

SAFARI: Adaptive Sequence Transformer for Weakly Supervised Referring Expression Segmentation

Sayan Nag^{1,2*}, Koustava Goswami², and Srikrishna Karanam²

¹ University of Toronto

`sayan.nag@mail.utoronto.ca`

² Adobe Research

`{koustavag, skaranam}@adobe.com`

Abstract. Referring Expression Segmentation (RES) aims to provide a segmentation mask of the target object in an image referred to by the text (i.e., referring expression). Existing methods require large-scale mask annotations. Moreover, such approaches do not generalize well to unseen/zero-shot scenarios. To address the aforementioned issues, we propose a weakly-supervised bootstrapping architecture for RES with several new algorithmic innovations. To the best of our knowledge, ours is the first approach that considers only a fraction of both mask and box annotations (shown in Figure 1 and Table 1) for training. To enable principled training of models in such low-annotation settings, improve image-text region-level alignment, and further enhance spatial localization of the target object in the image, we propose Cross-modal Fusion with Attention Consistency module. For automatic pseudo-labeling of unlabeled samples, we introduce a novel Mask Validity Filtering routine based on a spatially aware zero-shot proposal scoring approach. Extensive experiments show that with just 30% annotations, our model SAFARI achieves 59.31 and 48.26 mIoUs as compared to 58.93 and 48.19 mIoUs obtained by the fully-supervised SOTA method SeqTR respectively on RefCOCO+@testA and RefCOCO+testB datasets. SAFARI also outperforms SeqTR by 11.7% (on RefCOCO+testA) and 19.6% (on RefCOCO+testB) in a fully-supervised setting and demonstrates strong generalization capabilities in unseen/zero-shot tasks. Our project page can be found at https://sayannag.github.io/safari_eccv2024/.

Keywords: Weakly-Supervised Referring Expression Segmentation · Attention Consistency · Mask Validity Filtering

1 Introduction

We consider the problem of referring expression segmentation (RES) [3, 7, 14, 26, 28, 30, 38, 45, 50, 53, 54, 58] where given an image and a text sentence, the goal is to segment out parts of the image the text is referring to. Naturally, much existing

* Work done during internship at Adobe Research.

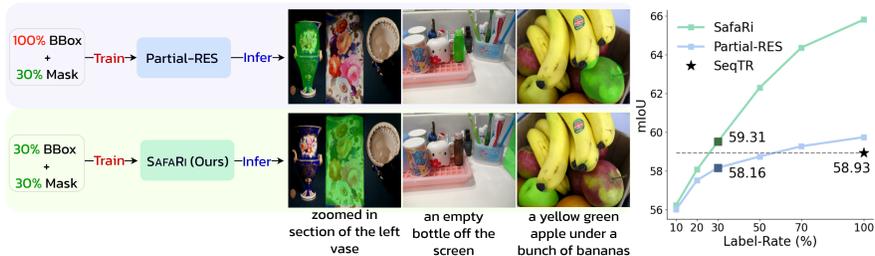


Fig. 1: SAFARI achieves state-of-the-art performance in both weakly- and fully-supervised RES tasks. Although unlike Partial-RES, SAFARI is not pretrained on fully-supervised REC task, with just 30% annotations, SAFARI achieves 59.31 mIoU whereas Partial-RES and fully-supervised SeqTR obtains 58.16 and 58.93 mIoUs (see mIoU vs Label-Rate plot). In the weak-supervision setting, the inclusion of X-FACT (with cross-modal fusion and AMCR components) and SPARC modules aids SAFARI to demonstrate excellent grounding capabilities under challenging scenarios where Partial-RES fails (see qualitative examples). Quantitative results are provided in Tables 2-3.

work in RES blends vision-language understanding [5, 9, 27, 35, 36, 42, 52, 57] and instance segmentation [1, 4, 6, 12, 31], allowing for object segmentation using free-form textual expressions rather than predefined categories [7, 10, 13, 15, 16, 21, 48, 51]. Existing RES methods are predominantly *fully-supervised* in nature [3, 7, 7, 10, 13, 15, 28, 30, 45, 50, 53, 58]. Hence, they necessarily require expensive-to-obtain and extensive manually-annotated segmentation Ground-Truth (GT) masks for training and typically do not generalize well to unseen/zero-shot scenarios.

As noted above, obtaining large-scale GT masks is very expensive and tedious which makes it challenging to scale up such methods. To address this, a recent work (Partial-RES) [41] propose a partially (weakly) supervised solution for RES task. However, they consider abundant (100%) bounding box annotations to begin with (Figure 1 and Table 1) and pretrain the model with these boxes on the Referring Expression Comprehension (REC) task (text referred box prediction) in a *fully-supervised* manner. Moreover, they use these 100% box annotations for filtering pseudo-labels. Their limitations are: (i) in a more pragmatic and realistic weakly-supervised setting *the percentage of boxes should equal the percentage of masks*, (ii) in their pretraining stage the model is already aware of the grounding information of the same dataset used in the weak-supervision stage, (iii) their approach does not account for the cross-modal region-level interaction between the image and language features which is crucial for localization tasks [24, 56].

On the contrary, we investigate a more realistic, challenging and unexplored problem of Weakly-Supervised Referring Expression Segmentation (WSRES) with limited human-annotated mask and box annotations, specifically where ***box % equals mask %*** (Figure 1, see Table 1 for comparisons). To tackle our proposed novel WSRES task, we present **SAFARI**, an auto-regressive contour-prediction-based RES method capable of demonstrating excellent performance under challenging scenarios with few available mask and box annotations. We

specifically choose segmentation as an auto-regressive point prediction task along the object’s contour instead of an independent pixel prediction approach. This is because the latter does not account for relationships between neighbouring pixels and therefore lacks structural information of the object being segmented, resulting in poor performances in the absence of abundant mask annotations [41].

In SAFARI, firstly, we introduce a novel Cross-modal (**X-**) **F**usion with **A**ttention **C**onsistency (**X-FACT**) module for facilitating excellent inter-domain alignment. Since we have only few annotated data, we make the cross-fusion component in X-FACT to be flexible and parameter efficient. Further, in X-FACT we actively leverage the cross-attention heatmaps by formulating a regularization technique that encourages these heatmaps to be consistent with the referred object in the image. This is enabled by confining the attended regions within the object contour, enhancing the fidelity of predicted masks. This is particularly beneficial for the limited annotation scenarios (in WSRES) where it explicitly helps features improving visual grounding, without relying on lots of data. Secondly, we systematically devise a novel bootstrapping strategy (Figure 2) which uses a handful of labeled masks (e.g., 10%) and iteratively trains the model by utilizing pseudo-masks obtained from a pseudo-labeling procedure. Finally, in contrast with conventional pseudo-labeling pipelines which directly use the model predictions as pseudo-labels, we design a new *Mask Validity Filtering (MVF)* routine that is responsible for selection of pseudo-masks (for unannotated data) by validating whether they spatially align with the boundaries (boxes) of the referred objects or not. Since we do not possess 100% bounding boxes, we propose **SPARC**, a novel REC technique with spatial reasoning capabilities for obtaining these boxes in an entirely *zero-shot* manner (Figure 3). Equipped with the above modules, unlike Partial-RES, our system learns meaningful, transferable and generalizable representations with rich semantic understanding, empowering it to reason and make accurate predictions on unseen data (Table 3).

We summarize our **main contributions** as: **(i)** To the best of our knowledge, ours is the first to consider an *accurate representation* of Weakly-Supervised Referring Expression (WS-RES) task by considering a novel, more practical and challenging scenario with limited box and mask annotations where box % equals mask %. **(ii)** The novel X-FACT module fosters prediction of high quality masks by improving *cross-modal alignment* quality, especially where abundant ground-truth annotations are not present. **(iii)** Utilizing SPARC, a novel zero-shot REC technique, the mask validity filtering stage together with the bootstrapping pipeline improve system’s *self-labeling capabilities*. **(iv)** Extensive experiments demonstrate the efficacy of SAFARI as it *significantly outperforms* baseline models on RES benchmarks. SAFARI also demonstrates *strong generalization capabilities* when evaluated on an *unseen* referring video object segmentation task in a *zero-shot* manner. Notably, with just 30% mask-annotated data, SAFARI achieves 59.31 and 55.83 mIoUs versus 58.93 and 55.64 mIoUs obtained by the fully-supervised SOTA method SeqTR [58] on RefCOCO+@testA and RefCOCOg@test sets (Table 2 and Figure 1).

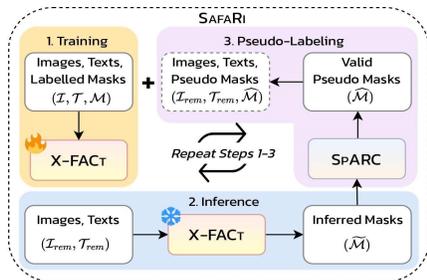


Fig. 2: Overview of weakly-supervised bootstrapping setup. It includes an initial training stage, followed by inference and pseudo-labeling steps. Filtered pseudo-masks are added to the initial dataset and model is retrained in an iterative manner.

2 Related Work

Fully-supervised RES. These methods chiefly focus on enhancing the quality of multi-modal fusion of extracted vision-language features [13, 28]. Recent efforts have been made to develop powerful multi-modal transformer-based methods [8, 21, 25, 30, 48, 58] showing significant improvements over previous baselines. **Weakly Supervised RES.** Weakly supervised RES methods have been primarily restricted to text-based weak supervision [20, 22, 29, 47]. Since they solely rely on text, they under-perform in the RES benchmarks. Also, they do not fall in our direct line of comparison. However, there exists one recent study (Partial-RES) (that is more closely related to ours) adopting partial or weak-supervision [41]. They specifically showed that employing a unified multi-modal transformer model, e.g., SeqTR [58] is more beneficial (as compared to MDETR [18]) in the presence of limited annotations. In their weakly supervised RES setting, Partial-RES [41] considered abundant bounding box annotations albeit limited mask annotations. They utilized a fully-supervised REC pretrained model to train on limited mask-annotations (Table 1). Furthermore, for filtering pseudo-masks, they again used the fully available (100%) bounding box annotations. However, since this involved 100% box annotations, our proposed setup is much more challenging and realistically weakly supervised since we consider *both box and mask annotations to be limited*. In their pseudo-labeling step on the data without masks (but with boxes), the model has already seen the boxes in the REC pretraining stage and thus the grounding information is already present.

3 SAFARI

We present SAFARI, an adaptive multi-modal sequence transformer with our novel components: (i) Cross-modal (**X-**) Fusion with **A**ttention **C**onsistency (**X-**

| Method | 100% GT Labels | | Fully-supervised | Filtering with |
|------------------|----------------|-------|------------------|----------------|
| | Mask? | BBox? | REC Pretraining? | 100% GT BBox? |
| Partial-RES [41] | ✗ | ✓ | ✓ | ✓ |
| SAFARI (Ours) | ✗ | ✗ | ✗ | ✗ |

Table 1: Comparison of SAFARI with Partial-RES [41]. Unlike Partial-RES, SAFARI *does not use 100% bounding boxes* and considers *mask % = box %* and it does not involve fully-supervised REC pretraining. Furthermore, as opposed to Partial-RES, for the filtering step, SAFARI involves a zero-shot approach to generate boxes instead of using 100% Ground Truth (GT) boxes. All ✗ evinces more challenging and realistic WSRES problem setting.

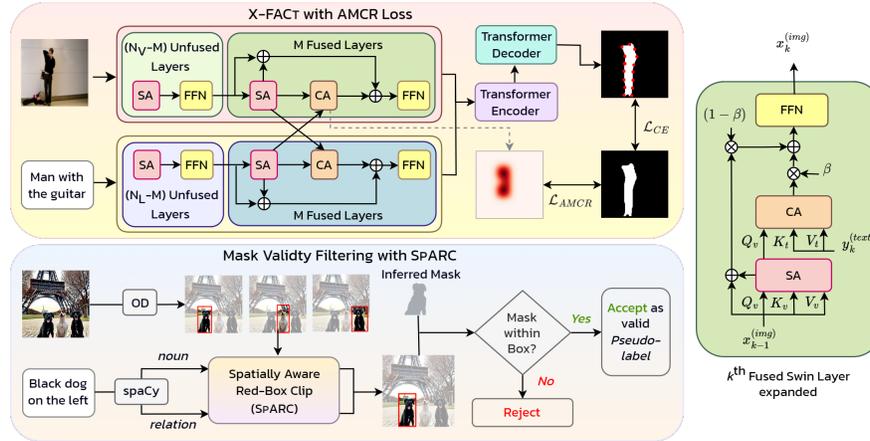


Fig. 3: Architectural components of SAFARI. (i) We introduce X-FACT, composed of normalized gated cross-attention based Fused Feature Extractors and Attention Consistency Mask Regularization (AMCR) for enhancing cross-modal synergy and spatial localization of target objects. The fused output is subsequently fed to Sequence Transformer for prediction of contour points. (ii) We design Mask Validity Filtering (MVF) strategy for choosing valid pseudo-masks using SPARC module which is a Zero-Shot REC approach with spatial reasoning capabilities.

FACT) module, (ii) bootstrapped Weak-Supervision with γ -Scheduling (WSGS), and (iii) Mask Validity Filtering (MVF) with SPARC.

3.1 X-FACT.

X-FACT consists of *Fused Feature Extractors* (with cross-modal fusion) and *Attention Mask Consistency Regularization* (AMCR) as discussed below:

3.1.1 Fused Feature Extractors. We use Swin transformer [33] and RoBERTa [32] as our image and text feature extractors. A simple concatenation-based encoding strategy fails to capture cross-modal interactions, leading to poor fine-grained multi-modal representations. To address this, previous Vision-Language studies employed cross-modal fusion modules [9, 23, 24] in their respective frameworks. However, such strategies introduce extra trainable fusion-specific layers. For instance, the early fusion scheme adopted in GLIP [24], used 6 vision-language fusion layers and 6 additional BERT layers for feature enhancement. However, such a scheme is not suitable in the presence of limited annotations. Therefore, as shown in Figure 3, we introduce a simple and lightweight fusion mechanism through normalized gated cross-attention into the layers of uni-modal feature extractors to learn high-quality language aware visual representations:

$$x_k = x_k + (1 - \beta) \cdot \text{S-MHA}(x_{k-1}) + \beta \cdot \text{C-MHA}(\text{S-MHA}(x_{k-1}), y_k) \quad (1)$$

$$x_k = x_k + \text{FFN}(x_k)$$

where x_{k-1} is the output from the $(k-1)^{th}$ layer, S-MHA and C-MHA represent Self- and Cross-Multi-Head Attention, FFN denotes cross-feed-forward network, and β is a learnable weighted gating parameter initialized from 0. The fused feature then goes into a sequence transformer for decoding the object contour points as outputs. It is worth noting that unlike previous works [9, 23] our gated methodology preserves the uni-modal embeddings (i.e., self-attention features), ensuring to learn the weighted representations during training. It also ensures to map linguistic semantic features to the localized parts of the images, making the system capable of understanding fine-grained visual features of the objects.

3.1.2 Attention Mask Consistency Regularization. Feature fusion through cross-attention strategy is highly effective in attending those parts of the image which the texts are referring to. However, in Figure 4, we observe that the attended regions are scattered across the objects including some background pixels (more pronounced in WSRES task). The lack of fidelity in being able to attend fine-grained information may be attributed to the lack of training data (in WSRES). It has also been found in [43] that vision transformers tend to demonstrate more uniform representations across all layers, enabled by early aggregation of global information and propagation of features from lower to higher layers. Moreover, in the *absence* of abundant masks (and boxes), the localization capability of the model needs to be further enhanced, especially, to yield superior pseudo-masks for the bootstrapping stage. Therefore, to address such issues, we formulate Attention Mask Consistency Regularization (AMCR) Loss to localize fine-grained cross-attention within the target object:

$$\mathcal{L}_{AMCR} = \sum_{b=1}^N \underbrace{\left(1 - \frac{\sum_{i,j} \mathcal{A}_{i,j} \mathcal{M}_{i,j}}{\sum_{i,j} \mathcal{A}_{i,j}} \right)}_{\text{Localization}} + \underbrace{\psi \text{KL}(\mathbf{U}_N \| \mathbf{Q}_N)}_{\text{Collapse-Reduction}} \quad (2)$$

$$\mathbf{Q}_N = \left\{ \left(\frac{n(\mathcal{A})}{\sum_{i,j} \mathcal{M}_{i,j}} \right)_b \mid \forall b \in N \right\}$$

where $\mathcal{A}_{i,j}$ and $\mathcal{M}_{i,j}$ represent the Cross-Attention map from the last layer and Mask, respectively, (i, j) represents pixel location, $n(\mathcal{A})$ is the number of non-zero elements in \mathcal{A} , KL denotes KL divergence loss, $\mathbf{U}_N(0, 1)$ denotes uniform distribution of minimum value 0 and maximum value 1, \mathbf{Q}_N represents computed normalized frequency distribution over batch with a size N . ψ is the loss-balancing term which we empirically set to 0.001. \mathcal{L}_{AMCR} consists of two terms: a Localization term and a Collapse-Reduction term. The localization term guarantees accurate localization and alignment of the generated cross-attention map within the mask of the object. The collapse-reduction term prevents collapsing of the cross-attention map within the mask. \mathcal{L}_{AMCR} gets added to the Cross-Entropy (CE) loss and the total loss becomes:

$$\mathcal{L}_{total} = \mathcal{L}_{CE} + \lambda \cdot \mathcal{L}_{AMCR} \quad (3)$$

where λ is responsible for weighting the \mathcal{L}_{AMCR} loss term. We provide more details regarding the formulation of \mathcal{L}_{AMCR} in the Supplementary.



Fig. 4: Qualitative differences between cross-attention maps and predicted masks in the presence and absence of AMCR. Without AMCR, some regions outside the object boundary are attended which affects the quality of predicted masks.

3.2 Weak-Supervision with γ -Scheduling.

It is worth reiterating that unlike Partial-RES [41], we *do not have box annotations for the full dataset*. Therefore, we do not *pretrain* our model on any box-prediction tasks such as Referring Expression Comprehension (REC) unlike Partial-RES. Our goal is to best adapt our architecture to the target RES task given limited mask annotations. Particularly, in a dataset consisting of \mathcal{N} training image-referring expression pairs, we consider mask (and box) annotations for $x\%$ data. These $x\% = \{10\%, 20\%, 30\%\}$ samples are selected randomly from the original training set. Our overall pipeline (Figure 2) is as follows:

- **Step 1: Initial RES Training.** In this initial step, we train SAFARI using Equation 3 with $x\%$ labeled mask data on the RES task. Subsequently, we obtain a trained *pseudo-labeler* with updated model parameters.
- **Step 2: Pseudo-labeling.** We infer masks by running inference on the remaining unlabeled $(100 - x)\%$ training samples using model trained from Step 1. Inferred masks are subsequently passed through the proposed Mask Validity Filtering (MVF) (Section 3.3) to verify the validity of these generated masks in a zero-shot fashion. Contour points are next sampled from the valid masks and added to the corresponding image-text pairs as *pseudo-masks*.
- **Step 3: Retraining with γ -Scheduling.** We retrain SAFARI (initialized from previous training) with the updated training dataset containing both the $x\%$ Ground Truth (GT) Masks (M) and Pseudo-Masks (\hat{M}) and minimize the final loss $\mathcal{L}_{\text{SAFARI}}$ using a Pseudo-Mask Loss weighting hyperparameter γ :

$$\mathcal{L}_{\text{SAFARI}} = \underbrace{\mathcal{L}_{\text{total}}^{\text{GT}}}_{\text{Loss with GT Masks}} + \gamma \cdot \underbrace{\mathcal{L}_{\text{total}}^{\text{Pseudo}}}_{\text{Loss with Pseudo-Masks}} \quad (4)$$

The role of γ is to balance the amount of GT masks against that of pseudo-masks during the retraining step. We introduce a γ -scheduling strategy to systematically change the value of γ over a set of iterations with the starting (minimum) value of γ being $\gamma_0 = 0.9$ (found empirically through ablation studies) and maximum value being 1.0. The intuition is to provide more weighting to pseudo-masks as the model becomes more confident of the predictions as the training steps progress. We outline the detailed steps in Algorithm 1 in Supplementary.

3.3 Mask Validity Filtering with SPARC.

Mask Validity Filtering (MVF) is the process of identifying which predicted (inferred) masks actually localize the object (present in the image) being referred to in the text (referring expression) [41]. Partial-RES [41] achieves this task by using ground-truth bounding-box annotations. However, as noted in Section 1, unlike Partial-RES, we consider a more realistic *weakly-supervised* scenario where the number of bounding box and mask annotations are equal. Therefore, to address this gap in the literature, we propose a zero-shot approach to obtain these boxes to be used in the MVF stage. This can also be formulated as a Zero-Shot Referring Expression Comprehension (ZS-REC) task. Our proposed MVF approach is composed of two parts - (i) obtaining bounding boxes using ZS-REC, and (ii) validating the inferred masks using the obtained bounding boxes.

3.3.1 ZS-REC with SPARC module. We introduce **S**patially **A**ware **R**ed-**B**ox **C**lip (**SPARC**), a ZS-REC module, comprising two components: a proposal scoring approach with visual prompts (red-box), and a spatial reasoning component accounting for spatial understanding of proposals in a rule-based manner.

Proposal scoring with red-box prompting: Our ZS framework considers an input image and the referring expression along with \mathcal{N}_B box proposals as generated by a pretrained object detector [44]. ZS-REC task can be conveniently converted to a ZS-retrieval task using an ensemble of image (visual) prompts with a single text as described below. Each of these visual prompts (images) must correspond to a single isolated proposal from the original image generating \mathcal{N}_B such images (visual prompts). Each proposal object region is isolated by employing a Gaussian blur on the background and by adding a red box (red border) of uniform thickness surrounding the object proposal. These image prompts along with the referring expression is passed through CLIP to obtain the CLIP score. Notably, using a red border acts as a *positive visual prompt* highlighting the target area, whereas blurring the background reduces the impact of weakly related information. However, these CLIP scores do not consider the spatial understanding of the objects in the different bounding boxes. Therefore, injecting spatial understanding is important which we discuss next.

Spatial Reasoning component: A major limitation of CLIP for the ZS-REC task is the lack of fine-grained spatial understanding. In Figure 3, the text "*black dog*" refers to the two black dogs in the picture. CLIP will generate high scores for both these proposal regions since both fall under the class "*black dog*". Furthermore, although addition of visual prompt improves detection performance [46], it neither incorporates nor enhances spatial understanding capabilities of the model. To mitigate this, we introduce a simple rule-based approach. We divide the referring expression into two constituents - noun chunks and their relations. In the previous example, it refers to the "*black dog*" in the picture which is the subject. The phrases connecting these noun chunks form the relationships (e.g., *right/east, left/west, smaller/tinier/further, bigger/larger/closer, between, within/inside, above/north/top, below/under/south, back/behind, and front*). In the presence of two boxes referring to the *black dog*, the probability of the proposal on the left will be the most (as decided based on the location of the

centroid of the box computed using the obtained box coordinates). To provide more clarity on this, we consider the example: "*black dog on the left of a white dog*", relation R between two nouns black dog (X) and white dog (Y) is "*left*". Coordinates of proposals X and Y , and relation R are the inputs of Spatial Reasoning component. If center point of box X is to the "*left*" of that of box Y , the output probability $\Pr[R(X, Y)] = 1$, otherwise 0. These probabilities are multiplied with the CLIP scores of the respective proposals. Note that in rare occasions where the referring texts in the RefCOCO datasets do not have any spatial relations, we resort to using only CLIP scores for scoring the proposals. More examples, with those of complex relations are given in Supplementary.

3.3.2 Validation of Inferred Masks with SPARC. In the Pseudo-labeling step, SAFARI predicts contour points which are connected to generate binary masks. We generate a bounding box from the outermost (top-most, bottom-most, right-most, left-most) points of each mask and compute Dice Similarity Coefficient (DSC) [39] between the generated box and the box obtained using the ZS-REC step using SPARC. We reject the noisy pseudo-masks with DSC value less than $\tau = 0.1$ (ablation in Supplementary). Contour points are resampled from the filtered pseudo-masks and added to the training set (Figure 2).

4 Experiments and Results

4.1 Datasets, Implementation Details and Metrics

4.1.1 Dataset. For *RES*, we conduct our experiments on the three major RES datasets: RefCOCO [55], RefCOCO+ [55], and RefCOCOg [37, 40]. For *Zero-shot Transfer to Referring Video Object Segmentation (ZS-R-VOS)*, we conduct zero-shot (ZS) evaluations on two popular R-VOS datasets: RefDAVIS17 [19] and JHMDB-Sentences [11].

4.1.2 Implementation. We use AdamW optimizer [34] with a batch size of 128. We use an initial learning rate (LR) of $5e-4$ with $(\beta_1, \beta_2) = (0.9, 0.98)$, $\epsilon = 10^{-9}$, and Multi-Step LR Warmup of 5 epochs, decay steps of 75, and a decay ratio of 0.1. We use A100 40GB GPUs for all experiments.

4.1.3 Metrics. For *RES*, we report mean Intersection-over-Union (mIoU) values. For *ZS-R-VOS*, we report mean of region similarity (\mathcal{J}) and contour accuracy (\mathcal{F}) denoted as $\mathcal{J}\&\mathcal{F}$ on RefDAVIS17, and mIoU scores on JHMDB-Sentences. Additional dataset and implementation details are provided in Supplementary.

4.2 Main Results

4.2.1 Referring Expression Segmentation (RES). We compare SAFARI with the state-of-the-art methods on the RefCOCO/+g validation and test sets and report the mean IoU (mIoU) in the Table 2 under different label-rates.

Full Supervision: SAFARI outperforms existing benchmark methods on each split of the three datasets by a significant margin (Table 2). For a fair comparison with PolyFormer [30], we additionally train SAFARI on a combined dataset

| Label-Rate Mask BBox | Method | Vis. Backbone | RefCOCO | | | RefCOCO+ | | | RefCOCog | | | |
|--|---|-------------------------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|-------|
| | | | val | testA | testB | val | testA | testB | val-g | val-u | test-u | |
| <i>Fully Supervised Models</i> | | | | | | | | | | | | |
| 100% | LTS [17] | DN53 | 65.43 | 67.76 | 63.08 | 54.21 | 58.32 | 48.02 | - | 54.40 | 54.25 | |
| | VLT [7] | DN56 | 65.65 | 68.29 | 62.73 | 55.50 | 59.20 | 49.36 | 49.76 | 52.99 | 56.65 | |
| | ResTR [21] | ViT-B | 67.22 | 69.30 | 64.45 | 55.78 | 60.44 | 48.27 | 54.48 | - | - | |
| | SeqTR [58] | DN53 | 67.26 | 69.79 | 64.12 | 54.14 | 58.93 | 48.19 | - | 55.67 | 55.64 | |
| | SAFARI (Ours) | Swin-B | 73.35 | 75.02 | 70.71 | 63.03 | 65.81 | 57.64 | - | 62.42 | 62.74 | |
| | $\Delta_{\text{Ours}} - \text{SeqTR}$ | - | 6.09 \uparrow | 5.23 \uparrow | 6.59 \uparrow | 8.89 \uparrow | 6.88 \uparrow | 9.45 \uparrow | - | 6.75 \uparrow | 7.10 \uparrow | |
| | PolyFormer [53] [†] | Swin-B | 75.96 | 77.09 | 73.22 | 70.65 | 74.51 | 64.64 | - | 69.36 | 69.88 | |
| | SAFARI[†] (Ours) | Swin-B | 77.21 | 77.83 | 75.72 | 70.78 | 74.53 | 64.88 | - | 70.48 | 71.06 | |
| $\Delta_{\text{Ours}} - \text{PolyFormer}^{\dagger}$ | - | 1.25 \uparrow | 0.74 \uparrow | 2.50 \uparrow | 0.13 \uparrow | 0.02 \uparrow | 0.24 \uparrow | - | 1.12 \uparrow | 1.18 \uparrow | | |
| <i>Weakly Supervised Models</i> | | | | | | | | | | | | |
| 30% | 100% 30% | Partial-RES [41] [♣] | Swin-B | 66.24 | 68.39 | 63.57 | 54.37 | 58.16 | 47.92 | - | 54.69 | 54.81 |
| | SAFARI (Ours) | Swin-B | 67.04 | 69.17 | 64.23 | 54.98 | 59.31 | 48.26 | - | 55.72 | 55.83 | |
| | $\Delta_{\text{Ours}} - \text{Partial-RES}$ | - | 0.80 \uparrow | 0.78 \uparrow | 0.66 \uparrow | 0.61 \uparrow | 1.15 \uparrow | 0.34 \uparrow | - | 1.03 \uparrow | 1.02 \uparrow | |
| 20% | 100% 20% | Partial-RES [41] [♣] | Swin-B | 65.20 | 67.43 | 62.85 | 53.78 | 57.52 | 47.39 | - | 53.94 | 54.02 |
| | SAFARI (Ours) | Swin-B | 65.88 | 67.96 | 63.24 | 54.23 | 58.07 | 47.67 | - | 54.45 | 54.61 | |
| | $\Delta_{\text{Ours}} - \text{Partial-RES}$ | - | 0.68 \uparrow | 0.53 \uparrow | 0.39 \uparrow | 0.45 \uparrow | 0.55 \uparrow | 0.28 \uparrow | - | 0.51 \uparrow | 0.59 \uparrow | |
| 10% | 100% 10% | Partial-RES [41] [♣] | Swin-B | 64.01 | 65.89 | 61.68 | 52.85 | 56.01 | 46.27 | - | 52.73 | 52.68 |
| | SAFARI (Ours) | Swin-B | 64.02 | 65.91 | 61.76 | 52.98 | 56.24 | 46.48 | - | 52.91 | 52.94 | |
| | $\Delta_{\text{Ours}} - \text{Partial-RES}$ | - | 0.01 \uparrow | 0.02 \uparrow | 0.08 \uparrow | 0.13 \uparrow | 0.23 \uparrow | 0.21 \uparrow | - | 0.18 \uparrow | 0.26 \uparrow | |

Table 2: Comparison with the state-of-the-arts on the RES task. SAFARI substantially outperforms SOTA SeqTR [58] in the fully-supervised benchmark. SAFARI also yields significant gains over baseline Partial-RES [41] even without using 100% box annotations in the WSRES task. \uparrow means trained on extra data combining RefCOCO datasets [30]. \clubsuit indicates our reimplementations of Partial-RES with Swin-B backbone where we get better mIoUs than their reported values [41]. Differences shown in blue.

after removing all data from validation and test sets [30]. Unlike PolyFormer, SAFARI is not pretrained on the REC task. However, SAFARI still outperforms PolyFormer on all three datasets. Note that we report fully-supervised results to show an upper-bound on the performance of the weakly-supervised models.

Weak Supervision: We consider label rates of 10%, 20%, and 30% for the WSRES task (Table 2). Despite methods such as Partial-RES using 100% box annotations, SAFARI shows substantial improvements over Partial-RES (note that Partial-RES originally used DN-53 backbone and we re-implemented with Swin-B backbone for fair comparisons) for all three label-rates. With just 30% labeled data and without any REC pretraining, SAFARI often surpasses the fully-supervised performance of SeqTR on RefCOCO/+g validation sets (e.g., 59.31 vs 58.93 on RefCOCO+@testA, 55.83 vs 55.64 on RefCOCog@test sets). Weakly-supervised SAFARI significantly outperforms other fully-supervised baselines (e.g., VLT [7] and LTS [17]). These results bespeak the role of our proposed X-FACT (containing cross-modal fusion together with AMCR) and SPARC modules in obtaining state-of-the-art results in the limited annotation scenarios.

4.2.2 Zero-shot Referring Video Object Segmentation (ZS-R-VOS).

We conduct evaluation on the R-VOS datasets (RefDAVIS17 and JHMDB) as a ZS transfer task. We consider the video frames as a sequence of images without involving any temporal information while predicting masks. As shown in Ta-

| Method | Vis. Backbone Eval. | | | RefDAVIS17 | JHMDB | Label-Rate | Cross-Attn. | AMCR | RefCOCO | | |
|--|---------------------|----|----------------------------------|------------|----------------------------------|------------|--------------|--------------|--------------|--------------|--------------|
| | | | | $J&F$ | mIoU | | | | val | testA | testB |
| <i>Fully Supervised Models (Label-Rate = 100%)</i> | | | | | | | | | | | |
| ReferFormer [49] | Swin-L | FT | 60.5 | - | - | 100 | \times | \times | 68.84 | 70.59 | 65.85 |
| MTTR [2] | Vid-Swin-T | FT | - | - | 36.6 | | \checkmark | \times | 72.41 | 74.65 | 69.96 |
| ReferFormer [49] | Vid-Swin-B | FT | <u>61.1</u> | - | <u>43.7</u> | | \checkmark | \checkmark | 73.35 | 75.02 | 70.71 |
| PolyFormer [30] | Swin-B | ZS | 60.9 | - | 42.4 | 30 | \times | \times | 60.91 | 64.06 | 57.73 |
| SeqTR [58] | DN-53 | ZS | 53.5 | - | 34.9 | | \checkmark | \times | 64.72 | 67.55 | 60.91 |
| SAFARI-100 (Ours) | Swin-B | ZS | 61.3 | - | 43.2 | | \checkmark | \checkmark | 67.04 | 69.17 | 64.23 |
| $\Delta_{\text{Ours}} - \text{SeqTR}$ | - | ZS | 7.8 \uparrow | - | 8.3 \uparrow | | \checkmark | \checkmark | | | |
| <i>Weakly Supervised Models</i> | | | | | | | | | | | |
| Partial-RES-30 [41] | Swin-B | ZS | 52.3 | - | 34.8 | 20 | \times | \times | 58.84 | 61.79 | 57.02 |
| SAFARI-30 (Ours) | Swin-B | ZS | 55.3 | - | 38.1 | | \checkmark | \times | 61.25 | 64.68 | 59.47 |
| $\Delta_{\text{Ours}} - \text{Partial-RES-30}$ | - | ZS | 3.0 \uparrow | - | 3.3 \uparrow | | \checkmark | \checkmark | 65.88 | 67.96 | 63.24 |
| Partial-RES-10 [41] | Swin-B | ZS | 51.5 | - | 34.3 | 10 | \times | \times | 54.06 | 55.64 | 52.61 |
| SAFARI-10 (Ours) | Swin-B | ZS | 53.1 | - | 36.4 | | \checkmark | \times | 60.11 | 62.18 | 57.95 |
| $\Delta_{\text{Ours}} - \text{Partial-RES-10}$ | - | ZS | 1.6 \uparrow | - | 2.1 \uparrow | | \checkmark | \checkmark | 64.02 | 65.91 | 61.76 |

Table 3: Evaluation performance of SAFARI on RefDAVIS17 and JHMDB validation datasets. SAFARI achieves significant gains in zero-shot settings for both datasets. FT and ZS indicate Fine-Tuned and Zero-Shot, respectively. Differences in blue.

Table 4: Impact of Cross-Attention and AMCR (in the X-FACT module) on SAFARI’s performance evaluated on RefCOCO val, testA, and testB sets for 100%, 30%, 20%, and 10% label-rates. Best results are highlighted.

ble 3, SAFARI achieves state-of-the-art results with just the spatial information, displaying strong generalization capabilities.

We also conduct evaluations on the *ZS-REC* task with our proposed SPARC module which significantly outperforms SOTA RedCircle [46] achieving 39.3% and 16.2% gains on RefCOCO@testB and RefCOCO+@testB sets respectively (see Supplementary). More comparisons are provided in Supplementary.

4.3 Ablation Study

4.3.1 Impact of X-FACT. We show the impact of X-FACT containing Cross-Attention (CA) based fusion and the AMCR components in Table 4. Inclusion of CA is shown to improve mIoU values consistently across varying label-rates. The impact of AMCR is more pronounced in the cases of limited annotations. For instance, with a label-rate of 100%, addition of AMCR boosts mIoU by 0.94 points on RefCOCO@val, whereas with a label-rate of 10%, this difference increases substantially to 3.91. Additionally, in Figure 4 we demonstrate that incorporating AMCR qualitatively improves both the cross-attention maps and the predicted masks, highlighting the efficacy of AMCR in our pipeline. We further assess the importance of AMCR loss balancing factor (λ) in the Figure 5b when evaluated on RefCOCO@val. Increasing λ improves mIoU initially and the maximum is achieved at 0.4, beyond which the performance drops significantly. Moreover, in Table 7 we show that despite being lightweight, our proposed fusion scheme in X-FACT is superior in performance as compared to (conventional) ALBEF-like fusion which contains extra fusion-specific layers.

4.3.2 Retraining with MVF using SPARC. In the presence of limited mask (and box) annotations, retraining with γ scheduling leads to a marked improve-

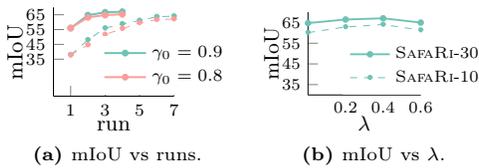


Fig. 6: Impact of γ -scheduling under different initial values (γ_0) and AMCR balancing factor (λ) when evaluated on RefCOCO@val at 30% and 10% mask labels.

| Red-box | Spatial | Blur | Crop | Ref@val | Ref+@val | Refg@val |
|---------|---------|------|------|--------------|--------------|--------------|
| x | x | x | ✓ | 63.11 | 47.48 | 49.43 |
| x | x | ✓ | x | 63.78 | 49.66 | 51.48 |
| ✓ | x | x | x | 64.35 | 51.16 | 52.87 |
| ✓ | ✓ | x | x | 65.29 | 52.51 | 54.22 |
| x | ✓ | x | ✓ | 64.12 | 50.19 | 51.80 |
| x | ✓ | ✓ | x | 65.94 | 52.83 | 54.13 |
| ✓ | ✓ | ✓ | x | 67.04 | 54.98 | 55.72 |

Table 5: mIoUs with various prompts for SPARC at 30% label-rate. Best results are obtained when both red-box and blurring are used with the spatial reasoning component.

| Label-Rate | L2 | AMCR | Ref@val | Ref+@val | Refg@val |
|------------|----|------|--------------|--------------|--------------|
| 30 | ✓ | — | 63.48 | 51.31 | 52.16 |
| | — | ✓ | 67.04 | 54.98 | 55.72 |
| 10 | ✓ | — | 59.23 | 48.29 | 48.83 |
| | — | ✓ | 64.02 | 52.98 | 52.91 |

Table 6: Comparison of proposed AMCR with L2 loss for 10% and 30% annotations.

| Fusion Scheme | Label-Rate | RefCOCOg@val | Δ parameter increase (%) \downarrow |
|---------------|------------|--------------|--|
| X-FACT (Ours) | 30 | 55.72 | 1.82 |
| | 100 | 62.42 | |
| ALBEF-like | 30 | 54.29 | 29.65 |
| | 100 | 60.17 | |

Table 7: X-FACT vs ALBEF-like fusion for 30% and 100% annotations.

ment in the mIoU values as depicted in Figure 5a. Here, we consider label-rates of 30% (solid lines) and 10% (dashed lines) with two different γ_0 values denoting the starting values of the γ variable. We note that as the number of retraining steps (runs) increases, the mIoU values (for both γ_0 values) in each label-rate increases. However, with $\gamma_0 = 0.9$ we achieve slightly better performance. Additionally, we evaluate the effectiveness of the MVF stage with SPARC in Supplementary. Results indicate that without MVF, generated pseudo-labels cannot be validated which negatively impacts model’s performance. In contrast, when pseudo-labels are validated using MVF, the mIoU values on RefCOCO@val set are shown to improve by 4.36 mIoU, demonstrating the importance of MVF.

4.3.3 Impact of various components of SPARC. In Table 5 we assess the impact of individual components of SPARC. We note that red-box prompting and blurring together boost performance. Adding spatial awareness via the rule-based reasoning module also substantially improves mIoU on the WSRES task.

4.3.4 AMCR vs L2 loss in X-FACT. To evaluate the impact of the AMCR component in X-FACT, we analyze SAFARI’s performance after replacing the proposed formulation of AMCR with a L2 loss. As shown in Table 6, we observe that using L2 (in place of AMCR) leads to substantially poor performance, as compared to when AMCR is used in X-FACT. This corroborates the strong localization capabilities introduced in the system by the AMCR formulation which is not achieved with a L2 loss.

4.3.5 Impact of Label Rates. We notice that as the label-rate increases from 10% to 100%, the performance of the model improves significantly across all the datasets (see Table 2 and Figure 1). Remarkably, with just 30% annotated data,

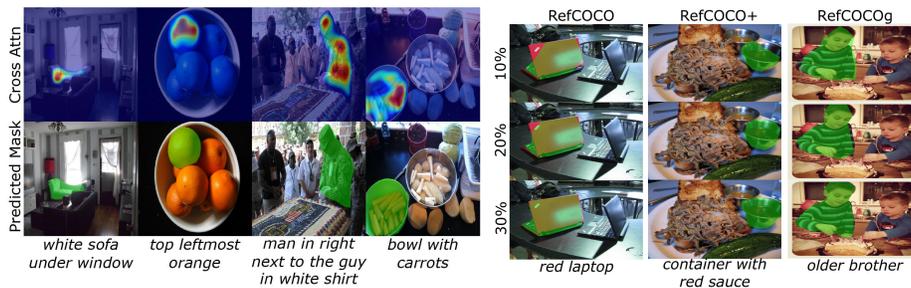


Fig. 7: Cross-attention Maps and cor- **Fig. 8: Predictions with varying**
responding predictions showing strong **label-rates**. With increasing mask an-
cross-modal alignment learned by SAFARI. notations %, prediction quality improves.

SAFARI achieves 59.31 mIoU versus 58.93 mIoU obtained by the fully-supervised SeqTR on RefCOCO+@testA set.

We provide additional ablations on the impacts of MVF, gated cross-attention parameter, and DSC thresholding parameter in Supplementary.

4.4 Qualitative Assessment and Error Analysis

4.4.1 Backbone Cross-Attentions and Predicted Masks. In Figure 7, we show how cross-modal attention maps attend to different objects in the images, guided by the referring texts. The cross-attention scores between the image and the associated expression are extracted and bilinearly interpolated to match the image dimension and superimposed on the original image. Notably, strong cross-modal interactions aid in predicting high-quality masks as displayed in Figure 7. In Figure 10, we also provide high-quality segmentation masks generated by SAFARI for the ZS-R-VOS tasks without any finetuning on the respective datasets.

4.4.2 Varying Label-Rates. In Figure 8, we display segmentation masks under varying label-rates using SAFARI. We clearly notice that the quality of masks improve substantially with an increase in the ground-truth annotations.

4.4.3 Varying Retraining Steps. In Figure 9 we show that with increasing retraining steps (runs), SAFARI becomes more confident which results in significant qualitative improvements in the predictions. For example, in Figure 9 although SAFARI could not recognize *the magazine partially behind laptop* in the first step, it was successful in predicting the mask accurately in the final step - this shows the efficacy of the retraining stage.

4.4.4 Comparisons with Partial-RES. We provide qualitative examples for 30% label-rate in Figure 1. It is evident from Figure 1 that SAFARI succeeds in challenging cases requiring extensive linguistic understanding (attained with X-FACT and SPARC modules) where Partial-RES [41] fails.

We provide extended versions of these qualitative visualizations in Supplementary. Additionally, we show examples of cases in Supplementary where SAFARI occasionally fails to focus on the referred objects. These cases typically

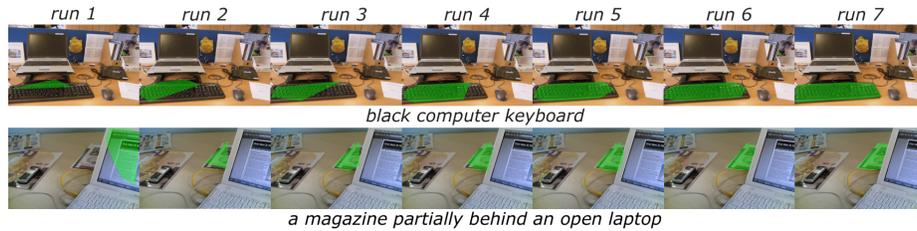


Fig. 9: Examples of masks with increasing WSRES bootstrapping runs (steps) for 10% annotations. We see significant improvements in localization capabilities with an increase in retraining steps illustrating the efficacy our approach.

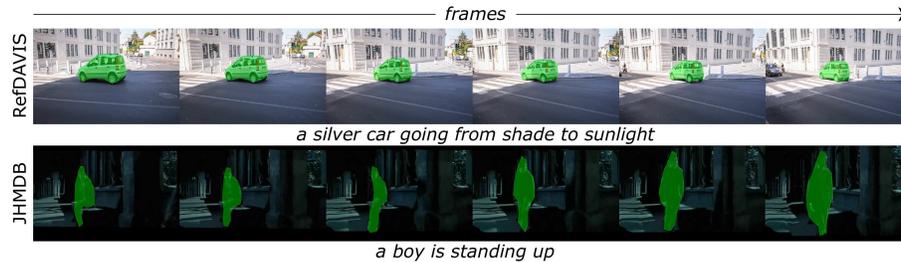


Fig. 10: Predicted masks on RefDAVIS17 and JHMDB datasets in a zero-shot setting with SAFARI-30 (trained with 30% annotations).

consists of vastly hindered objects in cluttered environments, especially in low-light blurry conditions and tiny objects.

5 Conclusion

We present a weakly-supervised learning framework for RES considering limited mask (and box) annotations and employing a contour-based sequence prediction approach. Unlike Partial-RES [41], we do not consider box annotations to be fully available and therefore we do not pretrain our model on the fully-supervised REC task. We incorporate lightweight gated cross-modal attention in the feature backbones along with an Attention Mask Consistency Regularization module to facilitate strong cross-modal alignment and improve the quality of predicted masks. We introduce a bootstrapping pipeline with self-labeling capabilities where pseudo-labels are validated using our proposed Mask Validity Filtering approach. We conduct extensive experiments to demonstrate that SAFARI consistently achieves state-of-the-art performance across all RES datasets. Finally, we demonstrate excellent generalization capabilities of SAFARI on *zero-shot* referring video-object segmentation task. Extending our approach to multi-image and video settings can be looked upon as a promising future work.

References

1. Bolya, D., Zhou, C., Xiao, F., Lee, Y.J.: YOLACT: Real-time instance segmentation. In: ICCV. pp. 9157–9166 (2019)
2. Botach, A., Zheltonozhskii, E., Baskin, C.: End-to-end referring video object segmentation with multimodal transformers. In: CVPR. pp. 4985–4995 (2022)
3. Chen, D.J., Jia, S., Lo, Y.C., Chen, H.T., Liu, T.L.: See-through-text grouping for referring image segmentation. In: ICCV. pp. 7454–7463 (2019)
4. Chen, K., Pang, J., Wang, J., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Shi, J., Ouyang, W., et al.: Hybrid task cascade for instance segmentation. In: CVPR. pp. 4974–4983 (2019)
5. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: ECCV. pp. 104–120 (2020)
6. Dai, J., He, K., Sun, J.: Instance-aware semantic segmentation via multi-task network cascades. In: CVPR. pp. 3150–3158 (2016)
7. Ding, H., Liu, C., Wang, S., Jiang, X.: Vision-language transformer and query generation for referring segmentation. In: ICCV. pp. 16321–16330 (2021)
8. Ding, H., Liu, C., Wang, S., Jiang, X.: Vision-language transformer and query generation for referring segmentation. In: ICCV. pp. 16321–16330 (2021)
9. Dou, Z.Y., Xu, Y., Gan, Z., Wang, J., Wang, S., Wang, L., Zhu, C., Zhang, P., Yuan, L., Peng, N., et al.: An empirical study of training end-to-end vision-and-language transformers. In: CVPR. pp. 18166–18176 (2022)
10. Feng, G., Hu, Z., Zhang, L., Lu, H.: Encoder fusion network with co-attention embedding for referring image segmentation. In: CVPR. pp. 15506–15515 (2021)
11. Gavriluk, K., Ghodrati, A., Li, Z., Snoek, C.G.: Actor and action video segmentation from a sentence. In: CVPR. pp. 5958–5966 (2018)
12. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV. pp. 2961–2969 (2017)
13. Hu, R., Rohrbach, M., Darrell, T.: Segmentation from natural language expressions. In: ECCV. pp. 108–124 (2016)
14. Hu, Z., Feng, G., Sun, J., Zhang, L., Lu, H.: Bi-directional relationship inferring network for referring image segmentation. In: CVPR. pp. 4424–4433 (2020)
15. Huang, S., Hui, T., Liu, S., Li, G., Wei, Y., Han, J., Liu, L., Li, B.: Referring image segmentation via cross-modal progressive comprehension. In: CVPR. pp. 10488–10497 (2020)
16. Hui, T., Liu, S., Huang, S., Li, G., Yu, S., Zhang, F., Han, J.: Linguistic structure guided context modeling for referring image segmentation. In: ECCV. pp. 59–75 (2020)
17. Jing, Y., Kong, T., Wang, W., Wang, L., Li, L., Tan, T.: Locate then segment: A strong pipeline for referring image segmentation. In: CVPR. pp. 9858–9867 (2021)
18. Kamath, A., Singh, M., LeCun, Y., Synnaeve, G., Misra, I., Carion, N.: Mdetrm: modulated detection for end-to-end multi-modal understanding. In: ICCV. pp. 1780–1790 (2021)
19. Khoreva, A., Rohrbach, A., Schiele, B.: Video object segmentation with language referring expressions. In: ACCV. pp. 123–141. Springer (2019)
20. Kim, D., Kim, N., Lan, C., Kwak, S.: Shatter and gather: Learning referring image segmentation with text supervision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15547–15557 (2023)

21. Kim, N., Kim, D., Lan, C., Zeng, W., Kwak, S.: Restr: Convolution-free referring image segmentation using transformers. In: CVPR. pp. 18145–18154 (2022)
22. Lee, J., Lee, S., Nam, J., Yu, S., Do, J., Taghavi, T.: Weakly supervised referring image segmentation with intra-chunk and inter-chunk consistency. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21870–21881 (2023)
23. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. In: NeurIPS (2021)
24. Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., et al.: Grounded language-image pre-training. In: CVPR. pp. 10965–10975 (2022)
25. Li, M., Sigal, L.: Referring transformer: A one-step approach to multi-task visual grounding. In: NeurIPS (2021)
26. Li, R., Li, K., Kuo, Y.C., Shu, M., Qi, X., Shen, X., Jia, J.: Referring image segmentation via recurrent refinement networks. In: CVPR. pp. 5745–5753 (2018)
27. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., et al.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: ECCV. pp. 121–137. Springer (2020)
28. Liu, C., Lin, Z., Shen, X., Yang, J., Lu, X., Yuille, A.: Recurrent multimodal interaction for referring image segmentation. In: ICCV. pp. 1271–1280 (2017)
29. Liu, F., Liu, Y., Kong, Y., Xu, K., Zhang, L., Yin, B., Hancke, G., Lau, R.: Referring image segmentation using text supervision. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22124–22134 (2023)
30. Liu, J., Ding, H., Cai, Z., Zhang, Y., Satzoda, R.K., Mahadevan, V., Manmatha, R.: Polyformer: Referring image segmentation as sequential polygon generation. In: CVPR. pp. 18653–18663 (2023)
31. Liu, S., Qi, L., Qin, H., Shi, J., Jia, J.: Path aggregation network for instance segmentation. In: CVPR. pp. 8759–8768 (2018)
32. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
33. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV. pp. 10012–10022 (2021)
34. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
35. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilmert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In: NeurIPS (2019)
36. Lu, J., Goswami, V., Rohrbach, M., Parikh, D., Lee, S.: 12-in-1: Multi-task vision and language representation learning. In: CVPR. pp. 10437–10446 (2020)
37. Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A.L., Murphy, K.: Generation and comprehension of unambiguous object descriptions. In: CVPR. pp. 11–20 (2016)
38. Margffoy-Tuay, E., Pérez, J.C., Botero, E., Arbeláez, P.: Dynamic multimodal instance segmentation guided by natural language queries. In: ECCV. pp. 630–645 (2018)
39. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 3DV. pp. 565–571. IEEE (2016)
40. Nagaraja, V.K., Morariu, V.I., Davis, L.S.: Modeling context between objects for referring expression understanding. In: ECCV. pp. 792–807 (2016)

41. Qu, M., Wu, Y., Wei, Y., Liu, W., Liang, X., Zhao, Y.: Learning to segment every referring object point by point. In: CVPR. pp. 3021–3030 (2023)
42. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763 (2021), <https://proceedings.mlr.press/v139/radford21a.html>
43. Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., Dosovitskiy, A.: Do vision transformers see like convolutional neural networks? *NeurIPS* **34**, 12116–12128 (2021)
44. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS* **28** (2015)
45. Shi, H., Li, H., Meng, F., Wu, Q.: Key-word-aware network for referring expression image segmentation. In: ECCV. pp. 38–54 (2018)
46. Shtedritski, A., Rupprecht, C., Vedaldi, A.: What does clip know about a red circle? visual prompt engineering for vlms. arXiv preprint arXiv:2304.06712 (2023)
47. Strudel, R., Laptev, I., Schmid, C.: Weakly-supervised segmentation of referring expressions. arXiv preprint arXiv:2205.04725 (2022)
48. Wang, Z., Lu, Y., Li, Q., Tao, X., Guo, Y., Gong, M., Liu, T.: CRIS: Clip-driven referring image segmentation. In: CVPR. pp. 11686–11695 (2022)
49. Wu, J., Jiang, Y., Sun, P., Yuan, Z., Luo, P.: Language as queries for referring video object segmentation. In: CVPR. pp. 4974–4984 (2022)
50. Yang, S., Xia, M., Li, G., Zhou, H.Y., Yu, Y.: Bottom-up shift and reasoning for referring image segmentation. In: CVPR. pp. 11266–11275 (2021)
51. Yang, Z., Wang, J., Tang, Y., Chen, K., Zhao, H., Torr, P.H.: LAVT: Language-Aware Vision Transformer for Referring Image Segmentation. In: CVPR. pp. 18155–18165 (2022)
52. Yang, Z., Gan, Z., Wang, J., Hu, X., Ahmed, F., Liu, Z., Lu, Y., Wang, L.: Unitab: Unifying text and box outputs for grounded vision-language modeling. In: ECCV. pp. 521–539 (2022)
53. Ye, L., Rochan, M., Liu, Z., Wang, Y.: Cross-modal self-attention network for referring image segmentation. In: CVPR. pp. 10502–10511 (2019)
54. Ye, L., Rochan, M., Liu, Z., Wang, Y.: Cross-modal self-attention network for referring image segmentation. In: CVPR. pp. 10502–10511 (2019)
55. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: ECCV. pp. 69–85 (2016)
56. Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding. *NeurIPS* **35**, 36067–36080 (2022)
57. Zhang, P., Li, X., Hu, X., Yang, J., Zhang, L., Wang, L., Choi, Y., Gao, J.: Vinvl: Revisiting visual representations in vision-language models. In: CVPR. pp. 5579–5588 (2021)
58. Zhu, C., Zhou, Y., Shen, Y., Luo, G., Pan, X., Lin, M., Chen, C., Cao, L., Sun, X., Ji, R.: Seqtr: A simple yet universal network for visual grounding. In: ECCV. pp. 598–615 (2022)