

# Semantics-Aware Dynamic Localization and Refinement for Referring Image Segmentation

Zhao Yang<sup>1</sup>, Jiaqi Wang<sup>2</sup>, Yansong Tang<sup>3</sup>, Kai Chen<sup>2</sup>, Hengshuang Zhao<sup>4</sup>, Philip H.S. Torr<sup>1</sup>

<sup>1</sup>University of Oxford,

<sup>2</sup>Shanghai AI Laboratory,

<sup>3</sup>Tsinghua-Berkeley Shenzhen Institute, Tsinghua University,

<sup>4</sup>The University of Hong Kong

## Abstract

Referring image segmentation segments an image from a language expression. With the aim of producing high-quality masks, existing methods often adopt iterative learning approaches that rely on RNNs or stacked attention layers to refine vision-language features. Despite their complexity, RNN-based methods are subject to specific encoder choices, while attention-based methods offer limited gains. In this work, we introduce a simple yet effective alternative for progressively learning discriminative multi-modal features. The core idea of our approach is to leverage a continuously updated query as the representation of the target object and at each iteration, strengthen multi-modal features strongly correlated to the query while weakening less related ones. As the query is initialized by language features and successively updated by object features, our algorithm gradually shifts from being localization-centric to segmentation-centric. This strategy enables the incremental recovery of missing object parts and/or removal of extraneous parts through iteration. Compared to its counterparts, our method is more versatile—it can be plugged into prior arts straightforwardly and consistently bring improvements. Experimental results on the challenging datasets of RefCOCO, RefCOCO+, and G-Ref demonstrate its advantage with respect to the state-of-the-art methods.

## 1 Introduction

Given an image and a natural language expression that describes an object from the image, the task of referring image segmentation is to predict a pixel-wise mask that delineates that object (Cheng et al. 2014; Hu, Rohrbach, and Darrell 2016). It has applications in a broad range of areas such as image editing, augmented reality, robotics, *etc.* Different from the more conventional semantic/instance segmentation tasks where targets fall into a pre-defined set of categories, this task requires an algorithm to not only predict an accurate mask for the target, but infer what the target is from a free-form language expression.

Taking advantage of a fully convolutional architecture (Long, Shelhamer, and Darrell 2015), current state-of-the-art methods (Liu et al. 2017; Li et al. 2018; Chen et al. 2019; Luo et al. 2020a; Liu et al. 2021a; Ding et al. 2021; Yang et al. 2022) jointly address both requirements by developing

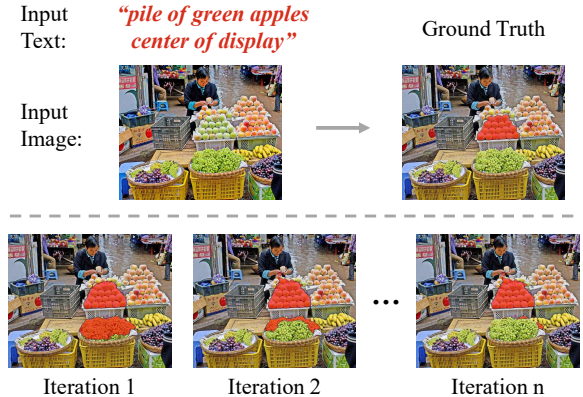


Figure 1: Top half: Referring image segmentation aims to predict a mask delineating a target area described by a language expression. Bottom half: Our method tackles this problem in iterations, where it continuously improves upon its previous prediction based on accumulated object context, finally producing a highly accurate mask.

powerful cross-modal feature fusion methods, which align linguistic meanings and visual cues in a common feature space. Intuitively, learning accurate cross-modal alignments can be a daunting task as abundant noise is present in both the language and the vision inputs. In this context, iterative learning becomes a characteristic approach that is adopted by many methods, which mitigates that difficulty: By performing cross-modal feature fusion in several rounds, alignments can be incrementally established and refined, subsequently leading to higher segmentation accuracy.

Existing state-of-the-art methods generally perform iterative learning via recurrent neural networks (RNNs) (Liu et al. 2017; Li et al. 2018; Chen et al. 2019) or stacked attention layers (Luo et al. 2020a). RNN-based methods exploit the sequential property of the input data, either that of a sentence encoded step by step by a recurrent language model (Liu et al. 2017), or that of an image encoded into a feature pyramid by a hierarchical vision backbone network (Li et al. 2018; Chen et al. 2019). Though having the potential of capturing helpful dependencies, they have several drawbacks. Recurrent language models (*e.g.*, (Sutskever, Vinyals, and Le 2014)) are disadvantaged against the preva-

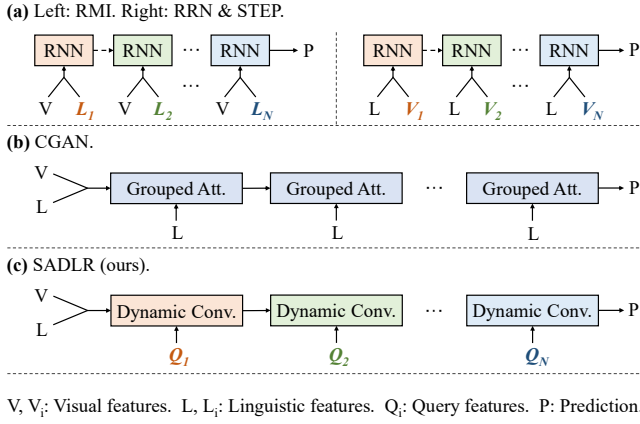


Figure 2: A comparison of iterative learning schemes for referring image segmentation. (a) The recurrent modeling approach exploits the sequential property of sentence features or the hierarchical property of pyramidal image features. (b) The cascade attentions approach employs stacked attention layers where each layer enhances the vision-language mapping from the previous layer. (c) Motivated differently, SADLR performs semantics-aware dynamic convolutions by leveraging continuously updated object context.

lent Transformer-based language models (*e.g.*, (Devlin et al. 2019)), while recurrently integrating multi-scale visual semantics only indirectly addresses cross-modal alignment. Moreover, neither strategy is orthogonal to many of the other state-of-the-art fusion methods (such as ones proposed in VLT (Ding et al. 2021) and LAVT (Yang et al. 2022)). The recent advance by Luo *et al.* introduces cascade grouped attentions (Luo et al. 2020a), which refine cross-modal features of the whole image but lack the ability to exploit prior evidence for helping focus on regions that may need refinement the most—local regions where targets live and mistakes occur frequently. As we observe in our experiments, this technique also does not bring improvements to baselines that adopt powerful Transformer-based fusion methods.

To address those issues, we propose a semantics-aware dynamic localization and refinement (SADLR) method for referring image segmentation. SADLR progressively strengthens target signals and suppresses background noise by leveraging a semantics-aware dynamic convolution module, which performs convolutions conditioning on a continuously updated representation of the target (what we denote as the “query”). To aid localization, in the first step, we initialize the query with a language feature vector summarized from the input expression, and predict a convolutional kernel from the query which operates on the multi-modal feature maps. This step strengthens features according to language information. In each of the following steps, we update the query with pooled object context obtained using the prediction from the previous step and perform dynamic convolution again. As more object context is incorporated along the way, SADLR is able to gradually pinpoint the target and refine segmentation to a high degree of accuracy.

When applied to three state-of-the-art referring image

segmentation methods, namely, LTS (Jing et al. 2021), VLT (Ding et al. 2021), and LAVT (Yang et al. 2022), the proposed SADLR approach is able to bring consistent performance improvements. On the challenging benchmark datasets of RefCOCO (Yu et al. 2016), RefCOCO+ (Yu et al. 2016), and G-Ref (Nagaraja, Morariu, and Davis 2016), our method combined with the most powerful LAVT obtains 74.24%, 64.28%, and 63.60% overall IoU on the validation sets, improving the state of the art for these datasets by absolute margins of 1.51%, 2.01%, and 2.36%, respectively.

## 2 Related Work

**Referring image segmentation** has attracted growing attention in the past decade. While prior arts focus on developing cross-modal feature fusion methods based on CNN backbones (Hu, Rohrbach, and Darrell 2016; Liu et al. 2017; Chen et al. 2019; Feng et al. 2021; Hu et al. 2020; Huang et al. 2020; Hui et al. 2020; Luo et al. 2020a; Margffoy-Tuay et al. 2018; Shi et al. 2018; Ye et al. 2019), recent developments leverage the Transformer architecture (Vaswani et al. 2017) for aggregating vision-language representations, including both at the feature encoding stage (Yang et al. 2022) and at the feature decoding stage (Ding et al. 2021).

With the purpose of mitigating difficulties in cross-modal context modeling, some methods, including RMI (Liu et al. 2017), RRN (Li et al. 2018), STEP (Chen et al. 2019), and CGAN (Luo et al. 2020a), adopt an iterative learning approach and refine cross-modal features in several rounds. Mimicking how humans tackle visual grounding tasks, RMI integrates linguistic features with visual features step by step in a word-reading order, an approach well motivated for RNN language models but ill-suited to the fully-connected Transformer language models. Conversely, RRN and STEP propose to integrate visual features with linguistic features stage by stage traversing through the image feature pyramid. This approach mainly aims at capturing multi-scale semantics from the visual features and has been shown to be inferior to some of the other multi-stage fusion methods such as LSCM (Hui et al. 2020) and CMPC (Liu et al. 2021a). The work most related to ours is CGAN, which leverages cascaded multi-head attention layers for iteratively refining cross-modal mappings of the whole image. In contrast to our method, it does not seek to exploit predictions as priors to help focus refinement on more important regions. Although in principle, the cascade attentions can be applied on top of existing models, we find in experiments that they only benefit the more conventional models which do not employ a powerful Transformer-based architecture for feature fusion.

Albeit our proposal is related to the above class of approaches, it offers important practical advantages. First, it does not rely on a specific choice of the language encoder or the multi-modal feature encoding scheme. Second, it is generalizable to a variety of top-performing models, including Transformer-based ones, as shown in our experiments.

**Dynamic convolution** is a versatile technique that can be applied for detecting user-defined patterns in a set of feature maps. Originally proposed in the context of short-range weather prediction (Klein, Wolf, and Afek 2015), it has recently found popularity in several query-based object detec-




Figure 3: A schematic illustration of the proposed SADLR approach. In the first iteration, the query used for dynamic convolution is initialized by a language feature vector and segmentation mainly exploits language semantics. In each of the following iterations, we continuously update our query with an object feature vector based on the previous iteration’s prediction, and segmentation exploits an increasing amount of object semantics. All parameters are shared across iterations.

tion systems (Sun et al. 2021; Fang et al. 2021). In such models, a set of learnable embeddings serve as latent object features, and dynamic weights are predicted from these queries for localizing and classifying objects. In comparison, we leverage dynamic convolution for solving the problem of segmentation from natural language.

Here, we make a distinction between the types of dynamic convolutions employed in this work and some earlier methods addressing the related task of referring video object segmentation (Gavrilyuk et al. 2018; Li et al. 2017; Wang et al. 2020). In those methods, the term dynamic convolution/filtering refers to predicting classification weights that have the same number of channels as the target feature maps, from which a foreground score map is generated. In contrast, our aim is to produce an enriched set of feature maps that highlight the semantics related to the dynamic query.

**Multi-modal reasoning** is a broader topic that covers many lines of work. For instance, Radford *et al.* focus on the large-scale pre-training of transferable representations using paired image-text data (Radford et al. 2021). Hu *et al.* propose the Unified Transformer model (Hu and Singh 2021) which jointly learns many vision-language tasks such as visual question answering (Antol et al. 2015) and visual entailment (Xie et al. 2019). Kamath *et al.* devise MDETR (Kamath et al. 2021), an end-to-end modulated detector conditioning on language queries, which can be fine-tuned to perform downstream tasks such as few-shot learning for long-tailed detection (Gupta, Dollar, and Girshick 2019).

## 3 Method

### 3.1 Overview

**Preliminaries.** Fig. 3 illustrates our method schematically. The input consists of an image and a language expression. We first extract a set of linguistic features,  $L \in \mathbb{R}^{C_l \times N}$ , from the input expression by leveraging a language encoder (such as BERT (Devlin et al. 2019) or an RNN (Cho et al. 2014)). Here,  $C_l$  denotes the number of channels and  $N$  denotes the number of words. Then, the input image and the linguistic features  $L$  go through a multi-modal feature encoding network, which can be instantiated by most any state-of-the-art referring segmentation model simply by removing its final classification layer. Fig. 3 illustrates our default choice of LAVT (Yang et al. 2022), for which we include more details in Sec. 3.2. This multi-modal feature encoding stage ensures that linguistic features and visual features are joined together in a single feature space that captures alignment information. We denote the output set of multi-modal feature maps as  $Y \in \mathbb{R}^{C \times H \times W}$ , where  $C$  is the number of channels and  $H$  and  $W$  denote the height and the width, respectively.

**Semantics-aware dynamic convolution.** In the ensuing stage, instead of directly producing the object mask via a classifier, we iteratively refine the multi-modal features via a semantics-aware dynamic convolution module, which we further motivate and detail in Sec. 3.3. Each layer of this module predicts a kernel from an input feature vector and convolves multi-modal feature maps with the predicted ker-




Figure 4: Illustration of the semantics-aware dynamic convolution module at iteration  $i$ . The input feature maps  $Y$  are strengthened by two layers of  $1 \times 1$  convolution (implemented by matrix multiplication here), with output  $Z_i$  capturing semantics encoded by the input query  $Q_i$ . ‘MatMul’ denotes matrix multiplication and ‘LN’ denotes layer normalization. This module is different from stacked (cross-)attention layers, for each of which, features at each output position would be computed as an aggregated sum of features at certain input positions based on pair-wise similarities (potentially with locality as a constraint).

nel. In this paper, we call this kernel-generating feature vector a “query,” to give relevance to its functionality as an exemplar that characterizes the notion we wish to highlight in the multi-modal feature maps.

**Pipeline of each iteration.** Assuming there are  $n$  iterations in our refinement procedure, where  $n$  is a hyper-parameter we study in Sec. 4.4, we denote the dynamic query at iteration  $i \in \{1, 2, \dots, n\}$  as  $Q_i \in \mathbb{R}^C$ , where  $C$  denotes the number of channels. At each iteration, we send the query  $Q_i$  and the multi-modal feature maps  $Y$  (from the encoding stage) into the aforementioned semantics-aware dynamic convolution module, where a kernel is derived from  $Q_i$  and convolved with  $Y$  for producing a set of feature maps that highlight notions relevant to  $Q_i$ . We denote the output feature maps as  $Z_i \in \mathbb{R}^{C \times H \times W}$ , where  $H$  and  $W$  denote the height and the width, respectively, and  $C$  denotes the number of channels. Next, a  $1 \times 1$  convolution projects  $Z_i$  into raw score maps, denoted as  $R_i \in \mathbb{R}^{2 \times H \times W}$ , where indices 0 and 1 along the channel dimension define the score maps for the “background” class and the “object” class, respectively. We obtain a binary object mask  $M_i \in \mathbb{R}^{H \times W}$  by taking *argmax* along the channel dimension of  $R_i$ .

**Query initialization and iterative update.** In different iterations, we use different queries with the purpose of highlighting different types of features. To start with, at iteration 1, we initialize  $Q_1$  as a sentence feature vector,  $S \in \mathbb{R}^C$ , which is obtained from  $L$  (from the language encoder) via average pooling across the word dimension followed by linear projection for channel reduction. Since  $Q_1$  contains pure linguistic information that summarizes the target object, the output from the semantics-aware dynamic convolution module  $Z_1$  highlights referring localization information. At iteration  $i$ ,  $i \in \{2, 3, \dots, n\}$ , we update  $Q_{i-1}$  with an object feature vector  $O_{i-1}$  pooled from  $Y$  using  $M_{i-1}$  to obtain a new query  $Q_i$  for the current iteration. This process can be described mathematically as follows

$$O_{i-1} = \text{AvgPool}(M_{i-1}, Y), \quad (1)$$

$$Q_i = Q_{i-1} + O_{i-1}, \quad (2)$$

where ‘AvePool’ denotes computing a weighted average of feature vectors across the spatial dimensions of  $Y$ , with the binary mask  $M_{i-1}$  being the weight map. This query update scheme aims to progressively integrate more object

context information into the query as the predicted mask  $M_{i-1}$  grows more accurate by each iteration. As mentioned earlier, dynamic convolution strengthens features related to the given query. Therefore, as the query grows from being purely linguistic to containing more object-centric features, the algorithm gradually shifts focus from referring localization to segmentation refinement. As shown in the qualitative analysis in Fig. 5, the predicted masks tend to grow more refined by each iteration.

### 3.2 Multi-modal Feature Encoding

By default, we adopt LAVT (Yang et al. 2022) as the multi-modal feature encoding network for its top performance in this task. LAVT leverages a powerful hierarchical vision Transformer (Liu et al. 2021b) to jointly embed vision and language information. At each stage of the vision backbone, a pixel-word attention module densely aligns linguistic features with visual features at each spatial location, and a language pathway directs integrated multi-modal cues into the next stage of the vision backbone. Our proposed method is not limited to a specific multi-modal feature encoding network. As we show in Table 3, when adopting LTS (Jing et al. 2021) or VLT (Ding et al. 2021) as the encoding network, SADLR consistently obtains improved results with respect to the vanilla model.

### 3.3 Semantics-Aware Dynamic Convolution

Much like a traditional convolution, dynamic convolution convolves an input set of feature maps with the same kernel across all locations. However, in dynamic convolution, the kernel is generated from a conditioning feature vector (what we call the query in this paper) instead of being fixed model parameters. It particularly suits our need in referring image segmentation, as the generated kernel is sample-specific.

Fig. 4 illustrates our semantics-aware dynamic convolution module. It consists of two layers of dynamic convolutions. At each iteration of SADLR, given query  $Q_i$ , a first linear function generates a dynamic kernel,  $K_i \in \mathbb{R}^{C \times C'}$ , where  $C$  and  $C'$  denote the input and the output numbers of channels, respectively. Then convolution is performed via matrix multiplication between the input feature maps  $Y$  and the dynamic kernel  $K_i$ , followed by layer normalization (Ba,



	RefCOCO			RefCOCO+			G-Ref		
	val	test A	test B	val	test A	test B	val (U)	test (U)	val (G)
DMN (Margffoy-Tuay et al. 2018)	49.78	54.83	45.13	38.88	44.22	32.29	-	-	36.76
RRN (Li et al. 2018)	55.33	57.26	53.93	39.75	42.15	36.11	-	-	36.45
MAttNet (Yu et al. 2018)	56.51	62.37	51.70	46.67	52.39	40.08	47.64	48.61	-
CMSA (Ye et al. 2019)	58.32	60.61	55.09	43.76	47.60	37.89	-	-	39.98
STEP (Chen et al. 2019)	60.04	63.46	57.97	48.19	52.33	40.41	-	-	46.40
BRINet (Hu et al. 2020)	60.98	62.99	59.21	48.17	52.32	42.11	-	-	48.04
CMPC (Huang et al. 2020)	61.36	64.53	59.64	49.56	53.44	43.23	-	-	49.05
LSCM (Hui et al. 2020)	61.47	64.99	59.55	49.34	53.12	43.50	-	-	48.05
CMPC+ (Liu et al. 2021a)	62.47	65.08	60.82	50.25	54.04	43.47	-	-	49.89
MCN (Luo et al. 2020b)	62.44	64.20	59.71	50.62	54.99	44.69	49.22	49.40	-
EFN (Feng et al. 2021)	62.76	65.69	59.67	51.50	55.24	43.01	-	-	51.93
BUSNet (Yang et al. 2021)	63.27	66.41	61.39	51.76	56.87	44.13	-	-	50.56
CGAN (Luo et al. 2020a)	64.86	68.04	62.07	51.03	55.51	44.06	51.01	51.69	46.54
LTS (Jing et al. 2021)	65.43	67.76	63.08	54.21	58.32	48.02	54.40	54.25	-
VLT (Ding et al. 2021)	65.65	68.29	62.73	55.50	59.20	49.36	52.99	56.65	49.76
CRIS (Wang et al. 2021)	70.47	73.18	66.10	62.27	68.08	53.68	59.87	60.36	-
LAVT (Yang et al. 2022)	72.73	75.82	68.79	62.14	68.38	55.10	61.24	62.09	60.50
Ours	<b>74.24</b>	<b>76.25</b>	<b>70.06</b>	<b>64.28</b>	<b>69.09</b>	<b>55.19</b>	<b>63.60</b>	<b>63.56</b>	<b>61.16</b>

Table 1: Comparison with other methods in terms of overall IoU. “U” and “G” represent UMD and Google, respectively.

Kiros, and Hinton 2016) and ReLU non-linearity (Nair and Hinton 2010). Another convolutional kernel,  $K'_i \in \mathbb{R}^{C' \times C'}$ , is subsequently generated from  $Q_i$  by a second linear function. Similarly, convolution is performed between the output from the previous layer and  $K'_i$  via matrix multiplication, followed by layer normalization and ReLU non-linearity.

### 3.4 Predicted Masks and the Loss Function

To produce segmentation masks, we upsample the raw score maps  $R_i$  via bilinear interpolation to the resolution of the input image, and take the *argmax* along the channel dimension of the upsampled score maps. The loss function used for training is the following,

$$L = \lambda_1 L_1 + \lambda_2 L_2 + \dots + \lambda_n L_n, \quad (3)$$

where  $n$  is the number of iterations,  $L_i$ ,  $i \in \{1, 2, \dots, n\}$ , denotes the individual loss from iteration  $i$ , and  $\lambda_i$ ,  $i \in \{1, 2, \dots, n\}$ , is the balancing weight for loss at iteration  $i$ . Each individual loss is computed as the average Dice losses (Milletari, Navab, and Ahmadi 2016) for the “object” class and the “background” class. During inference, mask from the last iteration is used as the prediction.

## 4 Experiments

### 4.1 Datasets and Evaluation Metrics

**Datasets.** We evaluate our proposed method on the datasets of RefCOCO (Yu et al. 2016), RefCOCO+ (Yu et al. 2016), and G-Ref (Mao et al. 2016; Nagaraja, Morariu, and Davis 2016). All three datasets collect images from the MS COCO dataset (Lin et al. 2014), and each object is annotated with multiple language descriptions from humans. RefCOCO and RefCOCO+ are similar in size, both containing about 20K images, 50K annotated objects, and 140K expressions. They

are annotated in a two-player game, where each player is motivated to provide a minimally sufficient description for the other to identify the target. As a result, expressions from these two datasets are succinct, averaging about 4 words per expression, and emphasize simple traits such as color, size, and position. A special property about RefCOCO+ is that location words (such as “left” and “right”) are banned, making it the harder of the two. G-Ref contains roughly 27K images, 55K annotated objects, and 105K expressions. Expressions in G-Ref tend to be longer, averaging around 8 words per expression, and provide more complete descriptions of the target objects. As the expressions can be elaborate, G-Ref is the most challenging dataset of the three. We evaluate our algorithm on both the UMD partition (Nagaraja, Morariu, and Davis 2016) and the Google partition (Mao et al. 2016).

**Evaluation metrics.** We adopt three sets of metrics: precision@K (P@K, where K indicates an IoU threshold), mean IoU (mIoU), and overall IoU (oIoU). The precision@K metric measures the percentage of test samples (*i.e.*, image-sentence pairs) that pass an IoU threshold, where the IoU is computed between the prediction and the ground truth. We evaluate precision at the thresholds of 0.5, 0.6, 0.7, 0.8, and 0.9. On a separate note, this metric should be more accurately called “recall@K.” The mean IoU metric is the average IoU of all test samples. This metric treats large and small objects equally. The overall IoU is the intersection accumulated on all test samples divided by the union accumulated on all test samples. It favors large objects.

### 4.2 Implementation Details

As our best SADLR model is based on LAVT, most training and inference settings follow those in (Yang et al. 2022). We adopt the BERT-base model from (Devlin et al. 2019) as the language encoder, and the Swin-B model from (Liu

	RefCOCO			RefCOCO+			G-Ref		
	val	test A	test B	val	test A	test B	val-U	test-U	val-G
RefTrans	74.34	76.77	70.87	66.75	70.58	59.40	66.63	67.39	-
Ours	<b>76.52</b>	<b>77.98</b>	<b>73.49</b>	<b>68.94</b>	<b>72.71</b>	<b>61.10</b>	<b>67.47</b>	<b>67.73</b>	<b>65.21</b>

Table 2: Comparison with Referring Transformer (Li and Leonid 2021) by mean IoU.

et al. 2021b) as the backbone network. BERT initialization weights are obtained from HuggingFace (Wolf et al. 2020), and Swin initialization weights are ones pre-trained on ImageNet22K from (Liu et al. 2021b). The rest of the parameters in our model are randomly initialized. The entire framework is trained end-to-end with the loss function defined in Eq. 3. We adopt an AdamW (Loshchilov and Hutter 2019) optimizer with initial learning rate  $5e-5$  and weight decay  $1e-2$ , and apply the “poly” learning rate scheduler (Chen et al. 2018). The numbers of channels  $C_l$  and  $C$  in Sec. 3 are 768 and 512, respectively. The default number of iterations ( $n$  in Sec. 3) is 3, for which the loss weights,  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$ , are 0.15, 0.15, 0.7, respectively. For each  $n$ , we search for optimal balancing weights as follows. While ensuring that the sum of the weights is 1, we experiment with large values ( $\geq 0.5$ ) on the last iteration, each time assigning equal weights to the rest of the iterations. For each dataset, the model is trained on the training set for 40 epochs with batch size 32, where each object is sampled exactly once in an epoch (with one of its text annotations randomly sampled). Images are resized to  $480 \times 480$  resolution and sentence lengths are capped at 20.

### 4.3 Comparison with Others

In Table 1, we evaluate our proposed method against the state-of-the-art methods on the popular benchmarks of RefCOCO, RefCOCO+, and G-Ref. Our approach obtains the highest overall IoU on all subsets of all benchmark datasets with various margins. On the validation, test A, and test B subsets of RefCOCO, our method surpasses the second-best LAVT by margins of 1.51%, 0.43%, and 1.27%, respectively. On RefCOCO+, our method establishes a relatively large 2.01% (absolute) lead with respect to the second-best method CRIS (Wang et al. 2021) on the validation set, while also comparing favorably against the second-best LAVT on the test A and test B subsets with small advantages. On both the validation and test sets of G-Ref (UMD partition), our approach improves the state of the art by solid margins of 2.36% and 1.47%, respectively. On the most challenging validation set of G-Ref (Google partition), our method is also able to achieve new state-of-the-art results with a small improvement of 0.66% (absolute). In Table 2, we report the mean IoU of our method for comparison with Referring Transformer (Li and Leonid 2021), as mean IoU instead of overall IoU is reported in (Li and Leonid 2021). Our method obtains better results on all three datasets.

### 4.4 Ablation Study

In this section, we evaluate the effectiveness of the SADLR method and study the effects of several key design choices.

Meth.	P@0.5	P@0.6	P@0.7	P@0.8	P@0.9	oIoU	mIoU	FLOPs	Ti.
LTS	82.10	77.51	70.94	58.54	27.14	70.84	71.92	133.3G	41
" + C	84.52	80.72	74.89	61.57	29.13	71.39	73.57	143.3G	49
" + S	<b>85.20</b>	<b>81.79</b>	<b>76.39</b>	<b>64.92</b>	<b>32.28</b>	<b>72.22</b>	<b>74.49</b>	139.4G	47
VLT	83.80	79.96	74.10	60.52	25.60	70.31	72.62	142.6G	42
" + C	80.33	76.33	69.94	56.32	22.80	66.41	69.62	152.7G	49
" + S	<b>85.00</b>	<b>81.85</b>	<b>76.17</b>	<b>62.44</b>	<b>27.81</b>	<b>71.64</b>	<b>73.90</b>	148.6G	48
LAVT	85.49	82.00	76.44	65.65	35.16	73.12	75.30	197.4G	46
" + C	85.14	81.37	75.98	64.57	34.35	72.56	75.05	207.5G	55
" + S	<b>86.90</b>	<b>83.68</b>	<b>78.76</b>	<b>67.93</b>	<b>37.36</b>	<b>74.24</b>	<b>76.52</b>	203.5G	52

Table 3: Comparison between our method, CGAN, and several baseline methods on the RefCOCO validation set. “+ C” or “+ S” denotes that we apply CGAN or SADLR on top of the baseline method represented by the ditto mark, respectively. As described in the text below, all experiments adopt BERT as the language encoder and Swin-B as the backbone network, and follow the same training settings.

The experiment setups, including the language encoder, the visual backbone network, the loss, *etc.*, are kept the same in all studies for fair comparisons. As a result, there is a difference between our reproduced overall IoU of LAVT (reported in Tables 3 and 4) and that of LAVT from the original paper (reported in Table 1). This difference mainly originates from the Dice loss that we adopt in all the experiments—the original LAVT adopts a cross-entropy loss.

**Comparison with CGAN and baseline methods.** In Table 3, we apply our iterative dynamic refinement procedure on three different state-of-the-art methods, namely, LTS (Jing et al. 2021), VLT (Ding et al. 2021), and LAVT (Yang et al. 2022), and show that our method can generalize to all three architectures. When using either the Transformer-based VLT or LAVT as the baseline method (the multi-modal feature encoding network in Sec. 3), our method consistently brings 1 to 2 absolute points of improvement on all of the seven evaluation metrics, including precision at the five threshold values, the overall IoU and the mean IoU. On the non-Transformer-based LTS, our proposal achieves still greater improvements. From P@0.5 to P@0.9, our method achieves 3.10%, 4.28%, 5.45%, 6.38%, and 5.14% absolute improvements, respectively. The gain generally increases as the threshold gets larger, which highlights the efficacy of our method in refining segmentation masks to a high degree of accuracy. Moreover, on the overall IoU and mean IoU, our method attains 1.38% and 2.57% absolute improvements, respectively. Overall, the results above highlight the generality and effectiveness of SADLR.

On the contrary, CGAN (Luo et al. 2020a) is only able to bring moderate improvement to the more conventional LTS model, which is the best model among ones not employing Transformers for cross-modal feature fusion. When the more effective Transformer architecture has been employed, applying CGAN leads to performance deterioration, as shown by the “VLT + CGAN” and “LAVT + CGAN” variants.

To analyze the computational overhead of SADLR, we report the FLOPs and the inference time (in milliseconds) of all models in Table 3. We measure the inference time by averaging over 500 forward passes using batch size 1

	P@0.5	P@0.6	P@0.7	P@0.8	P@0.9	oIoU	mIoU
<b>(a) number of iterations</b>							
0	85.49	82.00	76.44	65.65	35.16	73.12	75.30
1	85.44	82.01	76.68	65.92	35.86	73.22	75.41
2	86.63	83.63	78.48	67.83	37.04	73.99	76.32
3 (*)	<b>86.90</b>	<b>83.68</b>	<b>78.76</b>	67.93	<b>37.36</b>	<b>74.24</b>	<b>76.52</b>
4	86.64	83.12	77.77	<b>68.01</b>	36.87	73.56	76.09
<b>(b) structure of the semantics-aware dynamic convolution module</b>							
[128]	86.26	82.84	77.74	67.36	36.69	72.98	75.87
[256]	86.50	83.27	78.38	67.37	37.16	73.62	76.22
[64, 256]	<b>87.00</b>	<b>83.88</b>	78.33	67.75	37.02	73.76	76.34
[512]	86.33	83.21	78.27	67.53	36.73	73.56	76.02
[128, 512] (*)	86.90	83.68	<b>78.76</b>	<b>67.93</b>	<b>37.36</b>	<b>74.24</b>	<b>76.52</b>
[128, 512] × 2	86.47	83.27	78.09	67.19	36.87	73.70	76.12
<b>(c) query update method</b>							
sum (*)	<b>86.90</b>	<b>83.68</b>	<b>78.76</b>	<b>67.93</b>	<b>37.36</b>	<b>74.24</b>	<b>76.52</b>
replace	86.52	83.48	78.36	67.09	37.35	73.19	76.14

Table 4: Ablation studies on the RefCOCO validation set. Rows with (\*) indicate our default choices. In each experiment, we change only one variable from our default configuration and study its effect on the overall model.

at  $480 \times 480$  input resolution on an NVIDIA Quadro RTX 8000. SADLR brings 3.1%, 4.6%, and 4.2% relative increase of FLOPs and 13.0%, 14.6%, and 14.3% relative increase of inference time to LAVT, LTS, and VLT, respectively, all of which are smaller than those of CGAN.

**Number of iterations.** Next, we study the optimal number of iterations in SADLR. As shown in Table 4 (a), as the number of iterations grows, results continuously improve, with optimal ones obtained when there are three iterations. Adding a fourth iteration generally does not yield better results (except marginally so on P@0.8). Applying 0 iteration equals to using the baseline method (and in these experiments, LAVT). These observations are consistent with our intuition that refinement may eventually saturate but before saturation, applying more iterations is beneficial.

**Structure of the semantics-aware dynamic convolution module.** Our default implementation of this module consists of two convolution layers. In Table 4 (b), we study the effects of the number of convolution layers and the number of filters (the number of output channels) for each layer. We define  $[a]$  as one convolution layer with  $a$  filters and  $[a, b]$  as two convolution layers with  $a$  and  $b$  filters, respectively, where  $a$  and  $b$  are numbers. Similarly, we define  $[a, b] \times k$  as a stack of  $k$   $[a, b]$  sub-structures, where  $k$  is a number. As shown in Table 4 (b), the structure  $[128, 512]$  produces the best results in five of the seven metrics, and is chosen as our default implementation.  $[64, 256]$  comes as the overall second-best structure and produces the best P@0.5 and P@0.6. We note that when using only one layer of convolution, having 256 filters works better than having 512 filters (shown via a comparison between the “[256]” and “[512]” variants). Conversely, when using two layers of convolutions, having more filters prevails over having fewer filters (shown via a comparison between the “[128, 512]” and “[64, 256]” variants). This suggests that adopting more filters can be beneficial, but only likely so

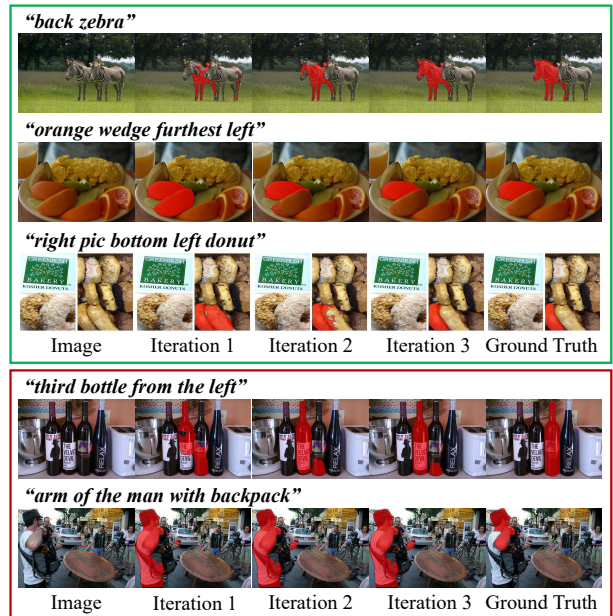


Figure 5: Visualizations of predictions from each iteration of SADLR (our default 3-iteration model). In the success cases (inside the green box), the first iteration predicts a mask that roughly locates the intended object and the following iterations correct the missing and/or extra parts progressively. In the failure cases (inside the red box), refinement fails as the prediction from the first iteration contains too much noise—that is, when localization is too bad. This is a fundamental challenge and solving it requires still more effort devoted to addressing cross-modal understanding.

when more layers are employed which facilitates learning. In addition, we show that repeating our best  $[128, 512]$  structure once (leading to a total of four convolution layers) does not bring further improvements.

**Query update method.** In Table 4 (c), we compare summation with an alternative replacement strategy when updating the query across iterations. This corresponds to changing Eq. 2 to  $Q_i = O_{i-1}$ , where the object feature vector replaces the previous query as the new one. Table 4 (c) shows that replacement leads to inferior results, which indicates that it is beneficial to retain language information in the query.

## 5 Conclusion

In this paper, we have proposed a semantics-aware dynamic localization and refinement (SADLR) method for the referring image segmentation problem, which iteratively refines multi-modal feature maps based on aggregated object context. Extensive experiments on three baseline methods demonstrate the effectiveness and generality of the proposed method. And evaluations on three standard benchmarks demonstrate its advantage with respect to the state-of-the-art methods. We hope that our method could inspire further interest in the development of iterative learning approaches for referring image segmentation and be extended to related tasks in the future.

## Acknowledgments

This work is supported by the UKRI grant: Turing AI Fellowship EP/W002981/1, EPSRC/MURI grant: EP/N019474/1, Shanghai Committee of Science and Technology, China (Grant No. 20DZ1100800), National Natural Science Foundation of China (Grant No. 62206153), HKU Startup Fund, HKU Seed Fund for Basic Research, and Shanghai AI Laboratory. We would also like to thank the Royal Academy of Engineering and FiveAI.

## References

- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. VQA: Visual Question Answering. In *ICCV*.
- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450*.
- Chen, D.-J.; Jia, S.; Lo, Y.-C.; Chen, H.-T.; and Liu, T.-L. 2019. See-Through-Text Grouping for Referring Image Segmentation. In *ICCV*.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *ECCV*.
- Cheng, M.-M.; Zheng, S.; Lin, W.-Y.; Vineet, V.; Sturgess, P.; Crook, N.; Mitra, N. J.; and Torr, P. 2014. ImageSpirit: Verbal Guided Image Parsing. In *TOG*.
- Cho, K.; van Merriënboer, B.; Gülçehre, Ç.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*.
- Ding, H.; Liu, C.; Wang, S.; and Jiang, X. 2021. Vision-Language Transformer and Query Generation for Referring Segmentation. In *ICCV*.
- Fang, Y.; Yang, S.; Wang, X.; Li, Y.; Fang, C.; Shan, Y.; Feng, B.; and Liu, W. 2021. Instances As Queries. In *ICCV*.
- Feng, G.; Hu, Z.; Zhang, L.; and Lu, H. 2021. Encoder Fusion Network With Co-Attention Embedding for Referring Image Segmentation. In *CVPR*.
- Gavrilyuk, K.; Ghodrati, A.; Li, Z.; and Snoek, C. G. 2018. Actor and Action Video Segmentation from a Sentence. In *CVPR*.
- Gupta, A.; Dollar, P.; and Girshick, R. 2019. LVIS: A Dataset for Large Vocabulary Instance Segmentation. In *CVPR*.
- Hu, R.; Rohrbach, M.; and Darrell, T. 2016. Segmentation from Natural Language Expressions. In *ECCV*.
- Hu, R.; and Singh, A. 2021. UniT: Multimodal Multitask Learning with a Unified Transformer. In *ICCV*.
- Hu, Z.; Feng, G.; Sun, J.; Zhang, L.; and Lu, H. 2020. Bi-Directional Relationship Inferring Network for Referring Image Segmentation. In *CVPR*.
- Huang, S.; Hui, T.; Liu, S.; Li, G.; Wei, Y.; Han, J.; Liu, L.; and Li, B. 2020. Referring Image Segmentation via Cross-Modal Progressive Comprehension. In *CVPR*.
- Hui, T.; Liu, S.; Huang, S.; Li, G.; Yu, S.; Zhang, F.; and Han, J. 2020. Linguistic Structure Guided Context Modeling for Referring Image Segmentation. In *ECCV*.
- Jing, Y.; Kong, T.; Wang, W.; Wang, L.; Li, L.; and Tan, T. 2021. Locate Then Segment: A Strong Pipeline for Referring Image Segmentation. In *CVPR*.
- Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. MDETR – Modulated Detection for End-to-End Multi-Modal Understanding. In *ICCV*.
- Klein, B.; Wolf, L.; and Afek, Y. 2015. A Dynamic Convolutional Layer for Short Range Weather Prediction. In *CVPR*.
- Li, M.; and Leonid, S. 2021. Referring Transformer: A One-step Approach to Multi-task Visual Grounding. In *NIPS*.
- Li, R.; Li, K.; Kuo, Y.-C.; Shu, M.; Qi, X.; Shen, X.; and Jia, J. 2018. Referring Image Segmentation via Recurrent Refinement Networks. In *CVPR*.
- Li, Z.; Tao, R.; Gavves, E.; Snoek, C. G.; and Smeulders, A. W. 2017. Tracking by Natural Language Specification. In *CVPR*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In Fleet, D.; Pajdla, T.; Schiele, B.; and Tuytelaars, T., eds., *ECCV*.
- Liu, C.; Lin, Z.; Shen, X.; Yang, J.; Lu, X.; and Yuille, A. 2017. Recurrent Multimodal Interaction for Referring Image Segmentation. In *ICCV*.
- Liu, S.; Hui, T.; Huang, S.; Wei, Y.; Li, B.; and Li, G. 2021a. Cross-Modal Progressive Comprehension for Referring Segmentation. In *TPAMI*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows. In *ICCV*.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully Convolutional Networks for Semantic Segmentation. In *CVPR*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Luo, G.; Zhou, Y.; Ji, R.; Sun, X.; Su, J.; Lin, C.-W.; and Tian, Q. 2020a. Cascade Grouped Attention Network for Referring Expression Segmentation. In *ACMMM*.
- Luo, G.; Zhou, Y.; Sun, X.; Cao, L.; Wu, C.; Deng, C.; and Ji, R. 2020b. Multi-Task Collaborative Network for Joint Referring Expression Comprehension and Segmentation. In *CVPR*.
- Mao, J.; Huang, J.; Toshev, A.; Camburu, O.; Yuille, A. L.; and Murphy, K. 2016. Generation and Comprehension of Unambiguous Object Descriptions. In *CVPR*.
- Margffoy-Tuay, E.; Pérez, J. C.; Botero, E.; and Arbeláez, P. 2018. Dynamic Multimodal Instance Segmentation Guided by Natural Language Queries. In *ECCV*.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In *3DV*.
- Nagaraja, V. K.; Morariu, V. I.; and Davis, L. S. 2016. Modeling Context Between Objects for Referring Expression Understanding. In *ECCV*.



Nair, V.; and Hinton, G. E. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *ICML*.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*.

Shi, H.; Li, H.; Meng, F.; and Wu, Q. 2018. Key-Word-Aware Network for Referring Expression Image Segmentation. In *ECCV*.

Sun, P.; Zhang, R.; Jiang, Y.; Kong, T.; Xu, C.; Zhan, W.; Tomizuka, M.; Li, L.; Yuan, Z.; Wang, C.; et al. 2021. Sparse R-CNN: End-to-End Object Detection with Learnable Proposals. In *CVPR*.

Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to Sequence Learning with Neural Networks. In *NeurIPS*.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. In *NeurIPS*.

Wang, H.; Deng, C.; Ma, F.; and Yang, Y. 2020. Context Modulated Dynamic Networks for Actor and Action Video Segmentation with Language Queries. In *AAAI*.

Wang, Z.; Lu, Y.; Li, Q.; Tao, X.; Guo, Y.; Gong, M.; and Liu, T. 2021. CRIS: CLIP-Driven Referring Image Segmentation. *arXiv:2111.15174*.

Wolf, T.; Chaumond, J.; Debut, L.; Sanh, V.; Delangue, C.; Moi, A.; Cistac, P.; Funtowicz, M.; Davison, J.; Shleifer, S.; et al. 2020. Transformers: State-Of-The-Art Natural Language Processing. In *EMNLP*.

Xie, N.; Lai, F.; Doran, D.; and Kadav, A. 2019. Visual Entailment: A Novel Task for Fine-Grained Image Understanding. In *arXiv:1901.06706*.

Yang, S.; Xia, M.; Li, G.; Zhou, H.; and Yu, Y. 2021. Bottom-Up Shift and Reasoning for Referring Image Segmentation. In *CVPR*.

Yang, Z.; Wang, J.; Tang, Y.; Chen, K.; Zhao, H.; and Torr, P. H. S. 2022. LAVT: Language-Aware Vision Transformer for Referring Image Segmentation. In *CVPR*.

Ye, L.; Rochan, M.; Liu, Z.; and Wang, Y. 2019. Cross-Modal Self-Attention Network for Referring Image Segmentation. In *CVPR*.

Yu, L.; Lin, Z.; Shen, X.; Yang, J.; Lu, X.; Bansal, M.; and Berg, T. L. 2018. MAttNet: Modular Attention Network for Referring Expression Comprehension. In *CVPR*.

Yu, L.; Poirson, P.; Yang, S.; Berg, A. C.; and Berg, T. L. 2016. Modeling Context in Referring Expressions. In *ECCV*.