

Shatter and Gather: Learning Referring Image Segmentation with Text Supervision

Dongwon Kim^{1*} Namyup Kim^{1*} Cuiling Lan² Suha Kwak¹
¹POSTECH ²Microsoft Research Asia

¹{kdwon, namyup, suha.kwak}@postech.ac.kr ²culan@microsoft.com

Abstract

Referring image segmentation, the task of segmenting any arbitrary entities described in free-form texts, opens up a variety of vision applications. However, manual labeling of training data for this task is prohibitively costly, leading to lack of labeled data for training. We address this issue by a weakly supervised learning approach using text descriptions of training images as the only source of supervision. To this end, we first present a new model that discovers semantic entities in input image and then combines such entities relevant to text query to predict the mask of the referent. We also present a new loss function that allows the model to be trained without any further supervision. Our method was evaluated on four public benchmarks for referring image segmentation, where it clearly outperformed the existing method for the same task and recent open-vocabulary segmentation models on all the benchmarks.

1. Introduction

Referring image segmentation is the task of segmenting the referent corresponding to a natural language expression given as a query. Unlike the conventional semantic segmentation that aims at segmenting a pre-defined set of classes, referring image segmentation enables segmentation of any arbitrary entities described in free-form texts and thus opens up a wide variety of applications such as human-computer interaction [8, 59] and text-based image editing [7, 17].

Thanks to the development of deep neural networks, recent studies have demonstrated remarkable results on referring image segmentation in the supervised learning setting [14, 22, 25, 27, 70, 71, 72]. However, obtaining manual annotation of training data for referring image segmentation is prohibitively costly since the data requires two types of labels for each image, *i.e.*, natural language expressions describing the entities that appear in the image and segmentation labels corresponding to them. This typically demands tremendously more manpower than annotation for semantic

*Equal contribution.

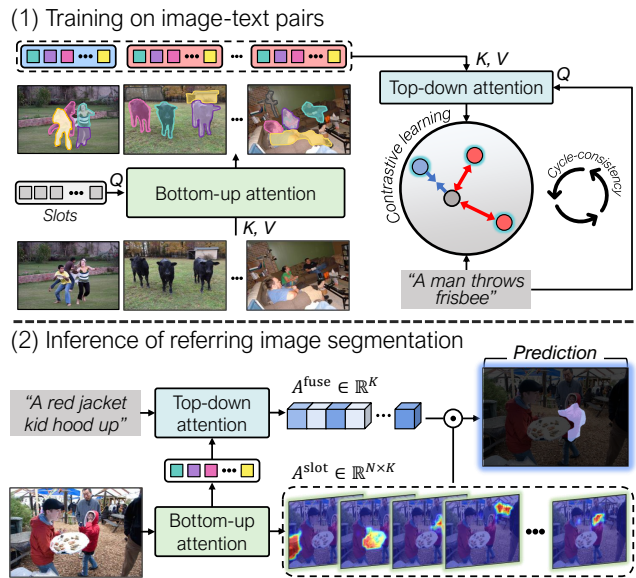


Figure 1: Our model consists of two different attention modules: bottom-up attention that identifies individual entities in an image, and top-down attention that combines the entities based on a referring expression. It is trained in a contrastive manner to ensure consistency between matching image-text pairs. In inference, a segmentation mask is predicted by combining the entities found by bottom-up attention with weights derived by top-down attention.

segmentation, leading to lack of labeled data for training.

A solution to this issue is weakly supervised learning, which trains a model on a dataset providing weaker forms of supervision than the conventional ones. In this paper, we particularly focus on learning referring image segmentation only with natural language expressions describing the associated images. Since these labels are readily available in many vision-language datasets [28, 48], the issue on lack of labeled training data can be alleviated by our approach. However, learning referring image segmentation in this setting introduces another challenge: Since natural language expressions often contain relations between visual entities such as instances and object parts (*e.g.*, “The cat next to the

table leg”), a model to be trained should be aware of the inter-object relations to exploit the expressions as supervision. Thereby we need a method to discover individual entities and infer their relations without any supervision but the referring expression.

To fulfill this requirement, we introduce a new *bottom-up* and *top-down* attention framework, which is illustrated in Fig. 1. Individual entities existing in an image are first discovered by *bottom-up* attention, which solely exploits visual information. For the bottom-up attention, an entity discovery module progressively refines a set of embedding vectors named *slots* to capture distinct visual entities, following slot attention [43]. We moreover propose a novel slot formulation named *entity slot*, which enables fine-grained entity discovery in real-world images. As a result, entity slots after the refinement capture individual visual entities without any segmentation supervision. These entities are considered primitive units for composing the segmentation mask of a referent during inference.

Afterward, based on the relational and structural information of referring expression, top-down attention selectively attends to the referred entities and combines them into the predicted mask. Top-down attention is implemented by a modality fusion module consists of cross-attention transformer [15]. This attention scheme enables the module to infer the relevance between discovered entities and the referring expression instead of relying on the cosine similarity or heuristic similarity functions used in previous work [18, 55, 67, 76].

For training, we propose the contrastive cycle-consistency (C^3) loss that enforces cycle-consistency between image-text pairs under the contrastive learning framework [48]. To establish such consistency, the model is trained to preserve the textual information after the fusion of entities and textual features. By doing so, we observe that latent relevance between the discovered entities and referring expression automatically emerges in the top-down attention without any explicit mask supervision.

One may wonder the difference between our task and open-vocabulary segmentation [18, 67, 76], *i.e.*, semantic segmentation of arbitrary categories without explicit mask supervision. The key difference is that a referent to segment in referring image segmentation is given by a complex free-form text, which introduces additional challenges.

The proposed method was evaluated and compared on four public benchmarks for referring image segmentation, where it substantially outperformed the previous weakly supervised learning method [55]. Moreover, we reproduced recent open-vocabulary segmentation models [67, 76] and evaluated them on referring image segmentation benchmarks; the results show that the proposed model surpassed them even without pre-training on large-scale image-text data [5, 48, 56].

In summary, our contribution is three-fold as follows:

- We present a new weakly supervised learning model for referring image segmentation. Our model employs bottom-up and top-down attention to discover individual entities and infer their relations, which enables to exploit natural language expressions as supervision effectively during training.
- We propose a new loss function called contrastive cycle-consistency loss, which allows latent relevance between discovered entities and referring expressions to emerge without requiring further supervision.
- Our model clearly outperformed the existing method for the same task and recent open-vocabulary segmentation models on all the four benchmarks.

2. Related Work

Referring image segmentation: Early approaches [22, 39, 41, 53] obtain multimodal feature maps by simple concatenation-convolution operation. With this operation as a basis, most existing studies have suggested attention mechanisms and multi-level feature aggregations [6, 16, 23, 25, 47, 71]. Various approaches have been employed for calculating fine interactions between the modalities, including multimodal feature fusion [14, 29, 33] using transformers [15], a large-scale vision-language pretrained framework [60], and holistically multimodal fusion [70]. The existing referring image segmentation methods have been developed in a fully supervised learning setting. Recently, a weakly supervised learning method [55] using only image-text pairs for training is introduced to alleviate the high annotation cost. However, this study relies on a heuristic similarity function to infer relations between patches and texts, which suffers from comprehending the relational and structural information from image and text modalities. Therefore, we propose a new weakly supervised model consisting of bottom-up and top-down attention to discover visual entities and infer their relations effectively.

Phrase grounding: Phrase grounding [12, 24, 29, 50, 69, 72] is the task of learning the association between image regions and words in caption describing the image. Weakly supervised approaches have also been proposed for this task [10, 20, 42, 58, 66], where the correspondence between words and regions are not given. It is worth noting that the contrastive learning framework introduced in [20] is closely related to the proposed C^3 loss by enforcing mutual information between regions and words. However, these studies rely on an off-the-shelf object detector [19], where the final prediction is restricted by the predefined set of classes with which the detector is trained. On the contrary, our method discovers visual entities in an image without an auxiliary model trained with external labeled data, which is more suitable for referring image segmentation.




Figure 2: Illustration of the overall architecture of our model along with its behavior during training and inference. (a) Feature extractors: The visual and textual features are extracted by transformer-based encoders of the two modalities respectively. (b) Entity discovery module: A set of visual entities are discovered from the visual features through the bottom-up attention. (c) Modality fusion module: The visual entities and a referring expression are fused considering their relevance, which is estimated by the top-down attention. Our model is trained with the proposed contrastive cycle-consistency (C^3) loss between the cross-modal embeddings and textual features. In inference, the segmentation mask is predicted by jointly considering the visual entities A^{slot} from the entity discovery module and the relevance scores A^{fuse} from the modality fusion module.

Weakly supervised learning for segmentation: Weakly supervised learning has become a favored solution for addressing the lack of annotated data in training, especially for semantic segmentation [1, 2, 21, 26, 34, 37, 68, 74, 79] and instance segmentation [30, 31, 49, 78]. These methods refine coarse localization cues obtained by CAMs [52, 62, 75], and use the results as pseudo labels for self-training. However, they cannot be directly applied to the referring image segmentation due to the following challenges. First, there are no predefined classes in this task, making it difficult to learn shared semantics within CAMs by multiple instance learning. Second, there is no mechanism to discover visual entities and combine them with textual information.

Slot attention: Slot attention [43] is proposed for the learning object-centric representation. It can discover the objects within an image in an unsupervised manner and be used in the various downstream task such as segmentation [67, 77] and retrieval [32, 63]. In specific, slot attention introduces an iterative attention mechanism that operates on the slots, which are the set of embedding vectors corresponding to the objects that appear in the input data. In previous work, slots are initialized as a set of learnable embeddings [32, 63, 67, 77] or a set of vectors sampled from a learnable Gaussian distribution [43, 54]. However, we empirically observed that both types of initialization are inapplicable to visual entity discovery for referring image segmentation, which we will elaborate on in Sec. 3.2.

Open-vocabulary segmentation: The open-vocabulary segmentation is proposed to address the limitation of semantic segmentation that it is constrained to the pre-

defined set of classes. Unlike semantic segmentation, open-vocabulary segmentation models are evaluated with classes unseen during training. The early approaches propose a generative model [4] and a projection model [65] to generate or project pixel-level semantic features from visual features for the unseen classes through word embedding space. Recent studies [18, 38, 67, 76] largely improve the quality of semantic features for the unseen classes, and further introduce the methods that do not require segmentation labels for seen classes using large-scale image-text pretrained models (e.g. CLIP [48], ALIGN [28]). However, since they are trained to predict masks for semantic categories, not for free-form texts describing arbitrary entities in an image, their applicability is still limited. We thus propose a new method for referring image segmentation that only requires image-text pairs during training.

3. Proposed Method

As shown in Fig. 2, our model consists of three main components: feature extractors, entity discovery module, and modality fusion module. The visual and textual features are extracted from the input using the backbone networks of the two modalities respectively. (Sec. 3.1). The visual features are then given as input to the entity discovery module, which aggregates the features into multiple entity embeddings using the bottom-up attention (Sec. 3.2). The entity embeddings correspond to distinct entities appearing in the image, which are considered as the smallest units of building the segmentation mask. The modality fusion module fuses the entity embeddings and the textual feature into

cross-modal embeddings according to their relevance using the top-down attention (Sec. 3.3). The produced cross-modal embedding is used only for training, with a learning objective enforcing cycle-consistency and reconstruction of visual features (Sec. 3.4). The segmentation mask is predicted by combining the entities discovered by the entity discovery module with the relevance scores estimated by the modality fusion module (Sec. 3.5).

3.1. Visual and Textual Feature Extraction

Here, we provide the details of the transformer-based backbone network for each modality.

Visual feature extraction: We employ the vision transformer (ViT) [15] as the visual backbone network. An input image is first split into N patches and then linearly projected. The patches are processed by the multiple self-attention blocks of ViT, where we refer to the output tokens of the last block as the visual features $\mathbf{x}^V \in \mathbb{R}^{N \times D}$.

Textual feature extraction: We consider the transformer-based text encoder, BERT [13]. A sentence is first tokenized into the token sequence, where the [CLS] and [SEP] tokens are added to the beginning and the end of the sequence. The activation of the last layer of the transformer for the [CLS] token is employed as the textual feature $\mathbf{x}^T \in \mathbb{R}^D$.

3.2. Bottom-up Attention: Entity Discovery Module

The entity discovery module applies a bottom-up attention mechanism to aggregate patch-based visual features into multiple entity embeddings that represent instances or object parts in the image. The discovered entities are then treated as primitive units of the segmentation mask, where the mask is predicted by selectively combining them. This approach is useful since the textual information is often built upon the relations between such visual entities.

Specifically, the entity discovery module consists of two blocks: aggregation block and interaction block. The *entity slots*, which are embedding vectors that corresponds to various visual entities, are updated in an iterative process alternating between these two blocks. These slots eventually become entity embeddings. Let $\Phi : \mathbb{R}^{N \times D} \mapsto \mathbb{R}^{K \times D}$ be the entity discovery module that takes visual features as input and outputs the set of K entity embeddings. The entity embeddings are obtained by iteratively updating initial entity slots $\mathbf{S}^0 \in \mathbb{R}^{K \times D}$ T -times:

$$\begin{aligned} \Phi(\mathbf{x}^V) &:= \mathbf{S}^T, \text{ where} \\ \mathbf{S}^{t+1} &= \text{Inter}(\text{Agg}(\mathbf{x}^V; \mathbf{S}^t)), \end{aligned} \quad (1)$$

and Agg and Inter denote the aggregation and interaction blocks, respectively.

In the literature, the initial slots \mathbf{S}^0 are often defined as either vectors sampled from a learnable Gaussian distribution [43, 54] or learnable embeddings [32, 63, 67, 77],

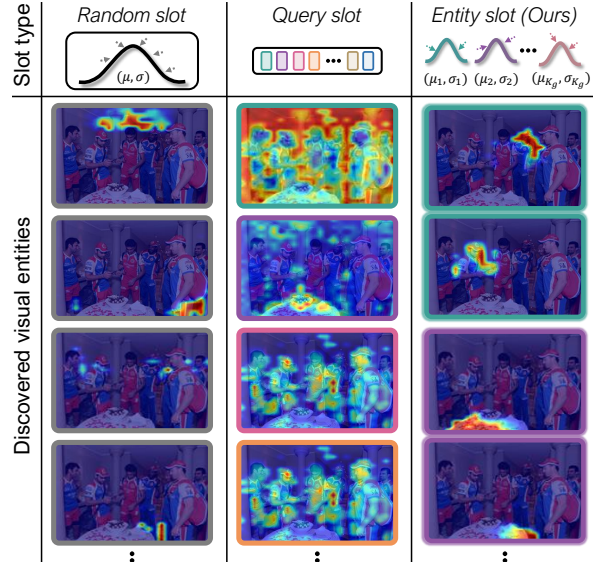


Figure 3: Comparison between the proposed entity slot and previous slot initialization methods. The boundary color indicates the Gaussian distribution that each slot is sampled from. The entity slot successfully discovers fine-grained entities such as head and food, while the others often fail to capture individual entities of specific semantics.

where such types of slots are referred as *random slot* and *query slot* in the remainder of this paper, respectively. However, we have empirically found that both of these initialization strategies are unsuitable for our task as shown in Fig. 3. The random slot suffers from the lack of semantic specificity since every slot is randomly sampled from a single distribution. This can lead to suboptimal performance since entities in real-world images are diverse in terms of semantics. The query slot, on the other hand, suffers from the limitation of being bound to specific semantic categories, such as human, animal, and vehicle. This limitation leads to entity discovery that is not sufficiently fine-grained, e.g., given an image of multiple people, a single slot that binds to “human” captures all of them and thus failed to identify the individual specified by the text query. This is problematic particularly in our model since entities act as the smallest units for producing segmentation masks. It is highly desired to have a mechanism that could guide the differentiation of slots to be semantically specific.

To overcome these limitations, we introduce a new initialization strategy for the *entity slot*, which enables the discovery of fine-grained visual entities while learning diverse semantics. Formally, the initial entity slots \mathbf{S}^0 are obtained by sampling K_s slots from each of K_g different Gaussian distributions, resulting in a total of $K = K_s \cdot K_g$ slots:

$$\begin{aligned} \mathbf{S}^0 &:= [\mathbf{S}_1^0; \dots; \mathbf{S}_{K_g}^0] \in \mathbb{R}^{K \times D}, \text{ where} \\ \mathbf{S}_i^0 &\sim \mathcal{N}(\mu_i, \text{diag}(\sigma_i)) \in \mathbb{R}^{K_s \times D}. \end{aligned} \quad (2)$$

The idea is to initialize slots by sampling them from different Gaussian distributions, each specifying a distinct semantic property to be captured. Individual entities sharing the same semantics are differentiated by slots sampled from the same distribution, which enables fine-grained entity discovery.

From now, we present the aggregation block and interaction block at t -th iteration. As described in Eq. (1), the slots of previous iteration \mathbf{S}^{t-1} are fed to the aggregation block with the visual features \mathbf{x}^V . The role of the aggregation block is to let each slot aggregate visual features of distinct visual entities, and it is achieved by employing the competitive attention mechanism introduced in the slot attention [43]. Let $q(\cdot)$, $k(\cdot)$, and $v(\cdot)$ are linear projection with layer normalization [3] that map inputs to dimension D_h . Then, the attention map $A^{\text{slot}} \in \mathbb{R}^{N \times K}$ between \mathbf{x}^V and \mathbf{S}^{t-1} is computed as follows:

$$A_{i,j}^{\text{slot}} = \frac{e^{M_{i,j}}}{\sum_{j=1}^K e^{M_{i,j}}}, \text{ where } M = \frac{k(\mathbf{x}^V)q(\mathbf{S}^{t-1})^\top}{\sqrt{D_h}}. \quad (3)$$

Unlike the cross-attention transformer [57] that normalized attention over the *keys*, we normalize attention map A^{slot} over the *slots*. This normalization scheme encourages slots to compete for aggregating the visual features, which results in the set of slots binding to different visual entities in the input data. It is worth noting that A^{slot} provides the localizations of entities discovered, which is used as primitive units of the segmentation mask during inference. Using A^{slot} , the slots are updated in a similar manner to the transformer:

$$\begin{aligned} \widehat{\mathbf{S}}^t &= \text{Agg}(\mathbf{x}^V; \mathbf{S}^t) \\ &= \text{normalize}(A^{\text{slot}})^\top v(\mathbf{x}^V)W_o + \mathbf{S}^{t-1}, \end{aligned} \quad (4)$$

where $\text{normalize}(\cdot)$ denotes ℓ_1 -normalization of the column vectors and $W_o \in \mathbb{R}^{D_h \times D}$ is linear projection. Finally, slots are further refined by $\widehat{\mathbf{S}}^t \leftarrow \text{MLP}(\widehat{\mathbf{S}}^t) + \widehat{\mathbf{S}}^t$, where MLP is multi-layer perceptron (MLP) with layer normalization.

The following interaction block consists of multiple self-attention transformers [57]. The weights of these transformers are shared, and each transformer is applied to slots sampled from the same distribution. We denote the transformer layer as $\text{Transformer}(Q; K; V)$, where $(Q; K; V)$ are for query, key, and value. The interaction block is denoted as follows:

$$\text{Inter}(\widehat{\mathbf{S}}^t) = [\text{SA}(\widehat{\mathbf{S}}_1^t); \dots; \text{SA}(\widehat{\mathbf{S}}_{K_g}^t)], \quad (5)$$

where $\text{SA}(X)$ is for the $\text{Transformer}(X; X; X)$. This self-attention scheme models information exchange between visual entities sharing the same semantic property, which helps slots within a single semantic category to capture different entities without redundancy.

3.3. Top-down Attention: Modality Fusion Module

The modality fusion module is proposed for top-down attention mechanism, which combines discovered visual entities based on the textual feature. Let $\Psi: \mathbb{R}^{K \times D} \times \mathbb{R}^D \mapsto \mathbb{R}^D$ be the modality fusion module, which takes entity embeddings $\Phi(\mathbf{x}^V)$ and textual feature \mathbf{x}^T as inputs. Then, the output cross-modal embedding $\mathbf{z} \in \mathbb{R}^D$ is computed using cross-attention transformer:

$$\mathbf{z} = \Psi(\Phi(\mathbf{x}^V); \mathbf{x}^T) = \frac{\text{CA}(\Phi(\mathbf{x}^V); \mathbf{x}^T)}{\|\text{CA}(\Phi(\mathbf{x}^V); \mathbf{x}^T)\|_2}, \quad (6)$$

where $\text{CA}(X; Y)$ is for the $\text{Transformer}(Y; X; X)$ and $\|\cdot\|_2$ denotes ℓ_2 norm. Moreover, we can obtain the attention map $A^{\text{fuse}} \in \mathbb{R}^K$ as a byproduct in cross-attention:

$$A_i^{\text{fuse}} = \frac{e^{M_i}}{\sum_{i=1}^K e^{M_i}}, \text{ where } M = \frac{k(\Phi(\mathbf{x}^V))q(\mathbf{x}^T)}{\sqrt{D_h}}. \quad (7)$$

Using cross-attention, the module effectively reasons the relevance between entities and the textual feature. Inferred relevance is described in the A^{fuse} , which is the probability of each entity being referred to in the textual feature.

3.4. Training

Our model is trained with two different loss functions: contrastive cycle-consistency (C^3) and reconstruction loss.

3.4.1 Contrastive cycle-consistency (C^3) Loss

The output of our model is the cross-modal embedding, which is obtained by fusing information from visual and textual modalities. The fusion should selectively combine discovered entities based on textual information, but a segmentation label for which entities are referred to is not given in our weakly-supervised setting. To resolve this issue, we add an additional constraint to the contrastive learning objective: establishing *cycle-consistency* between each ground-truth pair of image and text.

Let $\{(\mathbf{x}_i^V, \mathbf{x}_i^T)\}_{i=1}^B$ represent a batch of visual and textual features, and $\mathbf{z}_{ij} = \Psi(\Phi(\mathbf{x}_i^V); \mathbf{x}_j^T)$ denote the cross-modal embedding produced with \mathbf{x}_i^V and \mathbf{x}_j^T . For simplicity, we refer to the cross-modal embedding produced using the ground-truth image-text pair as matching embedding, and all other cross-modal embeddings as non-matching embeddings. Then C^3 loss is computed as follows:

$$\mathcal{L}_{C^3} = -\frac{1}{B} \sum_{j=1}^B \log \frac{e^{\langle \mathbf{z}_{jj}, \text{sg}(\mathbf{x}_j^T) \rangle}}{\sum_{i=1}^B e^{\langle \mathbf{z}_{ij}, \text{sg}(\mathbf{x}_j^T) \rangle}} \quad (8)$$

where $\langle \cdot, \cdot \rangle$ and $\text{sg}(\cdot)$ denote inner product and stop-gradient operation, respectively. Stop-gradient operation is used to prevent the collapsing solutions where every embedding becomes similar, as discussed in [9].

The proposed loss function encourages the matching embeddings to be similar to the input textual feature, while making the non-matching embeddings to be dissimilar. This approach ensures that the modality fusion module preserves the inherent information of the referring expression only when it is paired with the matching image (cycle-consistency), achieved by attending to the referred entities that may be absent in non-matching images. Therefore, an attention map A^{fuse} is trained to reflect the referring expression accurately, where relevance between the expression and entities automatically emerges.

3.4.2 Reconstruction Loss

If we can reconstruct the original data using only the set of slots, then the slots will correspond to distinctive individual entities in the input image. Therefore, we leverage the reconstruction objective $\mathcal{L}_{\text{recon}}$, which is given by

$$\mathcal{L}_{\text{recon}} = \frac{1}{B} \sum_{i=0}^B \|\text{sg}(\mathbf{x}_i^{\mathcal{V}}) - f_{\text{dec}}(\Phi(\mathbf{x}_i^{\mathcal{V}}))\|^2, \quad (9)$$

where $f_{\text{dec}}(\cdot)$ is a MLP based decoder. Unlike the original slot attention [43], our reconstruction is performed in the feature space, rather than the pixel space. This feature reconstruction scheme enables an efficient decoder design. Detailed descriptions are provided in the supplementary material.

3.5. Inference

The segmentation mask is computed with the two attention maps: $A^{\text{slot}} \in \mathbb{R}^{N \times K}$ and $A^{\text{fuse}} \in \mathbb{R}^K$. A^{slot} is patch-wise attention map that contains information about the discovered entities (Eq. (3)), while A^{fuse} is entity-wise attention map denoting relevance between the referring expression and the entities (Eq. (7)). To compute the segmentation mask, we first compute $A^{\text{slot}} \cdot A^{\text{fuse}}$, which results in a 1-dimensional vector of size N . We then reshape this vector into a 2-dimensional map. Next, we bilinearly interpolate the map to the original image size and rescale it using min-max normalization. Finally, we binarize the map using a threshold value τ to obtain the segmentation mask.

4. Experiments

4.1. Experimental Setting

Datasets: Our experiments utilize four datasets commonly used in referring image segmentation task: RefCOCO [73], RefCOCO+ [73], Gref [46], and PhraseCut [64]. RefCOCO, RefCOCO+, and Gref are obtained from COCO [40] dataset. RefCOCO contains 142,209 image-text pairs, while RefCOCO+ has 141,564 pairs and does not include expressions indicating orientation properties (*e.g.*, top, right). Gref includes 104,560 pairs for 54,822 objects. PhraseCut comprises 345,486 image-text pairs, which are derived from Visual Genome [36].

Implementation details: For visual and textual backbone, we employ ViT-S-16 pretrained on ImageNet-21K [11] and pretrained BERT [13], respectively. We make each modality feature have the same channel dimension D by applying a 1×1 convolution layer. We set D and D_h to 512 and 1024, respectively. In the aggregation block, the number of iteration T is set to 6, and the number of slots K is set to 36 except for PhraseCut, which is set to 60. For the entity slot, the number of slots sampled from each Gaussian distributions K_s is set to 3 in RefCOCO+ and Gref, while 2 in RefCOCO and PhraseCut. During training, we use AdamW [45] optimizer, with the initial learning rate of $1e-4$ scheduled with cosine annealing [44]. Batch size is set to 32 and model is trained for 50 epochs. We resize input images to 384×384 for both training and evaluation. Threshold value τ is set to 0.5 for all experiments.

Evaluation metric: Following [70], we employ the cumulative intersection-over-union (cIoU) and mean intersection-over-Union (mIoU) metric; cIoU divides the total intersections by the total unions across all test images, and mIoU computes IoU for each image and averages them. Also, we evaluate the accuracy at the IoU thresholds of $\{0.3, 0.5, 0.7\}$, denote as $\mathcal{A}@ \{0.3, 0.5, 0.7\}$ in Table 3.


4.2. Comparisons with the State of the Art

Weakly supervised methods: As presented in Table 1, our method outperforms TSEG [55], the previous work on weakly-supervised referring image segmentation, in all benchmarks. It shows that the proposed bottom-up and top-down attention framework is more effective than using patch-based visual features with a carefully designed similarity function of TSEG.

Open-vocabulary segmentation: In Table 1, we also compare our model with the recent open-sourced open-vocabulary segmentation models: GroupViT [67] and MaskCLIP [76], on the referring image segmentation benchmarks. Note that these methods are pre-trained with large-scale image-text datasets [5, 48, 56]. We also report their fine-tuned version for fair comparisons: models are additionally trained with the image-text pairs of each benchmark. Further details for the reproductions are provided in the supplementary material. Ours outperforms GroupViT and MaskCLIP in all settings and benchmarks, even without large-scale image-text pretraining. This result shows that open-vocabulary segmentation models [67, 76] have difficulty comprehending the complex free-from text involving the relation between entities, unlike ours. Interestingly, our model without fine-tuning, which pre-trained on the PhraseCut dataset and evaluated on other benchmarks, still achieved better results. This setting is identical to the open-vocabulary segmentation models, except it leverages a far smaller dataset. These results demonstrate the feasibility of our method in a zero-shot transfer setting.

Methods	Image-text pretraining dataset	Fine tuning	RefCOCO			RefCOCO+			Gref	PhraseCut
			val	testA	testB	val	testA	testB	val	val
GroupViT [67]	CC12M+YFCC (26M)	✗	7.99	6.16	10.51	8.49	6.79	10.59	10.68	5.82
		✓	10.82	11.11	11.29	11.14	10.78	11.84	12.77	9.41
MaskCLIP [76]	WIT (400M)	✗	11.52	11.85	12.06	11.87	12.01	12.57	12.74	12.13
		✓	19.45	18.69	21.37	19.97	18.93	21.48	21.11	23.80
TSEG [55]	-	✓	25.44	-	-	22.01	-	-	22.05	28.77
Ours	Visual Genome (0.31M)	✗	21.80	19.00	24.96	22.20	19.86	24.85	25.89	-
	-	✓	34.76	34.58	35.01	28.48	28.60	27.98	28.87	33.45

Table 1: Comparison with weakly-supervised method (TSEG) and open-vocabulary segmentation method (GroupViT and MaskCLIP). The results on four datasets are reported in mIoU (%). Fine-tuning ✓ means that the model is trained with the image-text pairs of the target benchmark; otherwise, it just relies on the pre-trained weight without further training.



Query: "Guy on left with beard and guitar"

Figure 4: Qualitative results of our framework with entity slot (Ours), random slot, and query slot on RefCOCO val set. In the case of entity slots, the boundary color indicates the Gaussian distribution that each slot is sampled from. For each slot type, we present 10 discovered entities from A^{slot} and the final predictions.

4.3. Ablation Study

We perform ablation studies on the val set of RefCOCO to analyze the contributions of the entity discovery module and the modality fusion module.

Entity discovery module: In Table 2(a), we ablate the entity discovery module and use visual feature directly, which leads to the drastical performance drop. This result demonstrates that contribution of our bottom-up attention framework is significant. Moreover, we ablate the interaction block, replace slot-attention of aggregation block to transformer, and train the model without reconstruction loss. The results show that all of the design choices for the module is important for effective entity discovery.

Modality fusion module: In Table 2(b), the modality fusion module is ablated, and the model is trained with image-text contrastive learning [48], where the image feature is obtained by pooling the entity embeddings. The result demonstrates that inferring relevance between entities and referring expression with modality fusion module, trained with C^3 loss, is crucial for accurate segmentation.

4.4. In-depth Analysis

Analysis of slot initialization: To verify our design choice for entity slots, we investigate variants of slots; query, random, and entity slot. As summarized in Table 3, our entity slot outperforms other slot types across the different

	cIoU	mIoU
(a) Analysis on entity discovery module		
Ours	30.40	34.76
w/o entity discovery module	5.40	5.91
w/o interaction block	29.59	33.69
aggregation block w/ transformer	11.29	12.61
w/o reconstruction loss	18.61	21.84
(b) Analysis on modality fusion module		
Ours	30.40	34.76
w/o modality fusion module and C^3 loss	13.81	14.65

Table 2: Ablation studies of model components and loss function on RefCOCO val set.

numbers of slots. Specifically, as the number of slots is increased, the performance of the query slot declines, and that of the random slot converges. This result implies that simply increasing the number of slots can not handle the discovery of fine-grained visual entities in real-world images, and thus gains from the proposed entity slot type are non-trivial. In Fig. 4, we also present several visual entities discovered using each slot type, which show the superiority of the entity slot against others as explained in Sec. 3.2. Specifically, the entity slots from the same distribution share the same semantic property (e.g., head, body, or window) while capturing different fine-grained entities. The pro-

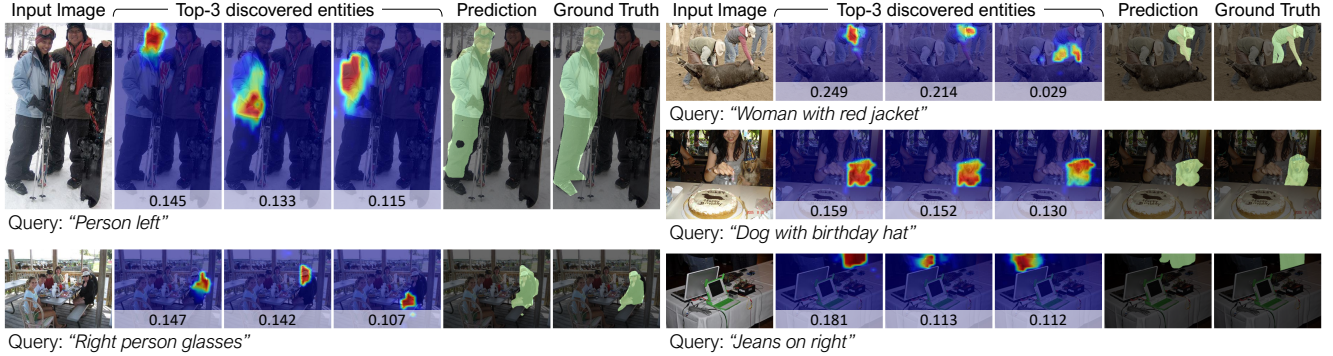


Figure 5: Qualitative results of our framework on RefCOCO *val* set. We present the discovered entities from A^{slot} and their relevance scores from A^{slot} . Top-3 entities in terms of relevance to query expression are presented.

Slot type	# slots	$\mathcal{A}@0.3$	$\mathcal{A}@0.5$	$\mathcal{A}@0.7$	cIoU	mIoU
Query	12	27.57	7.17	0.96	22.04	21.44
	24	36.73	10.61	1.93	23.59	24.78
	36	19.30	3.48	0.31	16.66	18.16
Random	12	36.30	11.27	2.69	24.20	26.65
	24	45.19	16.01	3.66	26.60	29.87
	36	45.62	16.88	3.76	26.55	30.15
Entity (Ours)	12	39.43	13.87	3.56	24.72	27.84
	24	52.33	22.97	6.23	29.22	33.54
	36	55.02	24.99	6.35	30.40	34.76

Table 3: Comparison between different slot types on RefCOCO *val* set, where # slots is the total number of slots.

posed entity slots effectively discover fine-grained visual entities, which results in an accurate prediction. In contrast, the prediction using random slots leads to a noisy mask, and using query slots leads to the mask of every head that appears in the image, respectively.

Effect of hyperparameters K_g and K_s : In Table 4(a), we investigate the impact of hyperparameters for the entity slots: K_g and K_s , which are the number of Gaussian distributions and slots sampled from each distribution, respectively. We keep the total number of entity slots to 36, only changing K_g and K_s . The optimal configurations are 18 and 2 for K_g and K_s , respectively. Interestingly, we observed that even entity slots with $K_s = 1$, where a single slot is sampled from each distribution, vastly surpass the query slots in Table 3; it shows that random sampling of slots prevents the bounding to specific semantic category.

Effectiveness of A^{fuse} : To verify the efficacy of A^{fuse} quantitatively, we evaluate our model with the following inference scheme variants: (1) AVG that replaces A^{fuse} with the uniform probability distribution, (2) MAX that picks a single attention map from A^{slot} with the maximum probability in A^{fuse} , and (3) MIN that picks the attention map with the minimum probability. As summarized in Table 4(b), MAX largely outperforms AVG even though it only uses a single attention map for the final prediction, while MIN results in a near-zero cIoU and mIoU. This tendency shows that A^{fuse}

K_g	K_s	cIoU	mIoU	Inference scheme	cIoU	mIoU
36	1	29.36	33.41	AVG	9.01	8.16
18	2	30.40	34.76	MAX	14.93	19.54
12	3	29.63	34.03	MIN	0.16	0.14
9	4	28.31	32.34	$A^{\text{fuse}} \cdot A^{\text{slot}}$ (Ours)	30.40	34.76
6	6	27.23	31.30			

(a)

(b)

Table 4: Impact of hyperparameters K_g and K_s (a) and inference schemes (b) on RefCOCO *val* set.

effectively assigns probability based on the relevance between reference and entities, even without direct supervision.

Qualitative Analysis: In Fig. 5, predicted segmentation mask and discovered entities with top-3 relevance are presented. Results show that the entity discovery module effectively discovers primitive units of segmentation masks, and the modality fusion module accurately infers their relevance to query referring expression.

5. Conclusion

In this paper, we address the problem of learning referring image segmentation with only image-text pairs. We present a new model based on a bottom-up and top-down attention framework, which first discovers entities in the input image and then combines such entities based on their relevance to the text query. Moreover, we present a new loss function enforcing cycle-consistency between image-text pairs, which enables the effective training of the model without any further supervision. As a result, the proposed method clearly outperforms an existing weakly-supervised method and recent open-vocabulary segmentation models. For future work, we will explore expanding our model to different modalities such as video and audio.

Acknowledgement. This work was supported by the NRF grant and the IITP grants funded by Ministry of Science and ICT, Korea (NRF-2021R1A2C3012728-20%, IITP-2020-0-00842-20%, IITP-2022-0-00290-20%, IITP-2022-0-00926-20%, IITP-2021-0-02068-10%, IITP-2019-0-01906-10%).

References

- [1] Jiwoon Ahn, Sunghyun Cho, and Suha Kwak. Weakly supervised learning of instance segmentation with inter-pixel relations. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [2] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [3] Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. In *ArXiv*, volume abs/1607.06450, 2016. 5
- [4] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [5] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6
- [6] Ding-Jie Chen, Songhao Jia, Yi-Chen Lo, Hwann-Tzong Chen, and Tyng-Luh Liu. See-through-text grouping for referring image segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [7] Jianbo Chen, Yelong Shen, Jianfeng Gao, Jingjing Liu, and Xiaodong Liu. Language-based image editing with recurrent attentive models. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [8] Shizhe Chen, Pierre-Louis Guhur, Cordelia Schmid, and Ivan Laptev. History aware multimodal transformer for vision-and-language navigation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2021. 1
- [9] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 5
- [10] Samyak Datta, Karan Sikka, Anirban Roy, Karuna Ahuja, Devi Parikh, and Ajay Divakaran. Align2ground: Weakly supervised phrase grounding guided by image-caption alignment. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: a large-scale hierarchical image database. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 6
- [12] Jiajun Deng, Zhengyuan Yang, Tianlang Chen, Wengang Zhou, and Houqiang Li. Transvg: End-to-end visual grounding with transformers. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2019. 4, 6
- [14] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 1, 2
- [15] 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. In *Proc. International Conference on Learning Representations (ICLR)*, 2021. 2, 4
- [16] Guang Feng, Zhiwei Hu, Lihe Zhang, and Huchuan Lu. Encoder fusion network with co-attention embedding for referring image segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [17] Tsu-Jui Fu, Xin Eric Wang, Scott T. Grafton, Miguel P. Eckstein, and William Yang Wang. M3l: Language-based video editing via multi-modal multi-level transformers. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1
- [18] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 2, 3
- [19] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 2
- [20] Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding. In *European Conference on Computer Vision*. Springer, 2020. 2
- [21] S. Hong, D. Yeo, S. Kwak, H. Lee, and B. Han. Weakly supervised semantic segmentation using web-crawled videos. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2224–2232, 2017. 3
- [22] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 1, 2
- [23] Zhiwei Hu, Guang Feng, Jiayu Sun, Lihe Zhang, and Huchuan Lu. Bi-directional relationship inferring network for referring image segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [24] Binbin Huang, Dongze Lian, Weixin Luo, and Shenghua Gao. Look before you leap: Learning landmark features for one-stage visual grounding. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [25] Shaofei Huang, Tianrui Hui, Si Liu, Guanbin Li, Yunchao Wei, Jizhong Han, Luoqi Liu, and Bo Li. Referring image segmentation via cross-modal progressive comprehension. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2
- [26] Zilong Huang, Xinggang Wang, Jiasi Wang, Wenyu Liu, and Jingdong Wang. Weakly-supervised semantic segmentation network with deep seeded region growing. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [27] Tianrui Hui, Si Liu, Shaofei Huang, Guanbin Li, Sansi Yu, Faxi Zhang, and Jizhong Han. Linguistic structure guided context modeling for referring image segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2020. 1
- [28] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *Proc. International*

- Conference on Machine Learning (ICML)*, 2021. 1, 3
- [29] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2021. 2
- [30] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele. Simple does it: Weakly supervised instance and semantic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [31] Beomyoung Kim, Youngjoon Yoo, Chae Eun Rhee, and Junmo Kim. Beyond semantic to instance segmentation: Weakly-supervised instance segmentation via semantic knowledge transfer and self-refinement. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [32] Dongwon Kim, Namyup Kim, and Suha Kwak. Improving cross-modal retrieval with set of diverse embeddings. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3, 4
- [33] Namyup Kim, Dongwon Kim, Cuiling Lan, Wenjun Zeng, and Suha Kwak. Restr: Convolution-free referring image segmentation using transformers. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [34] Alexander Kolesnikov and Christoph H. Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 3
- [35] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2011. 12, 13
- [36] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)*, 2017. 6
- [37] Suha Kwak, Seunghoon Hong, and Bohyung Han. Weakly supervised semantic segmentation using superpixel pooling network. In *Proc. AAAI Conference on Artificial Intelligence (AAAI)*, 2017. 3
- [38] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *Proc. International Conference on Learning Representations (ICLR)*, 2022. 3
- [39] Ruiyu Li, Kaican Li, Yi-Chun Kuo, Michelle Shu, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Referring image segmentation via recurrent refinement networks. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [40] 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 *Proc. European Conference on Computer Vision (ECCV)*, 2014. 6
- [41] Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. Recurrent multimodal interaction for referring image segmentation. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [42] Yongfei Liu, Bo Wan, Lin Ma, and Xuming He. Relation-aware instance refinement for weakly supervised visual grounding. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [43] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3, 4, 5, 6
- [44] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *Proc. International Conference on Learning Representations (ICLR)*, 2017. 6
- [45] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *Proc. International Conference on Learning Representations (ICLR)*, 2019. 6
- [46] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [47] Edgar Margffoy-Tuay, Juan C Pérez, Emilio Botero, and Pablo Arbeláez. Dynamic multimodal instance segmentation guided by natural language queries. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 2
- [48] 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 *Proc. International Conference on Machine Learning (ICML)*, 2021. 1, 2, 3, 6, 7
- [49] Tal Remez, Jonathan Huang, and Matthew Brown. Learning to segment via cut-and-paste. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 3
- [50] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 2
- [51] Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, Thomas Brox, and Francesco Locatello. Bridging the gap to real-world object-centric learning. In *Proc. International Conference on Learning Representations (ICLR)*, 2023. 13
- [52] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [53] Hengcan Shi, Hongliang Li, Fanman Meng, and Qingbo Wu. Key-word-aware network for referring expression image segmentation. In *Proc. European Conference on Computer Vision (ECCV)*, 2018. 2
- [54] Gautam Singh, Fei Deng, and Sungjin Ahn. Illiterate dalle learns to compose. In *Proc. International Conference on Learning Representations (ICLR)*, 2021. 3, 4
- [55] Robin Strudel, Ivan Laptev, and Cordelia Schmid. Weakly-supervised segmentation of referring expressions. *arXiv*

- preprint *arXiv:2205.04725*, 2022. 2, 6, 7, 12, 13
- [56] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 2016. 2, 6
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2017. 5
- [58] Liwei Wang, Jing Huang, Yin Li, Kun Xu, Zhengyuan Yang, and Dong Yu. Improving weakly supervised visual grounding by contrastive knowledge distillation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [59] Xin Wang, Qiuyuan Huang, Asli Celikyilmaz, Jianfeng Gao, Dinghan Shen, Yuan-Fang Wang, William Yang Wang, and Lei Zhang. Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [60] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. Cris: Clip-driven referring image segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [61] Nick Watters, Loic Matthey, Chris P. Burgess, and Alexander Lerchner. Spatial broadcast decoder: A simple architecture for disentangled representations in VAEs. In *ICLR Workshop on Learning from Limited Labeled Data*, 2019. 13
- [62] Yunchao Wei, Jiashi Feng, Xiaodan Liang, Ming-Ming Cheng, Yao Zhao, and Shuicheng Yan. Object region mining with adversarial erasing: A simple classification to semantic segmentation approach. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [63] Philippe Weinzaepfel, Thomas Lucas, Diane Larlus, and Yannis Kalantidis. Learning super-features for image retrieval. In *Proc. International Conference on Learning Representations (ICLR)*, 2022. 3, 4
- [64] Chenyun Wu, Zhe Lin, Scott Cohen, Trung Bui, and Subhransu Maji. Phrasecut: Language-based image segmentation in the wild. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6
- [65] Yongqin Xian, Subhabrata Choudhury, Yang He, Bernt Schiele, and Zeynep Akata. Semantic projection network for zero-and few-label semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8256–8265, 2019. 3
- [66] Fanyi Xiao, Leonid Sigal, and Yong Jae Lee. Weakly-supervised visual grounding of phrases with linguistic structures. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [67] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3, 4, 6, 7, 12
- [68] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Multi-class token transformer for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4310–4319, 2022. 3
- [69] Li Yang, Yan Xu, Chunfeng Yuan, Wei Liu, Bing Li, and Weiming Hu. Improving visual grounding with visual-linguistic verification and iterative reasoning. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [70] Zhao Yang, Jiaqi Wang, Yansong Tang, Kai Chen, Hengshuang Zhao, and Philip HS Torr. Lavt: Language-aware vision transformer for referring image segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 6
- [71] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2
- [72] 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 *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [73] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *Proc. European Conference on Computer Vision (ECCV)*, 2016. 6
- [74] Dong Zhang, Hanwang Zhang, Jinhui Tang, Xian-Sheng Hua, and Qianru Sun. Causal intervention for weakly-supervised semantic segmentation. In *Proc. Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [75] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [76] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *Proc. European Conference on Computer Vision (ECCV)*, 2022. 2, 3, 6, 7, 12
- [77] Yi Zhou, Hui Zhang, Hana Lee, Shuyang Sun, Pingjun Li, Yangguang Zhu, ByungIn Yoo, Xiaojuan Qi, and Jae-Joon Han. Slot-vps: Object-centric representation learning for video panoptic segmentation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3, 4
- [78] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [79] Yuliang Zou, Zizhao Zhang, Han Zhang, Chun-Liang Li, Xiao Bian, Jia-Bin Huang, and Tomas Pfister. Pseudoseg: Designing pseudo labels for semantic segmentation. In *Proc. International Conference on Learning Representations (ICLR)*, 2021. 3

Appendix

Symbol	Description
Feature extraction	
\mathbf{x}^V	The visual feature
\mathbf{x}^T	The textual feature
Entity discovery module	
Φ	The entity discovery module
Agg	The aggregation block
normalize	ℓ_1 normalization of column vectors
MLP	A multi-layer perceptron with layer normalization
\mathbf{S}^0	Initial slots
\mathbf{S}^t	Slots at t -th iteration
$\hat{\mathbf{S}}^t$	Output slots of the aggregation block at t -th iteration
\mathbf{S}_i^t	Slots sampled from the i -th distribution at t -th iteration
Inter	The interaction block
SA	A self-attention transformer
Modality fusion module	
Ψ	The modality fusion module
CA	A cross-attention transformer
Training	
$\langle \cdot, \cdot \rangle$	A inner product between two vectors
sg	A stop-gradient operation
f_{dec}	The reconstruction decoder
Inference	
A^{slot}	The patch-wise attention map from the entity discover module
A^{fuse}	The entity-wise attention map from the modality fusion module
Hyperparameters	
N	The number of image patches
K	The total number of slots
K_g	The number of different Gaussian distributions for entity slot
K_s	The number of slots sampled from each distribution for entity slot
D, D_h	The output and hidden dimension
T	The number of iterations in entity discovery module
τ	The threshold for predicting mask

Table A1: Summary of notations used in the main paper together with descriptions.

This supplementary material presents experimental results omitted from the main paper due to the space limit. We first summarize the notations in our paper in Table A1. Sec. A analyzes performance according to the number of iterations of the entity discovery module T and thresholding value τ . In Sec. B, we describe experimental details of reproducing open-vocabulary segmentation methods [67, 76]. Sec. C provides the details of the decoder used in reconstruction loss. We then present quantitative results with the DenseCRF [35] in Sec. D. Finally, Sec. E offers more qualitative results of our model on the RefCOCO *val* set. Since our model involves randomness when sampling the entity slots, all of the results in the main paper and supplementary material are obtained by averaging the results from 3 experiments, where the standard deviation across them is always less than 0.06.

A. Impact of hyperparameters T and τ

The number of iterations T : In Table A2, We investigate the impact of the number of iterations in the entity discovery module T . The model consistently achieves high IoU and accuracy when T is greater than 2. Notably, we ob-

T	$\mathcal{A}@0.3$	$\mathcal{A}@0.5$	$\mathcal{A}@0.7$	cIoU	mIoU
8	53.08	23.30	5.98	29.38	33.85
6	55.02	24.99	6.35	30.40	34.76
4	52.03	22.55	5.85	28.95	33.24
2	48.40	19.92	5.09	27.67	31.77
1	40.36	12.93	2.73	25.11	28.00

Table A2: Performance analysis according to the number of iteration T of the entity discovery module on RefCOCO *val* set. The setting used in the main paper is indicated with a grey-colored row.

τ	0.3	0.4	0.5	0.6	0.7
mIoU	32.58	34.39	34.76	33.57	30.72

Table A3: Performance analysis according to thresholding hyperparameter τ on RefCOCO *val* set.

serve a significant performance degradation when T equals 1 (3.77%p mIoU drop compared to when T is 2). These results highlight the significance of iteratively applying the aggregation and interaction blocks, which enables effective visual entity discovery through the progressive refinement of slots. Furthermore, it is worth noting that even the least-performing model with a single iteration still outperforms the previous work [55, 67, 76].

The threshold value τ : In Table A3, we present the mIoU of predictions from our model with the varying threshold value τ . Our model consistently attains high IoUs when τ ranges between 0.4 and 0.6, which is why we set τ to 0.5 for all experiments in the main paper. These findings indicate that our model is insensitive to the setting of τ .

B. Reproducing MaskCLIP and GroupViT

To evaluate referring image segmentation performance of the open-vocabulary segmentation methods, we consider two open-sourced models that do not require mask supervision during training: MaskCLIP [76] and GroupViT [67].

For MaskCLIP, the target referring query and image are input to the model, with inference made similarly to our approach, *i.e.*, obtaining prediction mask by thresholding the similarity map between all image patches and target query, using a threshold value of 0.5.

For GroupViT, the target referring query and image are input to the model along with an additional dummy query (*e.g.*, “A photo of a nothing”). This dummy query functions similarly to the background class in semantic segmentation methods. During inference, the model assigns image segments to the query with higher similarity, and segments assigned to the target query are considered the final

Methods	RefCOCO			RefCOCO+			Gref	PC
	<i>val</i>	<i>testA</i>	<i>testB</i>	<i>val</i>	<i>testA</i>	<i>testB</i>	<i>val</i>	<i>val</i>
GroupViT	10.82	11.11	11.29	11.14	10.78	11.84	12.77	9.41
MaskCLIP	19.45	18.69	21.37	19.97	18.93	21.48	21.11	23.80
TSEG	25.44	-	-	22.01	-	-	22.05	28.77
TSEG†	25.95	-	-	22.62	-	-	23.41	30.12
Ours	34.76	34.58	35.01	28.48	28.60	27.98	28.87	33.45
Ours†	35.75	35.52	36.03	29.30	29.41	28.67	30.02	35.67

Table A4: Comparison with other methods, including the post-processing by DenseCRF. [35]. The results are reported in mIoU (%). PC and † denote the PhraseCut dataset and the models post-processed by DenseCRF, respectively.

prediction. We have also tried a similar inference protocol with our model for a fair comparison, but it failed to produce meaningful performance.

For the fine-tuning, we noticed that the loss does not decline with the original training configurations. Therefore, we set the batch size and learning rate for both models to 32 and $1e-5$, respectively, matching the hyperparameters used for our model.

C. Decoder f_{dec} for the reconstruction loss

For the reconstruction decoder, we follow the model architecture of the spatial broadcast decoder [61] and its application to slot attention [51]. Specifically, each entity slot is broadcasted to the N sequence, which is the same length as the visual feature. Next, the broadcasted slots are augmented with sinusoidal positional encoding. These augmented slots are then fed into a four-layer multi-layer perceptron (MLP) with ReLU activation functions. The output dimension of the MLP is $D + 1$, with the additional dimension being used to compute weighting values via softmax. Finally, we compute the reconstruction of the visual feature by taking the weighted sum of all slots in each sequence position, using the calculated weighting values.

D. Additional quantitative results

Following the previous work that utilizes post-processing to enhance segmentation quality, we present the performance of our model with DenseCRF [35] in Table A4. The results demonstrate that employing DenseCRF yields performance improvements across all benchmarks. Notably, the performance of our model without DenseCRF surpasses that of the previous method, TSEG [55], even when TSEG uses DenseCRF. In the main paper, we only report the performance of our model without DenseCRF, as its computation is time-consuming and the resulting benefits are relatively marginal.

E. Additional qualitative results

In Fig. A1 and Fig. A2, qualitative results of our model on the RefCOCO *val* set are presented. The results demonstrate that our model, trained solely with image-text pair supervision, can successfully discover visual entities and integrate them into segmentation masks corresponding to free-form text queries. For instance, our model predicts accurate masks for referring expressions about non-human objects (row 2-3 in Fig. A1, A2), occluded objects (row 3 in Fig. A1, A2), and partially appeared objects (rows 5 in Fig. A1, A2). Furthermore, the results of the top-3 discovered entities reveal that the entity discovery module effectively identifies visual entities, and the modality fusion module accurately infers their relevance to the text query.

In Fig. A3, we present additional qualitative results of our framework, featuring three types of slots: entity (Ours), random, and query slots. For the entity slots, those sampled from the same distribution are represented with identical boundary colors. The results indicate that using random slots leads to noisy entity discovery due to insufficient semantic specificity. In contrast, utilizing query slots generates not adequately fine-grained entities, as they are bound to particular semantic categories. Our proposed entity slot effectively addresses these shortcomings. It produces accurate visual entities by maintaining an awareness of semantic properties and facilitating fine-grained entity discovery. Specifically, we can observe that the entity slots sampled from the same Gaussian distribution share the same semantic properties (*e.g.*, head, chair, and color) while capturing individual entities without redundancy.



Figure A1: Qualitative results of our framework on RefCOCO *val* set. We present the discovered entities from A^{slot} and their relevance scores from A^{fuse} . Top-3 entities in terms of relevance to query expression are presented.



Figure A2: Qualitative results of our framework on RefCOCO *val* set. We present the discovered entities from A^{slot} and their relevance scores from A^{fuse} . Top-3 entities in terms of relevance to query expression are presented.



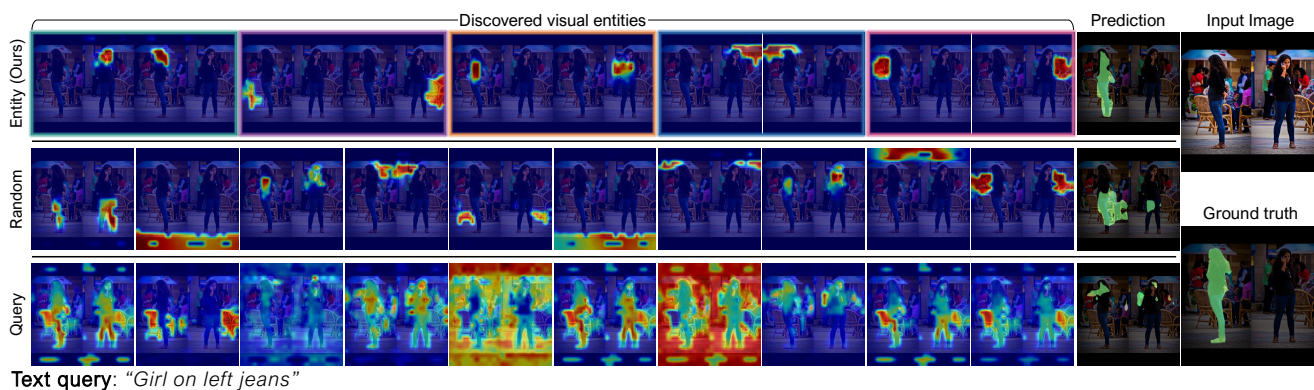



Figure A3: Qualitative results of our framework with entity slot (Ours), random slot, and query slot on RefCOCO *val* set. For each slot type, we present the 10 discovered entities from A^{slot} and final predictions. In the case of entity slots, the color of the boundaries indicates the Gaussian distribution that the slot sampled.