, MaskCLIP [75] uses a 1x1 convolution layer to extract CLIP’s final 2D features to obtain semantic information for different regions. SAN [64] trains a side network alongside CLIP to assist the model in local semantic perception. MaskCLIP [9] and ODISE [62] use attention masks to make CLIP focus more on local regions. These methods do not alter the weights of the CLIP model itself. RegionCLIP [74] generate region box-text pairs for local region and fine-tune CLIP model for box level recognition. MaskAdaptedCLIP[31] generates mask-text pairs for local masks through a pseudo-labeling process and fine-tunes the CLIP model to make it more adaptable to masked images. MaskQCLIP[65] fine-tunes attention layer for new mask [CLS] tokens to make it more fit for mask object classification. These two methods attempt to enhance CLIP’s ability to focus on local features and exclusively fine-tune CLIP on specific downstream datasets, resulting in poor generalization ability beyond detection or segmentation tasks.

Another approach is to *change the input image* by simply cropping or masking the image to leave only the foreground object. ReCLIP [54] and OvarNet [7] crop the original image using bounding box from object proposal network [70] and are applied on Referring Expression Comprehension and Open Attribute Recognition tasks. MaskAdaptedCLIP [31] sets the background area to pure color in pixel space and uses the masked image as input for open-

vocabulary segmentation. However, the valuable context information is lost except for using complex post-process proposed in ReCLIP [54]. Some other approaches prompt the CLIP by modifying the input image, guiding CLIP to focus on the area of interest. For example, Red-Circle [52], FGVP [66] use a circle or mask contour to tell CLIP where to focus. Overall, the quality of these approaches that *change the original content of input image* is heavily contingent upon the symbols in CLIP’s pre-training dataset. Another limitation is directing modification of images causes a domain gap with CLIP pertaining images. Unlike previous approaches that rely on *segmentation* or *changing the input image*, our Alpha-CLIP incorporates an additional alpha channel, which does not change the image content and preserves the generalization performance (cf. Fig. 2).

Region-level image annotation. Existing CLIP models are pretrained on large-scale datasets like LAION-400M [49] and LAION-5B [50], while fine-grained mask-level labels are not available due to high manual labor costs. Recently, Kosmos-2 [41] introduced a pseudo-labeling pipeline that uses the pre-trained GLIP [29] model to automatically generate fine-grained pseudo-labels of region boxes and their associate expressions. By using this pseudo-labeling baseline, Kosmos-2 releases the GRIT dataset and equips multimodal model [18] with local perception capabilities. Similarly, the All-Seeing [59] project also generates fine-grained text labels via the pseudo-labeling pipeline. Meanwhile, the recent SAM [22] model is trained on massive vision modality data with strong zero-shot abilities for downstream tasks like box-to-mask conversion and automatic mask generation. These developments have made it possible to generate pseudo-masks with region captions at a large scale and have opened up the potential for greater adjustments to CLIP for region-level recognition. Therefore, We build upon GRIT[41] and SAM [22] to propose a method for generating RGBA region-text pairs from grounding data.

CLIP in MLLM. At the age of Multi-modal Large Language Models (MLLMs) [1–4, 18, 26, 28, 33, 34, 40, 41, 71], CLIP [43] has been widely used as the vision backbone for its semantic representative feature and promis-

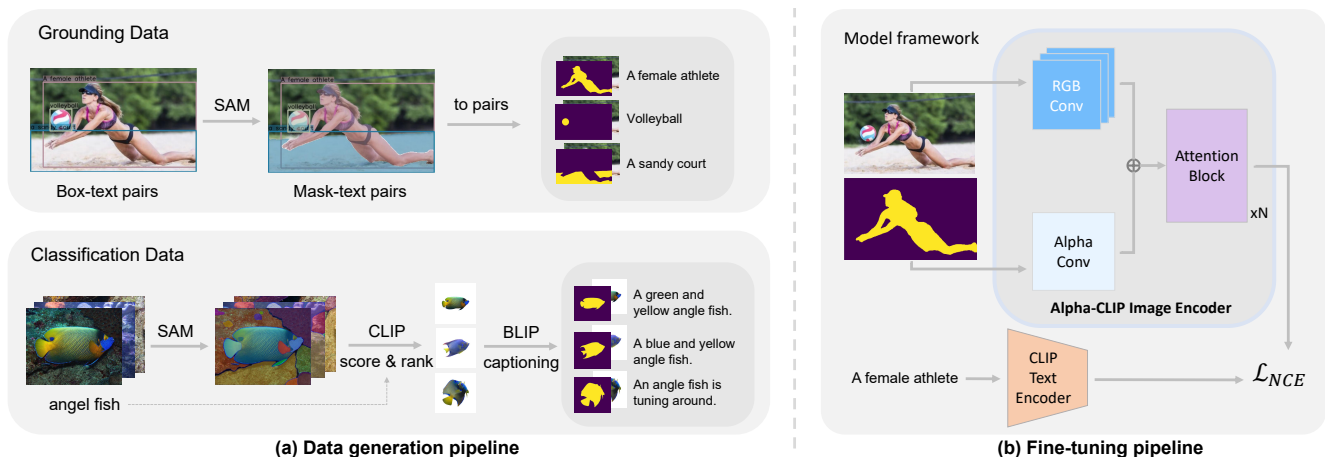


Figure 3. **The pipeline of our data generation method and model architecture.** (a) Our method generates millions of RGBA-region text pairs. (b) Alpha-CLIP modifies the CLIP image encoder to take an additional alpha channel along with RGB.

ing scalability. To make MLLM focus on the specific region, Kosmos-2 [41] uses millions of region-caption data to train the model with the guidance of box corner points. GPT4ROI [72] propose to apply the ROI Align [15] operator on the CLIP image feature to refer to the specific region. GLaMM [45] further adds an extra region encoder. Different from previous methods that only support box-level focusing and rely on training additional networks, our work achieves more fine-grained mask-level region focusing and merely uses the CLIP model.

CLIP in 2D image variation. CLIP image encoder is widely used in 2D image variation (*e.g.*, DALLE-2 [44], Diffusers [58] and IP-Adapter [68]) to achieve better quality or controllability. As for subject-driven image variation pioneered by DreamBooth [47], extraction of pure single object feature from the whole image is more important as the following method ELITE [60] proposes to use feature-level masking to eliminate background information to generate better subjects. Similarly, BLIP-Diffusion [27] uses text to extract the most relevant object features. All these subject-driven image variation methods require the image to have a single foreground object in the center of the image and cannot achieve variation by focusing on user-specified objects in more complex images while maintaining original context information. Such limitations highlight the importance of our Alpha-CLIP that enables subject-driven generation in complex scenes and achieves user-defined region focusing in image variation tasks.

CLIP in 3D generation. Some existing 3D object generation methods involve CLIP [43] model. In diffusion based 3D generation, Point-E [39] uses the point cloud diffusion model to generate the point cloud directly conditioned by the CLIP feature from a single view image or text. Another approach in the field of text-to-3D is pioneered by Dream Fields [20], which uses the CLIP model to provide supervision loss. Following works include PureCLIPNeRF [25], CLIP-Mesh [38], CLIP-Forge [48] and Dream3D [63] also

use CLIP image encoder to extract rendered image features. Our Alpha-CLIP can enhance CLIP in 3D object generation, enable Point-E with user-defined region focus ability and help optimization based text-to-3D models to yield high-quality generation results.

3. Method

This section describes the data pipeline and framework of Alpha-CLIP. As illustrated in Fig. 3, we first design a data pipeline to generate RGBA-region text pairs data (Sec. 3.1). Using our generated data, we then train our Alpha-CLIP with additional Alpha-channel inputs (Sec. 3.2).

3.1. RGBA Region-Text Pair Generation

To fine-tune the CLIP model with an additional alpha channel input, we first design a data generation pipeline (*cf.* Fig. 3a) to create millions of RGBA-region text pairs. Our pipeline consists of the following two components.

Grounding data pipeline. As depicted in the upper part of Fig. 3a, this branch is dedicated to generating region-text pairs, which include natural images with foreground alpha channels and corresponding referring expressions for specific regions. The natural images are from the GRIT dataset [41], which employs GLIP and CLIP to automatically extract labels of *box* region-text pairs. Building upon GRIT, we take a further step of generating *mask* region-text pairs. Specifically, we use SAM [22] to automatically generate high-quality pseudo-masks for each box region.

Classification data pipeline. As illustrated in the lower part of Fig. 3a, this branch is utilized for generating region-text pairs where the foreground objects are highlighted while the original background is removed. We employ the ImageNet [8] dataset for this purpose. Firstly, we use SAM to automatically generate several masks for each image in ImageNet. Subsequently, we crop the foreground object of each mask, center it, and enlarge it. CLIP is then used to

calculate scores with the corresponding class label of the image to which each mask belongs. Following this, we sort the masks by class based on their scores and select the top-ranked masks with the highest scores. Regarding the text component, to ensure that the caption for each mask is not merely the ImageNet [8] class label, we place the foreground object on a pure white background. Then we use BLIP-2 [28] to annotate these masks with captions. Finally, we merge the fine-grained ImageNet class label with the image-specific captions generated by BLIP-2 [28], resulting in millions of RGBA region-text pairs.

3.2. Alpha-CLIP

Model structure. Our Alpha-CLIP implements subtle structural modifications to the CLIP image encoder to preserve CLIP’s prior knowledge. In the CLIP image encoder’s ViT [11] structure, an RGB convolution is applied to the image in the first layer. As shown in Fig. 3b, we introduce an additional Alpha Conv layer parallel to the RGB Conv layer, which enables the CLIP image encoder to accept an extra alpha channel as input. The alpha channel input is set to range from [0, 1], where 1 represents the foreground and 0 indicates the background. We initialize the Alpha Conv kernel weights to zero, ensuring that the initial Alpha-CLIP ignores the alpha channel as input.

Training method. During training, we keep the CLIP text encoder fixed and entirely train the Alpha-CLIP image encoder. Compared to the first convolution layer that processes the alpha channel input, we apply a lower learning rate to the subsequent transformer blocks. To preserve CLIP’s global recognition capability for full images, we adopt a specific data sampling strategy during training. We set the sample ratio, denoted as $r_s = 0.1$ to occasionally replace our generated RGBA-text pairs with the original image-text pairs and set the alpha channel to full 1. Please refer to Appendix A.1 for ablation studies such as the number of unfreeze Transformer blocks and value of r_s .

Alpha-CLIP for downstream tasks. After the training, Alpha-CLIP possesses the capability to focus on a specified region and controlled editing. Alpha-CLIP can enhance CLIP’s performance on various baselines in a plug-and-play fashion, across various downstream tasks like recognition, MLLM, and 2D/3D generation (see Tab. 1 in Sec. 1).

4. Experiments

Data. We train Alpha-CLIP on RGBA region-text pairs using grounding data pipeline from GRIT-20m [41] for zero-shot ImageNet classification. We combine it with 460k RGBA region-text pair from ImageNet [8] using classification data pipeline to train Alpha-CLIP for other tasks including REC, OVD, region-level captioning, 2D image variation, and 3D generation. Ablation on data volume and mix-

Methods	ViT-B/16		ViT-L/14	
	Top-1	Top-5	Top-1	Top-5
Original CLIP [43]	66.48	88.90	73.48	91.60
MaskAdaptedCLIP [31]	57.86	79.12	63.50	86.34
Red Circle [52]	65.37	88.68	73.37	92.09
MaskCLIP* [75]	67.86	89.40	77.04	93.39
Alpha-CLIP(ours)	68.89	90.51	77.41	94.45

Table 2. **Zero-shot classification on ImageNet-S [12].** When given the foreground object on the alpha channel, our Alpha-CLIP significantly improves zero-shot classification and surpasses previous baselines such as MaskCLIP [75].

Model	Alpha Map	Top-1	Top-5
CLIP [43]	-	73.48	91.60
Alpha-CLIP	whole image	73.37	91.75
	rectangular box	75.62	93.34
	mask	77.41	94.45

Table 3. **Zero-shot classification on ImageNet-S [12] with different alpha map levels.** Alpha-CLIP is comparable to the original CLIP when the foreground mask is not available, and further boosts the performance with rectangular box or mask alpha maps.

ture of data are in Appendices A.2 and C

4.1. Alpha-CLIP in Image Recognition

Zero-shot image classification. We select the ImageNet-S [12] dataset for zero-shot classification analysis, which comprises 919 classes with semantic segmentation annotations selected from ImageNet-1k. We prepare the image-level semantic segmentation masks as the alpha channel input. We select representative baseline methods designed for making CLIP focus on the specific region: MaskCLIP [75], MaskAdaptedCLIP [31], and Red Circle [52]. Note that MaskCLIP is designed for mask generation rather than recognition. We make necessary modifications to MaskCLIP to adapt it for the recognition task (please refer to Appendix B for our implementation details). We use the mean of per-class accuracy as the evaluation metric.

Tab. 2 presents the zero-shot classification comparison on ImageNet-S *validation* set. This experiment effectively demonstrates that when provided with a foreground object mask through the alpha channel, our Alpha-CLIP generates visual features that are more focused on the foreground object, leading to better image-level classification compared to the original CLIP and other baseline approaches. It is worth noticing that Although MaskCLIP [75] achieves good results without needing to fine-tune the CLIP model, it is not directly compatible with methods that require the whole feature map instead of just the [CLS] token. This limitation is particularly relevant when considering methods like BLIP-2[28], BLIP-Diffusion [27], LLaVA [34] and PointE [39]. In contrast, our Alpha-CLIP is more general and can be applied to these approaches effectively.

We also evaluate Alpha-CLIP in scenarios where the

Method	RefCOCO			RefCOCO+			RefCOCog	
	Val	TestA	TestB	Val	TestA	TestB	Val	Test
CPT [67]	32.2	36.1	30.3	31.9	35.2	28.8	36.7	36.5
ReCLIP [54]	45.8	46.1	47.1	47.9	50.1	45.1	59.3	59.0
Red Circle [52]	49.8	58.6	39.9	55.3	63.9	45.4	59.4	58.9
Alpha-CLIP	55.7	61.1	50.3	55.6	62.7	46.4	61.2	62.0

Table 4. **Comparison with state-of-the-art on zero-shot REC.** We report top-1 accuracy (%). Replacing CLIP in ReCLIP [54] with Alpha-CLIP outperforms other zero-shot approaches on most datasets, including Red Circle [52], ReCLIP [54] and CPT [67].

foreground mask is unavailable. As shown in Tab. 3, when foreground prior is not available, we set alpha channel input to all one. We observe that the recognition ability of Alpha-CLIP (second row) remains on par with the original CLIP (top row). When provided foreground box (third row) or foreground mask (bottom row), Alpha-CLIP can significantly improve classification accuracy.

Zero-shot referring expression comprehension. In addition to the zero-shot image classification task, we also conducted experiments on zero-shot Referring Expression Comprehension (REC). zero-shot REC is the task of localizing objects in an image given a textual reference expression in a zero-shot manner. We follow previous works to select the RefCOCO [69], RefCOCO+ [69], and RefCOCog [36] datasets for evaluation. We select three representative approaches CPT [67], ReCLIP [54], and Red-Circle [52] as our baselines. We replace the CLIP model in this task with our Alpha-CLIP. Specifically, we use object proposals predicted by a pretrained detector [70] and employ SAM to obtain masks for each proposal. Instead of cropping the object by bounding box, we input the original image with an alpha map into our Alpha-CLIP. This modification has proven beneficial in preserving global contextual information as we find cropping only lead to worse result. Please refer to Appendix E for more implementation details.

As shown in Tab. 4, Alpha-CLIP achieves competitive zero-shot results on the REC task, surpassing ReCLIP and RedCircle by an average of 6.8% and 3.0% accuracy across RefCOCO, RefCOCO+ and RefCOCog benchmarks. The experimental results demonstrate that Alpha-CLIP enhances CLIP’s ability to focus on the relevant region and such enhancement is also beneficial for the REC task that requires image-text understanding and reasoning capabilities.

Open vocabulary detection. The Open-Vocabulary Detection (OVD) task aims to detect novel classes that are not available during training. Detic [76] is a pseudo-labeling baseline that proposes to use the ImageNet dataset for OVD. Specifically, Detic first trains the detector on the base classes of LVIS [14], then uses the detector to generate pseudo bounding boxes on ImageNet. These pseudo boxes may cover the novel objects and help improve the detector’s performance in novel classes. Such a semi-supervised pipeline is not data-efficient and Detic uses 1.2M images

Dataset	mAP _{novel}	mAP
Detic-ImageNet	24.6	32.4
MaskImageNet (ori CLIP)	27.9	32.5
MaskImageNet (Alpha-CLIP)	28.6	32.9

Table 5. **Open-vocabulary detection on OV-LVIS [14].** Using MaskImageNet and our Alpha-CLIP can significantly improve mAP_{novel} on novel classes.

from ImageNet in the OV-LVIS [14] benchmark.

To demonstrate the effectiveness of Alpha-CLIP on OVD, we transfer the top-ranked ImageNet (460K) into a collection, dubbed as MaskImageNet. Specifically, we apply our data generation pipeline, as detailed in Sec. 3.1 to generate pseudo-labeled bounding boxes and foreground masks for each image. We replace the ImageNet used in Detic’s pseudo-labeling steps with our MaskImageNet. We also remove the background category loss and adjust the blending ratios of LVIS and MaskImageNet. Experimental results are presented in Tab. 5. Compared to the Detic baseline using ImageNet (top row), The second row demonstrates that using our MaskImageNet already enhances OVD capabilities. Furthermore, our Alpha-CLIP (bottom row) further improves OVD performance. Remarkably, our method (460K in MaskImageNet) is more data efficient than Detic (1.2M in ImageNet).


4.2. Alpha-CLIP in MLLM

We replace CLIP used in BLIP-2 [28] and LLaVA-1.5 [33] with our Alpha-CLIP to make MLLM directly focus on user-defined region in vision-language tasks such as region level captioning and VQA.

Region level captioning. As shown in Fig. 4, simply replacing CLIP with Alpha-CLIP enables MLLM to generate captions more focused on user-defined areas. In the third row cases about the telephone and mushroom, the original CLIP generates the wrong caption. This error may arise due to the CLIP vision feature mixing different objects and their properties in images with too many foreground objects. Alpha-CLIP guides MLLM to generate the correct caption by providing the area to focus on. More visualization are in Appendices G.1 and G.2. We also visualize the CLIP attention map marked in the upper right to confirm our findings. More visualizations and implementation details of the attention map are in Appendix G.5.

Besides qualitative results, we also provide the quantitative region level captioning results of Alpha-CLIP with LLaVA-1.5 [33] on Visual Genome [23] and RefCOCog [36]. We fine-tune Alpha-CLIP+LLaVA-1.5 [33] with vicuna-7b [57] on these datasets with the same setting in [45, 72] and task prompts in [72] is adopted. Alpha-CLIP image encoder is kept frozen with LLM fine-tuned to adapt region caption format. Results are shown in Tab. 6. Alpha-CLIP+LLaVA-1.5 achieves competitive results over the baseline methods, even surpassing previous

Alpha-CLIP in BLIP-2



Alpha-CLIP in LLaVA

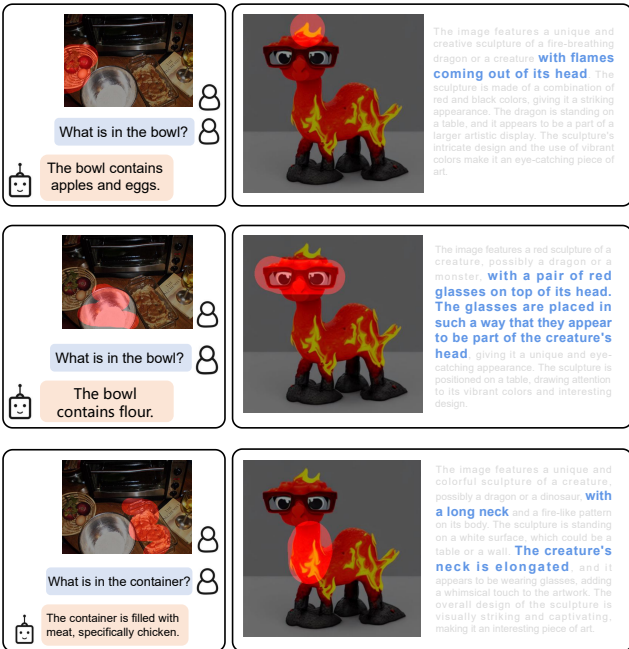


Figure 4. **Some results of Alpha-CLIP used in MLLMs** The upper half is image captioning result with BLIP-2 [28]. The first column is the original CLIP generated captions. Other columns represent the outcomes of Alpha-CLIP with highlighted region marked in red. The lower half is region focused VQA and image captioning result with LLaVA-1.5 [33].

expert models like GPT4ROI [72] and GLaMM [45] with ROI Align [15] or additional region encoder structure pre-trained on a large volume of region-text pairs.

Model	RefCOCOg		Visual Genome	
	METEOR	CIDEr	METEOR	CIDEr
GRIT [61]	15.2	71.6	17.1	142.0
Kosmos-2 [41]	14.1	62.3	-	-
GPT4RoI [72]	-	-	17.4	145.2
GLaMM [45]	16.2	105.0	18.6	157.8
Alpha-CLIP+LLaVA [33]	16.7	109.2	18.9	160.3

Table 6. **Performance of Alpha-CLIP in region level captioning.** We report METEOR and CIDEr metrics on Visual Genome and refCOCOg Datasets.

Region based VQA. MLLM can chat with users with simple reasoning. In this scenario, alpha channel input can act as the visual prompt defined by the user to highlight specific regions of interest. As shown in Fig. 4 and Fig. 1, the user can simply use stroke to tell MLLM the referring object or regions to focus on. More visualization results of VQA with Alpha-CLIP are in Appendix G.2.

4.3. Alpha-CLIP in image variation.

Alpha-CLIP can be used in most image variation models that use CLIP image encoder [27, 44, 58, 60, 68]. For example, BLIP-Diffusion bridges CLIP [43] and stable-diffusion [46] with Q-former to generate and edit 2D images controlled by text. Since BLIP-Diffusion [27] is a typical method that maintains subject information, we use BLIP-Diffusion to demonstrate the effectiveness of Alpha-CLIP. By introducing Alpha-CLIP, we can add an additional set of vision prompts to allow the model to focus on specified regions for 2D generation. We replace the ViT-L/14 model in BLIP-Diffusion [27] with Alpha-CLIP while keeping the other parts unchanged. We set the empty text prompt to make results irrelevant with semantics. As shown in Fig. 1, Alpha-CLIP with alpha map on highlighted areas enables BLIP-Diffusion to generate region-focused results. We also compare our Alpha-CLIP with other CLIP region-focused approaches such as image cropping, pixel-level image masking, red circle, and feature-level masking (Please refer to Appendix D for implementation details). As shown in Fig. 2, image cropping can not solve the occlusion problem. The red-circle solution will change the image content. Neither pixel-level nor feature-level masking can convey original background information. In contrast, our Alpha-CLIP that prompts CLIP with fine-grained region mask solves the above problems and generates cleaner results while maintaining original background information. More visualizations are in Appendix G.3

4.4. Alpha-CLIP in 3D Object Generation.

Alpha-CLIP can also apply to 3D Object Generation. We test it in two different approaches: 1) Point-E [39] that is a




Figure 5. **Results of Alpha-CLIP in 3D generation.** The top part shows 3D point clouds generation using Point-E [39]. The first row displays objects generated by the original CLIP. The second row illustrates the results of Alpha-CLIP, with highlighted areas in red. The bottom part shows 3D objects generated by PureCLIPNeRF [25]. The CLIP model is replaced with Alpha-CLIP, and tests are conducted with and without background augmentation.

diffusion-based method for image-to-3D, and 2) PureCLIPNeRF [25] that is an optimization-based approach for text-to-3D.

Alpha-CLIP in Point-E. Point-E [39] can achieve image-to-3D through conditioning diffusion model with CLIP image feature. We replace the CLIP ViT-L/14 image encoder of the Point-E base-40M model with our Alpha-CLIP. We demonstrate that Alpha-CLIP is helpful in two cases: 1) When Point-E generates the point cloud with some parts missing, users can highlight the missing part in the condition image to remind the diffusion model to pay more attention to that part and fix this missing parts problem. 2) Users can highlight the part that needs to be emphasized on the 2D image. Point-E will spend more points on the highlighted part (with 1024 points in total in the base model). The results are shown in Fig. 5 with more results in Appendix G.4.

Alpha-CLIP in PureCLIPNeRF. We input the rendered images with alpha channels obtained from density integration of NeRF [37] into Alpha-CLIP. When optimizing the

object with Alpha-CLIP, the gradient can flow back from the alpha channel to help generate better results. As shown in Fig. 5, we find that PureCLIPNeRF generates objects that closely align with the provided textual prompts (especially bolded text) in terms of shape and color when replacing CLIP with Alpha-CLIP. Furthermore, there is an enhancement in the overall coherence of the generated objects, coupled with notable aesthetic qualities. We attribute this phenomenon to Alpha-CLIP’s enhanced capability in optimizing density parameters of 3D representations directly and focusing only on the foreground area, which helps to generate an object that is more coherent and closely matches the input text.

Background augmentation in PureCLIPNeRF [25] inherited from Dream Fields [20] is a vital step to improve the consistency of objects, making them less diffuse compared to the first column in Fig. 5. However, this process is time-consuming as each augmented image has to go through CLIP to get optimization direction. We thus test the capabilities of Alpha-CLIP without background augmentations. Results are presented in the second column of Fig. 5. We observe that in most cases, using Alpha-CLIP without background augmentation produces objects that are clearer and better aligned with the given text than the original CLIP with 2x faster speed. Quantitative results and More visualizations are in Appendices F and G.4

5. Limitation and Future Direction

While Alpha-CLIP demonstrates effective performance in various scenarios requiring region focus, its current structure and training process limit its capability to focus on multiple objects or model relationships between different objects. Furthermore, the current training methodology restricts the alpha channel from generalizing beyond intermediate values, apart from the binary values of 0 and 1. As a result, users are unable to specify the amplitude of attention. Another limitation both lays in our Alpha-CLIP and original CLIP is low resolution, which hinder the way for Alpha-CLIP to recognize small object. We plan to address these limitations in future work and expand the CLIP input resolution. We believe these future directions are pathways to augment Alpha-CLIP’s abilities and broaden its utility across diverse downstream tasks.

6. Conclusion

In this work, We propose the Alpha-CLIP model, which introduces an additional alpha channel to specify the regions of interest. Trained on millions of RGBA region-text pairs, Alpha-CLIP not only exhibits excellent region-focus capabilities but also ensures its output space remains consistent with the original CLIP model. This consistency allows seamless replacement in various downstream applications of CLIP. We demonstrate that when prompted with specific regions of interest, Alpha-CLIP shows improved

zero-shot recognition abilities and verifies its usefulness in many downstream tasks. The applications of CLIP extend far beyond the scope of this article. We hope that Alpha-CLIP will be applicable in more scenarios when foreground regions or masks are available.

A. Training Detail

A.1. Hyperparameter

Basic hyperparameters. The training process utilizes a batch-size of 4096 for all scales of CLIP models. We use 8 A100-80G GPUs for ViT-B/16, 64 GPUs for ViT-L/14, and 128 GPUs for ViT-L/14@336px. The training process utilizes mixed-precision float16 for acceleration. The temperature coefficient τ for CLIP is fixed to the value obtained after the completion of the original CLIP training. The optimizer chosen is AdamW [35] with a weight decay of $2e-2$. Regarding learning rates, the learning rate for the convolutional kernels accepting alpha channel input is set to $2e-4$, while the rest of the layers have a learning rate of $2e-6$, employing a cosine learning rate scheduler. For GRIT-1m, the training lasts 6-8 epochs, whereas GRIT-20m is trained for 2 epochs.

Whole image sample ratio. Due to our desire to preserve the original CLIP’s recognition ability for the entire image, in the training of Alpha-CLIP on GRIT [41], we sample a portion of RGBA-text pairs and set the alpha channel to all 1 (indicating region over the entire image). The text is replaced with the original full-image caption. We use the ViT-B/16 model and train on GRIT for 4 epochs, varying the sampling ratio. Zero-shot classification on Imagenet-S is used as the evaluation metric. The experimental results, as shown in Tab. 7, indicate that training without sampling a proportion of the entire image-text pair performs worse than choosing a small number of images that require full-image attention. However, excessively emphasizing full-image attention during training significantly impairs the model’s capability. Based on these results, a sample ratio of 0.1 is chosen for the experiments in the main text.

Whole image perception setting. In accordance with the definition of transparency in the 2D image, setting alpha to

sample ratio r_s	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
top1-Acc	68.06	68.25	67.87	67.71	67.83	67.74	67.37	66.87	66.39	64.94	63.96

Table 7. **Sample ratio search experiment.** We search sample ratio with a step of 0.1. Test metric is zero-shot classification top1 accuracy on ImageNet-S [12]. As we find $r_s = 0.1$ produce best result.

unfreeze block nums	0	2	4	6	8	10	12	full-tuning on ori CLIP
top1-Acc	63.61	64.73	65.63	66.59	67.09	68.07	68.27	66.52(+0.04)

Table 8. **Number of unfreeze block search experiment.** We search number of learnable Transformer block number. Test metric is zero-shot classification top1 accuracy on ImageNet-S [12]. As we find that unfreeze the whole CLIP image encoder generate the best result.

	test with gt mask	all_0	all_1	ori_clip
whole_1	77.41	72.53	73.37	73.48
whole_0	74.99	73.45	73.27	

Table 9. **Different strategy for whole image perception experiment.** Test metric is zero-shot classification top1 accuracy on ImageNet-S [12]

all 1 is indicative of the requirement for CLIP to focus on the entire image. However, due to the absence of bias in the first layer convolution of CLIP, which consists only of weights, an input with an alpha channel of all 0 maintains CLIP’s original state, contrary to the definition of image transparency. To determine an optimal approach, we conducted training and testing under different configurations. During training, we utilized image-text pairs with the entire image set to all 0 and all 1, each with a sample ratio of 0.1. During testing, we cross-validated the classification accuracy with alpha channels set to all 0 and all 1. Experiment results are presented in Tab. 9. Our observations indicate that configuring the training and inference with the alpha channel set to all 1 achieves the best perception performance. Therefore, we consistently adopt this configuration in the main section of the paper, utilizing an all-1 alpha input when Alpha-CLIP is required to focus on the entire image.

Unfreeze block number. We search through the number of layers to unfreeze when training Alpha-CLIP on RGBA-text pairs that create the best result. we test on ViT-B/16(with 12 attention blocks) and train on GRIT-1m [41] for 4 epochs. We select the zero-shot classification top1 accuracy on ImageNet-S [12] as our test metric. The result is shown in Tab. 8. We also test full finetuning of original CLIP without alpha channel on GRIT [41] dataset, which only gets negligible improvement, proving that the improvement contributes to the input focus region through alpha channel instead of training data.




Figure 6. **Zero-shot ImageNet-S [12] classification accuracy w.r.t training data volume.** two different scale ViT model of CLIP [43] are tested.

We unfreeze the number of transformer blocks from 0 to 12 (full) with a step of 2. Results show that as the number of unfrozen blocks increases, the classification accuracy increases steadily. We also test LoRA [17], but it does not work well compared with full model finetuning. So we consider full model tuning in the main part of the paper.

A.2. Ablation on Data Volume

We examine the efficacy of data volume in enhancing the training of robust models through an ablation study. Our ablation involved the training of ViT-B/16 and ViT-L/14 using RGBA region-text pairs, with data quantities ranging from 1k to 10M. We use the zero-shot top-1 accuracy on ImageNet-S as our evaluation criterion. As illustrated in Fig. 6, increasing data volume corresponds to a concurrent improvement in the model’s classification accuracy. Notably, larger ViT models exhibited a more substantial performance boost in comparison to their smaller counterparts throughout this process.

B. Different implementation of MaskCLIP

To the best of our knowledge, two methods propose to use the attention mask to guide the CLIP visual model to pay more attention to the foreground area. We test these two methods respectively. Because these two methods can only do masking at the feature level as [H, W] ([14, 14] for ViT-B/16, [16, 16] for ViT-L/14), we first do max pooling on the binary mask M to make it match the size of the feature

model		ViT-B/16	ViT-L/14
Original CLIP	top1	66.48	73.48
	top5	88.90	91.60
MaskCLIP [9]	top1	53.61	65.28
	top5	76.47	87.25
Alpha-CLIP	top1	68.89	77.41
	top5	90.51	94.45

Table 10. **Zero-shot classification on ImageNet-S[12].** Comparison using method proposed in [9].

map.

$$m = \text{MaxPooling}(M) \quad (1)$$

Mask Area guided last attention. As [75] proposes using 1×1 conv layer to get feature space level 2D semantic classification map, we use the same idea of the attention mask to set [cls] token only to calculate attention with foreground area patches. In other words, we use m to guide the last attention calculation. We report this result in the main part of this paper as it shows better results than the original CLIP.

Relative Mask Attention. This method is proposed in [9], and is used in [62], which introduces “Mask Patch Tokens” to do the same attention as [CLS] token but only attach to those patches that contain foreground area. We test this method but it does not produce good results on ImageNet-S [12] as shown in Tab. 10 as it is used in the segmentation task to classify each semantic mask(area) of a whole image, but ImageNet-S [12] only cares about a single prominent foreground object in most cases. So we do not report this result in the main part of this paper.

C. Effectiveness of Classification Data

While Grounding data holds more promising prospects in the future, especially with the advent of more powerful grounding and segmentation models, we demonstrate that, at the current stage, leveraging large-scale manually annotated classification datasets like ImageNet [8] and constructing RGBA region-text pairs using the pipeline shown in Fig. 3 still significantly benefits Alpha-CLIP in achieving enhanced Region-Perception capabilities.

C.1. Zero-shot Classification on COCO

In addition to natural image classification tests, there are scenarios where there is a need to crop or mask objects in images[7] [54] [31]. Therefore, we conducted classification tests for Alpha-CLIP in such scenarios using the validation set of the Instance-COCO [32] dataset, which consists of 80 classes. We cropped objects using ground-truth bounding boxes and enlarged them by 1.5 times (referred to as

model		ViT-B/16	ViT-L/14
masking	CLIP	49.42	54.43
	Alpha-CLIP _g	49.27	56.45
	Alpha-CLIP _{g+c}	53.39	58.84
no masking	CLIP	64.21	67.65
	Alpha-CLIP _g	61.57	67.44
	Alpha-CLIP _{g+c}	71.08	77.56
ImageNet-S top1	CLIP	66.48	71.48
	Alpha-CLIP _g	68.89	77.41
	Alpha-CLIP _{g+c}	69.40	77.80

Table 11. **Zero-shot classification results on COCO.** Our Alpha-CLIP also achieve significant improvement on zero-shot Instance-COCO [32] classification tasks.

Data	RefCOCO			RefCOCO+			RefCOCog	
	Val	TestA	TestB	Val	TestA	TestB	Val	Test
GRIT-1M	56.1	63.4	48.9	55.1	62.6	45.1	60.3	60.6
GRIT-1M + IN	55.7	61.1	50.3	55.6	62.7	46.4	61.2	62.0

Table 12. **Zero-shot REC results of Alpha-CLIP with different pretraining data.** We compare the results of using only grounding data with adding classification data. We report top-1 accuracy (%).

COCO crop). We conduct tests in two scenarios: masking (setting the background to a solid color) and no masking (using the original background). To prevent results from being dominated by the most frequent classes, we use the mean of per-class accuracy as the evaluation metric. To ensure that Alpha-CLIP is adapted to images with backgrounds replaced by solid colors, we incorporate object-centric image data (from the lower branch of Fig. 3 into the training data for this scenario. This data is generated from the top 460k RGBA-region text pairs auto-generated from ImageNet-21k [8], and we include it in the pairs generated from GRIT-1M [41] as the training dataset for Alpha-CLIP. Results are shown in Tab. 11. We compare it with the baseline method trained on GRIT-1m only and find a huge improvement for cropped image classification. We also test its classification accuracy on ImageNet-S [12], and the result even surpasses models trained on GRIT-20m. We contribute this to the human annotation of fine-grained class labels of ImageNet [8] dataset.

C.2. Different version Alpha-CLIP in REC

To investigate the effectiveness of the classification data on REC, we conduct experiments comparing Alpha-CLIP pretrained solely on the grounding data with a combination of classification and grounding data. We use an ensemble of ViT-B/16 and ViT-L/14 backbones, with grounding data sourced from GRIT-1M [41] and classification data from ImageNet-21k [8]. As shown in Tab. 12, on the majority of benchmarks, using classification data yields better results compared to models that are not pretrained with it.

D. Other CLIP masking baselines

There are simple ways to make CLIP focus on user-specified regions without modifying the weights of the CLIP model. We test two possible approaches here. Namely Image-level masking and feature-level masking. We test on these simple baselines and make comparisons with Alpha-CLIP. As shown in Fig. 7. It is worth noticing that we only draw structure using the Q-former proposed by BLIP-2. But they can also be adapted to other VL systems that use CLIP image encoder as the visual backbone like LLaVA [33, 34] and miniGPT4-v2 [6].

Image Level Masking means simply masking the background region by directly setting these regions to pure color. We choose the color same as MaskAdaptedCLIP [31].

Feature Level Masking is another method that applies masking on the feature level. As shown in Fig. 7, we use max pooling to downsample the image level mask to feature level coarse-grained mask, and use this mask to do element-wise product to set the features that belong to the background to zero. This method has been proven useful in Ellite [60] when the object is in the center of the image and occupies a large space.

Results of the two masking methods are shown at the top of Fig. 7. We use the same settings as in the main section of the paper. BLIP-2 [28], CLIP-L/14+flant5xl is used for captioning, BLIP-Diffusion [27], CLIP-L/14+stable-diffusion [46] is used for Image generation. The first column represents the original image and the area that needs to be focused, the second column represents the results of using image-level masking, the third column represents the results of using feature-level masking, and the fourth column represents the results of our Alpha-CLIP. The first four lines are image captioning results, and the last four lines are image variation results. As can be seen from Fig. 7, image-level masking will lose the context information of the object, causing its semantics to be incorrect or blurred and cannot produce good results; while feature-level masking can sometimes produce better results, but rough masking directly on the feature level may cause unpredictable behaviors, such as generating pictures with completely irrelevant semantic information, or being dominated by the semantics of the main objects in the picture. In contrast, Alpha-CLIP can produce better results because it is pretrained on millions of RGBA-text pairs. These two feature masking methods also destroy the features of other areas of the image to a greater extent and completely lose information about the other part of the image, and therefore fail in simple reasoning problems that need to involve the relationship between objects and the environment. In the meantime, Alpha-CLIP can highlight the region that needs to be focused on with features of the remaining areas in the image better preserved.

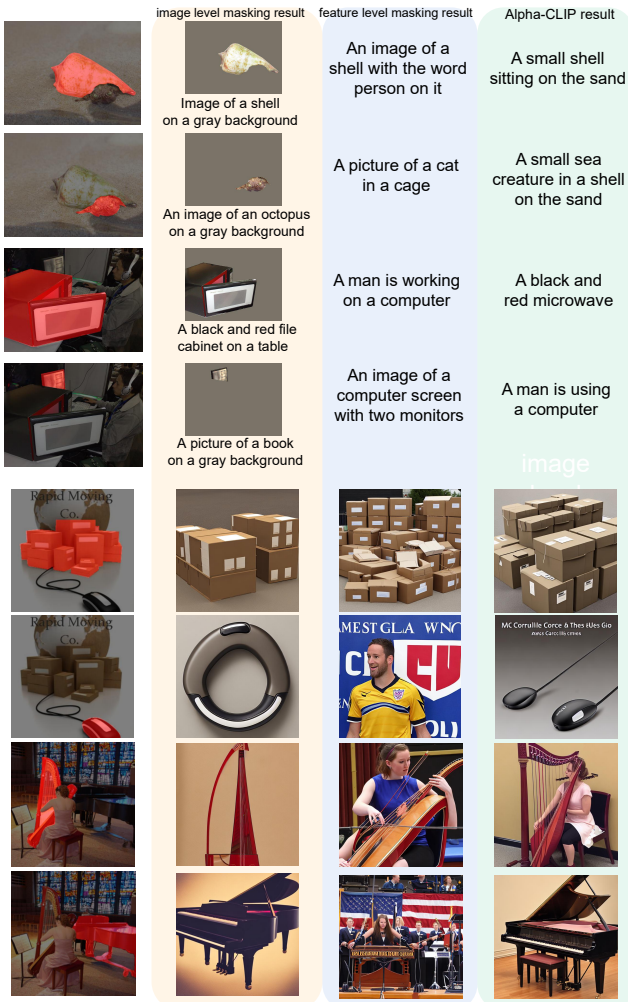
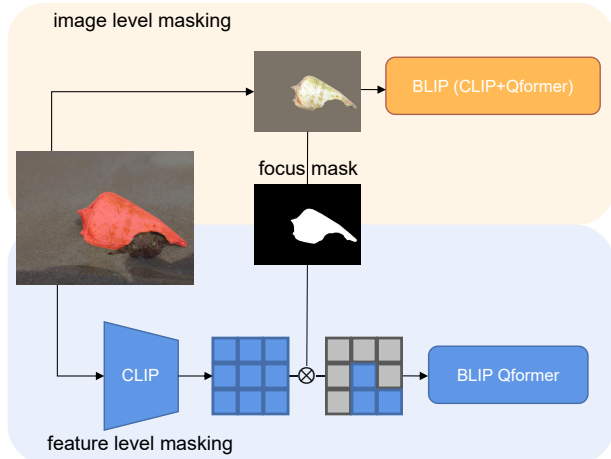


Figure 7. **Two baselines of image level masking and feature level masking and their comparison with Alpha-CLIP.** It is worth noticing that we use BLIP-2 [28] structure with Q-former as presented in the baseline pipeline. Alpha-CLIP and these two masking approaches can also adapt to structures that only have the projection layer, like LLaVA [34] and miniGPT-v2 [6]




Figure 8. **Model pipeline of Alpha-CLIP in Zero-shot REC Tasks.** We compared our model (as illustrated in a detailed flowchart in the lower part) with two other baselines [52, 54] (represented by a concise flowchart in the upper part).

PP	RefCOCO			RefCOCO+			RefCOCOg	
	Val	TestA	TestB	Val	TestA	TestB	Val	Test
B	56.8	63.7	49.4	56.2	63.6	45.9	59.9	61.7
B C	57.0	62.8	50.6	57.1	63.7	48.0	64.0	64.1
B C G	57.0	63.0	51.0	56.9	64.0	48.5	63.6	64.3

Table 13. **Zero-shot REC results of Alpha-CLIP with different image preprocessing methods.** We make comparisons across RefCOCO, RefCOCO+, and RefCOCOg datasets. **PP**: Preprocessing methods. “B” denotes original input and blurring operation. “C” denotes cropping. “G” denotes grayscaling.




Figure 9. **Visualization of different image preprocessing operations.** In our basic approach, we only utilize the original image and blurring. Additionally, we supplement the process with cropping and grayscaling operations.

E. Zero-shot REC with Alpha-CLIP Implementation Details

In Fig. 8 (c), we provide a detailed illustration of the Alpha-CLIP model pipeline in zero-shot Referring Expression Comprehension. We also compare our architecture with ReCLIP [54] and RedCircle [52], highlighting the differences and advantages. ReCLIP employs cropping and blurring operations to isolate image regions, which are obtained from box proposals from a detector [70]. However, as shown in Fig. 8 (a), cropping results in losing relationship

Method	Res+Iter	R-Precision	Time
PureCLIPNeRF †	168 ² +10k	85.62	~34min
α -PureCLIPNeRF	168 ² +10k	88.89	~36min

Table 14. **Quantitative Results of 3D Generation.** We compare the R-Precision of PureCLIPNeRF [25] model using original CLIP and Alpha-CLIP, as well as the time cost to generate a single object. † indicates our reimplementation.

information between objects, leading to decreased performance. While RedCircle draws red circles on specific regions across the entire image to avoid this issue, its prompt is still coarse-grained and it alters the original image, as presented in Fig. 8 (b). We use SAM [22] to generate fine-grained alpha maps with the aforementioned box proposals. We input both alpha maps and original or blurred images into Alpha-CLIP and calculate similarity with referring expressions. The remaining steps closely align with [54]. Our basic approach ensures the complete input of images and utilizes preprocessing methods as few as possible. It is efficient and achieves excellent performance across different benchmarks, as demonstrated in Tab. 4.

In addition to the preprocessing operations in our basic approach (original image and blurring), we further explore the cropping and grayscaling operations depicted in Fig. 8 (a) and (b). More specifically, for the blurring operation, our hyperparameter, namely the standard deviation (σ), is set to $\sigma = 100$ [54]. For the cropping operation, we pad the cropping box to be a square and fill the background at the image level with zeros (i.e., black color), as shown in Fig. 9. We use an ensemble of ViT-B/16 and ViT-L/14 Alpha-CLIP backbones to test, which are trained on GRIT-20M. The results, as shown in Tab. 13, indicate the additional benefits of incorporating these operations. This underscores the strong adaptability of our model, demonstrating its ability to adapt to diverse image inputs.

F. Quantitative Results of Neural Field Optimization based 3D Object Generation

We evaluate the quantitative results of Alpha-CLIP in PureCLIPNeRF, using the same test method proposed in Dreamfields [20], which includes 153 text prompts related to the COCO dataset. We measure the generated results using CLIP R-Precision. Under the same setting proposed in [25], we use the Alpha-CLIP ViT-B/16 model to optimize the generated object and compare our method with the original CLIP. We test R-Precision using CLIP ViT-B/32. The results are presented in Tab. 14, which shows our Alpha-CLIP can generate better objects than the original CLIP with negligible extra time consumption (test on V100-32G GPU).

G. More Qualitative Result Visualization

G.1. Region-focused Image Captioning

As described in Sec. 4.2, we replace original CLIP ViT-L/14 in BLIP-2 [28] with Alpha-CLIP without any post fine-tuning to generate captions. More results are shown in Fig. 10

G.2. Region-focused VQA and Detailed Image Description

As described in Sec. 4.2, we replace original CLIP ViT-L/14-336px in LLaVA-1.5 [33] without any post fine-tuning. Results are shown in Fig. 11

G.3. Region-focused Image Variation

As described in Sec. 4.3, we replace the original CLIP ViT-L/14 used in BLIP-Diffusion to make the condition image feature mainly focus on the user-specified area while maintaining background information. Results are shown in Fig. 12.

G.4. 3D Object Generation

Using the same setting in Sec. 4.4 Diffusion based object generation based on Point-E [39] base-40M are shown in Fig. 13, where Alpha-CLIP can achieve user-defined area focus to rectify missing part or emphasizing specific part. Neural field optimization based object generation based on PureCLIPNeRF [25] using ViT-B/16 for object optimization is shown in Fig. 14, where Alpha-CLIP generally generates better objects than original CLIP with or without background augmentation.

G.5. Attention map in Alpha-CLIP

Follow the spirit of DINO [5], we visualize Alpha-CLIP attention maps to check whether Alpha-CLIP pays more attention to user-defined highlighted areas at feature grid space. We check the attention map of [CLS] token in the last transformer block in the vision encoder. The model used for visualization is ViT-L/14 with 16 heads self-attention. For a fair comparison, we use the 5th and 16th heads attention maps for visualization, as we find that these two feature maps are most distinguishable among 16 heads. Results are shown in Fig. 15. This visualization verifies that Alpha-CLIP pays more attention to the area to focus on and more importantly, with no damage to the 2D location information preserved in the feature location of the original CLIP [43].









Figure 10. More results of Alpha-CLIP used in BLIP-2 [28]. The first row per three is the original BLIP-Diffusion generated images. Other rows represent the outcomes of Alpha-CLIP with highlighted regions marked in red. It is worth noticing that although we use ground truth mask as the highlighted region, Alpha-CLIP can actually take any region as alpha input. Users can easily get high-quality masks with the help of SAM[22].





What is in the bottle? 


 The bottle contains beer.


But I want to drink water now, what should I do? 


 There are two water bottles on the table, so you can choose one of them to drink water.


Where are they exactly? 


 The two water bottles are located on the table, next to the beer bottle.





What is this man doing in the image? 


 The man is playing baseball, specifically as a catcher, and is in the process of catching a ball during a game.





What's the meaning of this sign? 


 This sign indicates that there is a parking area nearby, possibly a designated parking zone.





What is this man doing in the image? 

 The man is standing at home plate, holding a baseball bat, and preparing to swing at an incoming pitch during a baseball game.




What's the meaning of this sign? 


 The sign in the image is a street sign that reads "Gladys." This sign indicates the name of the street where the pole is located.




The image features a statue of a woman holding scales, which is situated on top of a building. The statue is positioned in front of a large American flag, which is flying high on a pole. In the background, there is an airplane flying overhead, adding a sense of motion and activity to the scene. The combination of the statue, the flag, and the airplane creates a dynamic and visually interesting composition.



The image features a large clock mounted on the side of a building, with a statue of a woman holding scales nearby. The clock is prominently displayed, and it appears to be a central feature of the building. In the background, there is an airplane flying, adding a sense of motion and activity to the scene. The combination of the clock, statue, and airplane creates an interesting and dynamic visual composition.



The image depicts a group of people riding in a horse-drawn carriage on a street in front of the U.S. Capitol building. The carriage is being pulled by two horses, and there are several people visible in the scene, including a man riding the carriage and others standing nearby.



The image features a horse-drawn carriage with a man wearing a top hat driving the carriage. There are several people in the scene, including a man sitting on the carriage, a woman sitting on the back of the carriage, and other individuals standing nearby. The carriage is moving down a street, and the people appear to be enjoying the ride.

Figure 11. More results of Alpha-CLIP used in LLaVA-1.5 [33]. All cases shown here are made simply by replacing the original CLIP of LLaVA-1.5 [33] with a plug-in Alpha-CLIP without further tuning. Alpha-CLIP can achieve region-based VQA and region-focused detailed image descriptions.




Figure 13. **More results of Alpha-CLIP used in Point-E[39].** In each example, the results in the first row are 3D point clouds generated by the original CLIP, while the results in the second row are 3D point clouds generated with highlighted areas in red under the guidance of Alpha-CLIP. The left part shows Alpha-CLIP’s ability to rectify missing parts, and the right part shows emphasizing specific areas using Alpha-CLIP.




Figure 14. **More results of Alpha-CLIP used in PureCLIPNeRF[25].** In each example, the results in the last two columns are 3D objects generated with PureCLIPNeRF under the guidance of Alpha-CLIP and the original CLIP, while the results in the first two columns are objects generated by them respectively but without background augmentations.




Figure 15. **Alpha-CLIP Attention map visualization.** Last transformer block attention map of [CLS] token with other patch tokens. each first line per four is from original CLIP [43] and the other three lines are from Alpha-CLIP with user-defined focus regions marked in red. Prompting with region need focusing, Alpha-CLIP will focus on the part accordingly without compromising the original object location in feature grid.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [3](#)
- [2] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [3] Jinze Bai, Rui Men, Hao Yang, Xuancheng Ren, Kai Dang, Yichang Zhang, Xiaohuan Zhou, Peng Wang, Sinan Tan, An Yang, et al. Ofasys: A multi-modal multi-task learning system for building generalist models. *arXiv preprint arXiv:2212.04408*, 2022.
- [4] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. [2](#), [3](#)
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. [13](#)
- [6] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: large language model as a unified interface for vision-language multi-task learning, 2023. [11](#), [12](#)
- [7] Keyan Chen, Xiaolong Jiang, Yao Hu, Xu Tang, Yan Gao, Jianqi Chen, and Weidi Xie. Ovaret: Towards open-vocabulary object attribute recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23518–23527, 2023. [2](#), [3](#), [10](#)
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [4](#), [5](#), [10](#), [11](#)
- [9] Zheng Ding, Jie Wang, and Zhuowen Tu. Open-vocabulary universal image segmentation with maskclip. In *International Conference on Machine Learning*, 2022. [3](#), [10](#)
- [10] Xiaoyi Dong, Jianmin Bao, Yinglin Zheng, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Maskclip: Masked self-distillation advances contrastive language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10995–11005, 2023. [2](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [5](#)
- [12] Shanghua Gao, Zhong-Yu Li, Ming-Hsuan Yang, Ming-Ming Cheng, Junwei Han, and Philip Torr. Large-scale unsupervised semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 2022. [5](#), [9](#), [10](#), [11](#)
- [13] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *International Conference on Learning Representations*, 2021. [2](#)
- [14] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. [6](#)
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [4](#), [7](#)
- [16] MD Zakir Hossain, Ferdous Sohel, Mohd Fairuz Shiratud-din, and Hamid Laga. A comprehensive survey of deep learning for image captioning. *ACM Computing Surveys (CSUR)*, 51(6):1–36, 2019. [2](#)
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. [10](#)
- [18] Shaohan Huang, Li Dong, Wenhui Wang, Yaru Hao, Saksham Singhal, Shuming Ma, Tengchao Lv, Lei Cui, Owais Khan Mohammed, Qiang Liu, et al. Language is not all you need: Aligning perception with language models. *arXiv preprint arXiv:2302.14045*, 2023. [2](#), [3](#)
- [19] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip. https://github.com/mlfoundations/open_clip, 2021. [2](#)
- [20] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 867–876, 2022. [2](#), [4](#), [8](#), [13](#)
- [21] John R Kender, Parijat Dube, Zhengyang Han, and Bishwaranjan Bhattacharjee. G2l: A high-dimensional geometric approach for automatic generation of highly accurate pseudo-labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1093–1102, 2023. [2](#)
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. [2](#), [3](#), [4](#), [13](#), [14](#), [16](#)
- [23] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. [6](#)

- [24] S Chandeesh Kumar, M Hemalatha, S Badri Narayan, and P Nandhini. Region driven remote sensing image captioning. *Procedia Computer Science*, 165:32–40, 2019. [2](#)
- [25] Han-Hung Lee and Angel X Chang. Understanding pure clip guidance for voxel grid nerf models. *arXiv preprint arXiv:2209.15172*, 2022. [2](#), [3](#), [4](#), [8](#), [13](#), [17](#)
- [26] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*, 2023. [2](#), [3](#)
- [27] Dongxu Li, Junnan Li, and Steven CH Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *arXiv preprint arXiv:2305.14720*, 2023. [2](#), [3](#), [4](#), [5](#), [7](#), [11](#), [16](#)
- [28] Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, pages 19730–19742. PMLR, 2023. [2](#), [3](#), [5](#), [6](#), [7](#), [11](#), [12](#), [13](#), [14](#)
- [29] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10965–10975, 2022. [2](#), [3](#)
- [30] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23390–23400, 2023. [2](#)
- [31] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023. [2](#), [3](#), [5](#), [10](#), [11](#)
- [32] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [10](#), [11](#)
- [33] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. [2](#), [3](#), [6](#), [7](#), [11](#), [13](#), [15](#)
- [34] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. [2](#), [3](#), [5](#), [11](#), [12](#)
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [9](#)
- [36] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016. [6](#)
- [37] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [2](#), [8](#)
- [38] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. Clip-mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia 2022 conference papers*, pages 1–8, 2022. [2](#), [4](#)
- [39] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. [2](#), [3](#), [4](#), [5](#), [7](#), [8](#), [13](#), [17](#)
- [40] OpenAI. Gpt-4 technical report, 2023. [3](#)
- [41] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023. [2](#), [3](#), [4](#), [5](#), [7](#), [9](#), [11](#)
- [42] Yanyuan Qiao, Chaorui Deng, and Qi Wu. Referring expression comprehension: A survey of methods and datasets. *IEEE Transactions on Multimedia*, 23:4426–4440, 2020. [2](#)
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [2](#), [3](#), [4](#), [5](#), [7](#), [10](#), [13](#), [18](#)
- [44] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [2](#), [3](#), [4](#), [7](#)
- [45] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M. Anwer, Erix Xing, Ming-Hsuan Yang, and Fahad S. Khan. Glamm: Pixel grounding large multimodal model, 2023. [4](#), [6](#), [7](#)
- [46] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [7](#), [11](#)
- [47] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. [4](#)
- [48] Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. Clip-forge: Towards zero-shot text-to-shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18603–18613, 2022. [4](#)
- [49] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [3](#)
- [50] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo

- Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, 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. 3
- [51] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instant-booth: Personalized text-to-image generation without test-time finetuning. *arXiv preprint arXiv:2304.03411*, 2023. 2
- [52] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. What does clip know about a red circle? visual prompt engineering for vlms. *arXiv preprint arXiv:2304.06712*, 2023. 2, 3, 5, 6, 12
- [53] Wei Su, Peihan Miao, Huanzhang Dou, Yongjian Fu, and Xi Li. Referring expression comprehension using language adaptive inference. *arXiv preprint arXiv:2306.04451*, 2023. 2
- [54] Sanjay Subramanian, William Merrill, Trevor Darrell, Matt Gardner, Sameer Singh, and Anna Rohrbach. Reclip: A strong zero-shot baseline for referring expression comprehension. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5198–5215, 2022. 2, 3, 6, 10, 12, 13
- [55] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. EVA-CLIP: improved training techniques for CLIP at scale. *CoRR*, abs/2303.15389, 2023. 2
- [56] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative pretraining in multi-modality. *arXiv preprint arXiv:2307.05222*, 2023. 2
- [57] Vicuna. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. <https://vicuna.lmsys.org/>, 2023. 6
- [58] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022. 4, 7
- [59] Weiyun Wang, Min Shi, Qingyun Li, Wenhai Wang, Zhenhang Huang, Linjie Xing, Zhe Chen, Hao Li, Xizhou Zhu, Zhiguo Cao, Yushi Chen, Tong Lu, Jifeng Dai, and Yu Qiao. The all-seeing project: Towards panoptic visual recognition and understanding of the open world, 2023. 3
- [60] Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. ELITE: Encoding visual concepts into textual embeddings for customized text-to-image generation. *arXiv preprint arXiv:2302.13848*, 2023. 2, 4, 7, 11
- [61] Jialian Wu, Jianfeng Wang, Zhengyuan Yang, Zhe Gan, Zicheng Liu, Junsong Yuan, and Lijuan Wang. Grit: A generative region-to-text transformer for object understanding. *arXiv preprint arXiv:2212.00280*, 2022. 7
- [62] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. 2, 3, 10
- [63] Jiale Xu, Xintao Wang, Weihao Cheng, Yan-Pei Cao, Ying Shan, Xiaohu Qie, and Shenghua Gao. Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20908–20918, 2023. 4
- [64] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2945–2954, 2023. 2, 3
- [65] Xin Xu, Tianyi Xiong, Zheng Ding, and Zhuowen Tu. Masq-clip for open-vocabulary universal image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 887–898, 2023. 2, 3
- [66] Lingfeng Yang, Yueze Wang, Xiang Li, Xinlong Wang, and Jian Yang. Fine-grained visual prompting. *arXiv preprint arXiv:2306.04356*, 2023. 2, 3
- [67] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021. 6
- [68] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023. 2, 4, 7
- [69] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 69–85. Springer, 2016. 6
- [70] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. MATTNet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1307–1315, 2018. 2, 3, 6, 12
- [71] Pan Zhang, Xiaoyi Dong Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Shuangrui Ding, Songyang Zhang, Haodong Duan, Hang Yan, et al. InternLM-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*, 2023. 2, 3
- [72] Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601*, 2023. 4, 6, 7
- [73] Shiyu Zhao, Zhixing Zhang, Samuel Schuster, Long Zhao, BG Vijay Kumar, Anastasis Sathopoulos, Manmohan Chandraker, and Dimitris N Metaxas. Exploiting unlabeled data with vision and language models for object detection. In *European Conference on Computer Vision*, pages 159–175. Springer, 2022. 2
- [74] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. 2, 3

- [75] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, pages 696–712. Springer, 2022. [2](#), [3](#), [5](#), [10](#)
- [76] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *European Conference on Computer Vision*, pages 350–368. Springer, 2022. [2](#), [3](#), [6](#)
- [77] Yijie Zhou, Likun Cai, Xianhui Cheng, Zhongxue Gan, Xiangyang Xue, and Wenchao Ding. Openannotate3d: Open-vocabulary auto-labeling system for multi-modal 3d data. *arXiv preprint arXiv:2310.13398*, 2023. [2](#)