Pseudo-RIS: Distinctive Pseudo-supervision Generation for Referring Image Segmentation

Seonghoon Yu¹ Paul Hongsuck Seo^{2†} Jeany Son^{1†}

¹AI Graduate School, GIST, Korea

² Department of Computer Science and Engineering, Korea University, Korea seonghoon@gm.gist.ac.kr phseo@korea.ac.kr jeany@gist.ac.kr https://github.com/Seonghoon-Yu/Pseudo-RIS

Abstract. We propose a new framework that automatically generates high-quality segmentation masks with their referring expressions as pseudo supervisions for referring image segmentation (RIS). These pseudo supervisions allow the training of any supervised RIS methods without the cost of manual labeling. To achieve this, we incorporate existing segmentation and image captioning foundation models, leveraging their broad generalization capabilities. However, the naïve incorporation of these models may generate non-distinctive expressions that do not distinctively refer to the target masks. To address this challenge, we propose two-fold strategies that generate distinctive captions: 1) 'distinctive caption sampling', a new decoding method for the captioning model, to generate multiple expression candidates with detailed words focusing on the target. 2) 'distinctiveness-based text filtering' to further validate the candidates and filter out those with a low level of distinctiveness. These two strategies ensure that the generated text supervisions can distinguish the target from other objects, making them appropriate for the RIS annotations. Our method significantly outperforms both weakly and zero-shot SoTA methods on the RIS benchmark datasets. It also surpasses fully supervised methods in unseen domains, proving its capability to tackle the open-world challenge within RIS. Furthermore, integrating our method with human annotations yields further improvements, highlighting its potential in semi-supervised learning applications.

Keywords: Referring image segmentation · Pseudo-supervision

1 Introduction

Referring image segmentation (RIS) task aims to segment the object referred by a given natural language description and has substantial applications across various fields such as human-object interaction [55], image editing [6]. While RIS is particularly challenging because it requires the model to map and understand both visual data and textual descriptions simultaneously, recent fully supervised approaches [9,28,39,78,81,85] have shown remarkable advances. However, these

[†]Corresponding authors.



Fig. 1: Illustration of our distinctive pseudo-supervision generation: (a) a distinctive caption "a brown cow with a long tail" distinctively refers to a target mask, while a non-distinctive caption "a cow with a tail" causes misleading to a non-targeted object since there is another cow with a tail in an image, (b) two proposed methods for distinctive caption generation given segmentation masks: 1) multiple caption candidates generation with target-specific words; 2) filtering out the misleading captions among all candidates.

methods heavily rely on costly human annotated datasets which include pixelwise instance masks with their corresponding referring descriptions. Moreover, the models trained on heavily supervised datasets show limited generalization performances [91]. This issue is particularly amplified by the current mainstream RIS datasets that mainly focus only on a few object classes (*e.g.*, 80 classes in the RefCOCO variants).

A couple of recent works have proposed weakly supervised [27, 34, 40, 69] and zero-shot RIS approaches [53, 91] to handle costly human annotation. On one hand, weakly supervised methods attempt to tackle RIS using the image-text paired data where the spatial groundings of language descriptions are missing. On the other hand, zero-shot RIS approaches utilize large-scale vision-language pre-trained models [58,60] and use their discriminative power to localize language expressions within an input image. However, both the weakly supervised RIS models and the pre-trained models for the zero-shot methods are trained on the image-text paired datasets which lack grounding annotations, and furthermore the pre-trained models use scene-level captions that are often ambiguous to refer to a specific object. Therefore, the final models have limitations in precisely localizing references.

In this paper, we resolve these issues by introducing a novel framework where high-quality segmentation masks together with target-specific distinctive referring expressions are automatically generated. We achieve this by leveraging multiple foundation models with strong generalization powers. The automatically generated data then serve as pseudo RIS annotations that can train any supervised RIS methods at scale. Moreover, the broad domain coverage of the foundation models makes our pseudo-supervisions appropriate for the open-world nature of RIS.

In particular, we incorporate a segmentation foundation model, namely SAM [31], and a text-generative foundation model such as CoCa [90]. The former extracts precise class-agnostic segmentation masks; the latter then generates a

target-specific caption given the image region obtained from one of the generated masks as the input. However, although this naive approach may generate a valid caption for the input image region, the generated caption may not distinctively refer to the target object. For example, Fig. 1a shows a correct but non-distinctive caption in purple for the input image region of a generated instance mask. The caption "a cow with a tail" correctly describes the content of the input region but does not distinctively refer to the target region; there is another cow with a tail, making the generated caption not appropriate as a referring expression for the input region. In contrast, the other caption "a brown cow with a long tail" in red distinctively refers to the target region of the single cow.

To generate such a distinctive caption, we propose two techniques that allow distinctive caption generation, as in Fig. 1b: 1) *distinctive caption sampling* for generating candidate referring expressions with details in the target region, and 2) *distinctiveness-based text filtering* that further validates distinctiveness of generated captions. These two techniques enable the final captions to not only accurately describe the target, but also distinguish it from other objects within the same image, making the generated caption appropriate for RIS training. Furthermore, our method allows to naturally integrate the generalization powers of the strong foundation models into our generated dataset and consequently, the trained models on our framework achieve superior performances in various challenging RIS configurations.

Our contributions can be summarised as follows:

- We introduce a novel framework of automatically generating pseudo supervision for RIS that allows to train any supervised RIS methods with no human labeling cost.
- By incorporating with foundation models, our method not only generates precise pseudo supervisions for RIS but also successfully addresses generalization problems in supervised RIS.
- We propose distinctive caption sampling for generating caption candidates with target-specific expressions and distinctiveness-based text filtering for further validation of the generated referring expressions.
- Our extensive set of experiments demonstrates that trained models under our framework achieve state-of-the-art performances in various RIS configurations and show strong generalization performances in grounding unseen objects.

2 Related Work

2.1 Pseudo Supervision Generation.

Several works have explored generating pseudo supervision [8,14,26,43,45] without human labels to address expensive annotation costs in tasks like natural language video localization [52], video sentence localization [93], video corpus

moment retrieval [24], and referring expression comprehension [22]. In videorelated tasks [24,52,93], pseudo temporal event boundaries are generated within videos, followed by pseudo language queries using captioning models [36]. In referring expression comprehension task, Pseudo-Q [22] extract object bounding boxes with their category and attributes utilizing pre-trained detector [1] on Visual Genorm dataset [32], then they use heuristic algorithms to produce pseudo language queries based on the template comprising class name, attribute, and relationships. Unlike Pseudo-Q [22], instead of heuristic methods, we automatically generate pseudo language queries using an image captioning model [90] combined with our proposed decoding strategy.

2.2 Decoding Strategy.

Decoding strategy [30, 35, 48, 56, 63, 68, 84] is a vital in auto-regressive language modeling like GPT [3, 59] for progressing from previously generated words to the next work. The decoding strategies can be grouped into search [37, 54, 70] or sampling approaches [2, 15, 49]. In search-based approaches, beam search maximizes the multiple of the probabilities of each word in a sequence of words. Contrastive decoding [37] maximizes the word probability difference between larger and smaller language models. In sampling-based approaches, top-k and top-p [15] sampling sample words from the fixed top-k or top-p word probability distribution. Mirostate sampling [2] uses the statistical information of previously generated words to chose the next word. Our proposed strategy utilizes the word distributions of other images to sample the next word focused on a target image.

2.3 Referring Image Segmentation.

Fully supervised RIS approaches [4, 5, 7, 9, 11, 12, 16-18, 20, 21, 23, 28, 29, 33, 39, 41, 42, 46, 57, 64, 65, 72-74, 77, 78, 80-83, 85, 87-89, 94] have achieved remarkable performance by effectively fusing two different modalities. However, they highly depend on expensive human annotated data. To handle this, weakly and zeroshot RIS methods have emerged as potential solutions, recently. Weakly supervised methods [27, 34, 40, 69] use only image-text pairs for training. They mainly deal with overcoming the lack of localization for referred instances. *Lee et al.* [34] utilize Grad-CAM [62] for localization, and SaG [27] exploits slot attention [44] to capture individual entities. Zero-shot RIS methods [53, 71, 91] solve the RIS task without any additional training by leveraging the vision-language foundation models [58, 60] with the instance segmentation model [31, 75]. Unlike current weakly supervised and zero-shot RIS methods, we generate precise pseudo-supervisions for RIS to handle costly human annotation.

3 Method

In this section, we present Pseudo-RIS, the novel framework for RIS that automatically generates distinctive referring expressions paired with their corresponding target segmentation masks. We note that getting a distinctive referring



Fig. 2: Distinctive caption candidates generation of our Pseudo-RIS. Given segmentation masks, we generate multiple distinctive caption candidates on each mask using a frozen image captioning model with the proposed *distinctive caption sampling*. This method calibrates the word distribution of the target by utilizing that of others, and then samples the next word from a calibrated distribution of the target.

text for each target region is critical particularly in RIS, as the resulting mask should change when a non-distinctive caption is given. For instance, in Fig. 1(a), the mask for the non-distinctive caption should change to cover all cows that have a tail. The proposed Pseudo-RIS consists of three steps: 1) precise mask extraction, 2) distinctive caption candidate generation, and 3) distinctiveness-based text filtering. In the mask generation step, we first obtain precise instance-level segmentation mask using a segmentation foundation model such as SAM [31]. Then, a text-generative foundation model, *i.e.*, large-scale captioning model [90], is used to sample an extensive set of candidate referring expressions for each of the generated masks using our proposed distinctive caption sampling method. Finally, we further evaluate the captions generated for the given masks using our distinctiveness metric derived from CLIP [58] scores, and filter out ambiguous and imprecise captions [19]. In the following sections, we further elaborate on each step in more detail.

3.1 Mask Extraction

We first extract segmentation masks from an unlabeled image by exploiting a segmentation foundation model, SAM [31]. It allows us to obtain high-quality and class-agnostic masks in the wild, leading to valuable pseudo-supervisions appropriate for the open-world nature of RIS.

3.2 Distinctive Caption Candidates Generation

Given extracted segmentation masks within an image, we employ an image-text foundation model, such as CoCa [90], to generate multiple distinctive caption candidates for each mask. To generate distinctive captions, we propose a new decoding strategy, termed *distinctive caption sampling* based on the captioning model. In the proposed decoding strategy, the captioning model is encouraged to produce target-specific words capturing dicriminative features of the target image region, which allows the generated captions to distinctively refer only to the target region. For example, among several people as in Fig. 2, a target is unique identified by the description "man wearing a red tie" due to the targetspecific word "red". In contrast, more generic words like "tie" or "suit" do not provide the same level of distinctiveness. In this section, we demonstrate how to generate such distinctive captions using the frozen captioning model with our distinctive caption sampling.

Generating Caption for a Target Mask. Given a cropped patch for a given target mask as visual context, the image captioning model generates a sequence of m words $\mathbf{y} = \langle y_1, y_2, \ldots, y_m \rangle$ from the vocabulary $\mathcal{V} \ (y \in \mathcal{V})$, that form the description representing the image region of the target mask. Consequently, the frozen auto-regressive captioning model computes a joint probability distribution of the sequence of words, $P(\mathbf{y}|x_i)$, in a left-to-right and token-by-token manner:

$$P(\mathbf{y}|x_i) = \prod_{t=1}^{m} p(y_t|y_{< t}, x_i, C),$$
(1)

where x_i is a cropped patch for a target *i* with margins to capture surrounding contexts and *C* denotes a set of all cropped patches for all extracted masks in a given image. Each word in a description is generated by a specific decoding strategy, such as beam search, top-*k*, or top-*p* [15] sampling, which determines how the model progresses from the previously generated sequence of words $y_{<t} = \langle y_1, \ldots, y_{t-1} \rangle$ to the next word y_t at each *t* step. We use different values of *k* and *p*, along with different cropping margins for the patches, to generate diverse multiple caption candidates. For further details about these decoding strategies, which are commonly used in text generation tasks [10, 15], please refer to our supplementary material.

Distinctive Caption Sampling. This aims to sample the next word that is more specific to a target mask at each step. To this end, we calibrate the probabilities of the words for a target patch by penalizing high-probability words in other patches and re-warding low-probability words in other patches. In detail, for a given target cropped patch x_i , we calculate the probability distribution of the next word $P(y_t|y_{< t}, x_i)$, which is obtained by applying the softmax function to model output lozit z_t (unnormalized score) at t step. Simultaneously, we also compute this probability for all other cropped patches as $P(y_t|y_{< t}, x_j)$, where $x_j \in C \setminus \{x_i\}$. Notice that $y_{< t}$ is previously generated words for a target patch x_i . Then, the calibrated word probability at t step is given by:

$$P(y_t|y_{(2)$$

where $\sigma(\cdot)$ is the softmax function with a temperature, n = |C| - 1 denotes the number of the extracted masks except a target *i* in a given image, and s_{ij} is a cosine similarity between visual embeddings of x_i and x_j , extracted from the visual encoder of the image captioning model. s_{ij} modulates the influence of other patches, increasing it for patches similar to a target patch x_i , and decreasing it for those not similar, thereby aiding in sampling words that distinguish similar objects better. We sample the next word y_t from this calibrated distribution:

$$y_t \sim P(y_t | y_{\le t}, x_i, C). \tag{3}$$

The detailed illustration of the proposed distinctive caption sampling is shown in Fig. 2. In the original word distribution of a target, the generic words ("tie" and "suit") have a high probability, but after calibrating the word distribution of a target, generic words are suppressed, and target-focused word "red" has the highest probability, which is more likely to be selected for the next word y_t .

3.3 Distinctiveness-based Text Filtering

In this section, we present our distinctiveness-based text filtering method, that filter out ambiguous and imprecise captions given all the generated caption candidates for each mask, by utilizing the discriminative power of a foundation model, CLIP [58], to validate whether the captions correctly describes the target masks. The generated caption candidates may sometimes ambiguously refer to the unintended object or even incorrectly describe a non-targeted mask located around a target mask within a cropped patch. This is because, (1) we employ a captioning model trained only on noisy non-curated web datasets, causing noisy captions, (2) a cropped patch for a target could contain other masks, (e.g. a horse appears in a cropped patch for a girl, as in Fig. 3, thus a captioning model could describe these non-targeted masks. To tackle these issues, we introduce a distinctiveness score that evaluates how well the generated captions uniquely and correctly identify the target mask. Using this score, misleading caption candidates that have low distinctiveness scores can be effectively filtered out.

Uniqueness. First, we demonstrate how to measure a description **uniquely** referring to its target. We calculate the CLIP score between a caption and all masks and then divide the CLIP score for a target mask by the highest CLIP score among all other masks. This is formulated as:

$$\operatorname{UoS}_{i} = \frac{\theta(x_{i}, \mathbf{y})}{\max_{j \in C \setminus \{x_{i}\}} \theta(x_{j}, \mathbf{y})},\tag{4}$$

where $\theta(x, \mathbf{y})$ is a function calculating the CLIP score between inputs. x_i and \mathbf{y} are a cropped patch and a description for a target *i*. x_j is a cropped patch for other instance *j*, where $j \in C \setminus \{x_i\}$. By dividing the CLIP score for the target mask with a given description by the maximum CLIP score among other masks, the UoS score reflects the relevance of the caption to the target compared to



Fig. 3: Distinctiveness-based text filtering of our Pseudo-RIS. On each mask, we filter out non-distinctive caption candidates with a *distinctiveness score* below a threshold τ . This distinctiveness score measures how well a caption distinctively refers to a target mask, with examples illustrating: (a) incorrect caption describing a non-targeted horse, resulting in a low correctness despite its high uniqueness, (b) ambiguous caption referring to an unintended mask, yielding low uniqueness but high correctness, and (c) a target-specific distinctive caption, achieving high uniqueness and correctness scores, leading to high distinctiveness.

other masks. If this ratio is close to or lower than 1, it indicates that the caption may refer to at least one other unintended mask. However, the limitation of uniqueness is its inadequate ability to capture captions that correctly point to the target, caused by the simple calculation of the CLIP score between visual and textual embeddings without any in-depth consideration of whether the correct target object is being described. This issue, depicted in Fig. 3(a), highlights the need for another term to capture a description that correctly points to a target.

Correctness. To assess whether a caption describes the correct target or nontargeted object, we again employ the CLIP score. But this time, we focus on the masked region of objects, excluding the surrounding context, and on a target noun phrase that identifies the subject described in the caption. Formally, the correctness score is computed as follows:

$$\operatorname{CoS}_{i} = \theta(x_{i} \odot m_{i}, \mathcal{T}_{NP}(\mathbf{y})), \tag{5}$$

where \odot is a Hadamard product operation, m_i is a segmentation mask for a target *i*, and $\mathcal{T}_{NP}(\cdot)$ is a function that extracts a target noun phrase, as in [91].

Distinctiveness. We finally define distinctiveness scores by extending uniqueness scores using the correctness term in Eq (5) as follows:

$$\mathrm{DoS}_{i} = \frac{\mathrm{CoS}_{i} \times \theta(x_{i}, \mathbf{y})}{\max_{j \in C \setminus \{x_{i}\}} \mathrm{CoS}_{j} \times \theta(x_{j}, \mathbf{y})}.$$
(6)

In this manner, our distinctiveness score effectively measures how well a description **uniquely** and **correctly** refers to the target mask relative to other masks. We present examples of our distinctiveness score in Fig. 3. On each mask, we only keep the candidates whose distinctiveness score is above a threshold τ and these are paired with their corresponding mask to form our pseudo supervision.

4 Implementation Details

We utilize SAM [31] as a mask generator and CoCa [90] trained only on webcrawled datasets (LAION-2B [61] and CC3M [66]) as a captioning model. For distinctiveness-based text filtering, we adopt CLIP [58] with ResNet-50 [13] and set a threshold τ to 1.3. We generate our pseudo supervision from 16,992 unlabeled images in the training set of RefCOCO+. We evaluate our pseudo supervisions using two RIS models: CRIS [78] and ETRIS [85]. In all our experiments, we follow their configurations without modifications.

Mask Generation. We extract segmentation masks using SAM [31]. It generates an excessive number of masks and results in 2,026,942 masks from 16,992 raw images, making a significant computational challenge. To address this, we reduce the mask count from overwhelming \sim 120 SAM masks per image to a manageable range of \sim 5 masks by taking overlaid SAM masks with CutLER [76] masks, which is an unsupervised segmentation model. The visualization of this process and the performance using CutLER masks are reported in the supplementary.

Text Generation. For each segmentation mask, we generate 4 distinct types of cropped patches, each with a different margin setting which is 0%, 10%, 20%, and a special 0% with the masked region of objects, influencing the contents of the generated caption. For each cropped patch, we generate 11 caption candidates using a captioning model by employing three decoding strategies (Beam search and our distinctive caption sampling with fixed top-k or top-p words), each with different hyperparameters (top-k words by $k \in \{5, 7, 9, 11, 13\}$ or top-p words [15] by $p \in \{0.4, 0.5, 0.6, 0.7, 0.8\}$). Therefore, we generated a total of 44 caption candidates, and those captions with distinctiveness scores below the threshold τ are discarded. The temperature is used as a default value in CoCa.

5 Experiments

5.1 Datasets and Metrics.

In our experiments, we use RefCOCO [51], RefCOCO+ [51], RefCOCOg [25,47] and PhraseCut [79] datasets. RefCOCO, RefCOCO+, RefCOCOg, and PhraseCut provide 19,994, 19,992, 26,711, and 77,262 images and 50,000, 49,856, 54,822, and 354,496 annotated instances with multiple referring descriptions, respectively. In particular, PhraseCut offers a large number of instance categories (3101 classes), in contrast RefCOCO datasets which handle only 80 categories. The referring descriptions in each dataset have different characteristics. RefCOCO provides positional cues of objects within the image such as *left* and *right*, it is banned in RefCOCO+, and RefCOCOg has longer and more complex sentences than others. PhraseCut presents a structured text form that includes attributes, categories, and relationships. For evaluation, standard evaluation metrics in RIS are used: oIoU (overall Intersection over Union) and mIoU (mean Intersection over Union). oIoU tends to penalize more in cases of failure in larger instances,

Table 1: oIoU comparision with zero-shot and weakly-supervised RIS methods. Our pseudo supervised RIS models perform a zero-shot evaluation on RIS benchmark datasets. U and G denote UMD and Google splits of RefCOCOg, respectively. S and T are segmentation and text supervision by humans. † denotes the method utilizing multiple foundation models (*e.g.*, Stable diffusion, SAM, and CLIP) as ours.

Sup.	Human	Methods	Mask	RefCOCO			RefCOCO+			RefCOCOg		
Types	Label		Gen.	val	testA	testB	val	testA	testB	val(U)	$\operatorname{test}(U)$	val(G)
Weakly	Т	SaG [27]		30.40	29.75	31.24	25.27	24.46	26.04	27.12	-	-
		TRIS [40]		28.12	28.68	27.68	28.33	27.20	29.00	33.63	33.89	33.24
Zaro shot		Global-Local [91]	SAM	24.55	26.00	21.03	26.62	29.99	22.23	28.92	30.48	30.19
Zero-snot	-	$\operatorname{Ref-Diff}^{\dagger}[53]$	$_{\rm SAM}$	35.16	37.44	34.50	35.56	38.66	31.40	38.62	37.50	37.82
Pseudo	-	Ours + CRIS	SAM	35.45	39.69	30.37	37.55	41.52	31.71	37.90	39.64	39.32
		Ours + ETRIS	SAM	37.33	43.43	31.90	40.19	46.43	33.63	41.63	43.52	42.97
Fully	$^{\rm S+T}$	CRIS [78]		67.35	71.54	62.16	57.94	64.05	48.42	56.56	57.38	-
		ETRIS [85]		69.36	72.68	63.68	58.59	66.93	50.14	58.62	60.12	55.56

Table 2: mIoU comparision with zero-shot and weakly-supervised RIS methods.

Sup.	Human	Methods	Mask	RefCOCO			RefCOCO+			RefCOCOg		
Types	Label		Gen.	val	testA	testB	val	testA	testB	val(U)	test(U)	val(G)
Weakly	Т	Lee <i>et al.</i> [34]		31.06	32.30	30.11	31.28	32.11	30.13	32.88	-	-
		SaG [27]		34.76	34.58	35.01	28.48	28.60	27.98	28.87	-	-
		TRIS [40]		31.17	32.43	29.56	30.90	30.42	30.80	36.19	36.23	36.00
Zero-shot	-	Global-Local [91]	SAM	31.83	32.93	28.64	34.97	37.11	30.61	40.66	40.94	41.68
		$\operatorname{Ref-Diff}^{\dagger}[53]$	$_{\rm SAM}$	37.21	38.40	37.19	37.29	40.51	33.01	44.02	44.51	44.26
Pagado	-	Ours + CRIS	SAM	39.75	44.77	32.99	42.15	46.28	34.54	43.73	43.37	43.78
r seuuo		Ours + ETRIS	$_{\rm SAM}$	41.05	48.19	33.48	44.33	51.42	35.08	45.99	46.67	46.80
Fully	$^{\rm S+T}$	CRIS [78]		70.34	73.41	65.72	62.13	67.68	53.42	60.30	60.29	-
		ETRIS [85]		71.06	74.11	66.66	62.23	68.51	52.79	60.28	60.42	57.86

because it calculates the total intersection divided by the total union across all predictions and ground truths masks. In contrast, mIoU treats each instance equally, regardless of its size as it computes the IoU for each individual sample and then averages these values across all samples.

5.2 Results

Main Results. In Table 1 and 2, we compare our pseudo-supervised method with zero-shot, weakly and fully supervised methods, in terms of oIoU and mIoU, respectively. We perform a zero-shot evaluation on RefCOCO, RefCOCO+ and RefCOCOg using two RIS models, ETRIS [85] and CRIS [78], trained with our pseudo supervisions generated from 16,994 raw images in the training set of RefCOCO+. For a fair comparison with zero-shot methods, we re-implement Global-Local CLIP [91] using SAM. As shown in Table 1 and 2, our method significantly surpasses both zero-shot [53,91], and weakly supervised RIS methods [27,34,40] by large margins. We also report the scores of the fully supervised SoTA RIS models as the upper bound.

Cross-domain Results. The robustness of a model to work with consistent performance across domains is crucial for real-world applications. In this context,

Supervision Types Image Dataset RefCOCO RefCOCO+ RefCOCOg RefCOCO+ 64.54 62.23 60.35 Fully PhraseCut 24.94 -39.60 25.34 -36.89 26.41 -33.94 \overline{Pseudo} RefCOCO+ 41.0544.3345.99PhraseCut 40.65 -0.40 43.25 -1.08 45.92 -0.07 (Ours)

Table 3: Cross-domain mIoU comparison between our pseudo-supervised and fullysupervised ETRIS models in two different domains, *e.g.* RefCOCO and PhraseCut. Our pseudo-supervisions are generated using only train set images in the image dataset.

Table 4: Cross-domain and open-world results on PhraseCut. We generate our psuedosupervisions using only train set images in the image dataset. *Unseen* column denotes evaluation on a subset of object categories not included in RefCOCO datasets. † denotes the upper bound score of ETRIS where the training and testing domains are the same.

Supervision	Imago	PhraseCut Test Set						
Turper	Detect	n	nIoU	oIoU				
Types	Dataset	All	Unseen	All	Unseen			
Zero-shot [91]	N/A	29.73	28.69-1.04	31.25	$31.30_{+0.05}$			
	RefCOCO	16.28	14.97-1.31	15.18	14.26-0.92			
E. H.	RefCOCO+	16.14	$14.94_{-1.20}$	14.93	$14.11_{-0.72}$			
Fully	RefCOCOg	13.03	11.92-1.11	11.80	$10.96_{-0.84}$			
	PhraseCut	38.10 [†]	-	41.62 [†]	-			
Pseudo	RefCOCO+	30.02	30.11+0.09	28.68	$29.14_{+0.46}$			
(Ours)	PhraseCut	33.70	$33.95_{\mathbf{+0.25}}$	32.75	$33.52_{\pm 0.77}$			

we report the performances of both our pseudo-supervised and fully supervised ETRIS models trained on two different domains (*e.g.* RefCOCO variants and PhraseCut) in Table **3** and **4** for RefCOCO and PhraseCut datasets, respectively. Our pseudo-supervisions are generated from raw images in a train set of the image dataset. We observe that the performances of the supervised model trained with human-labeled data from PhraseCut or RefCOCO datasets drop significantly when evaluated on different domains. This performance degradation is due to domain gaps [67] caused by differences in textual styles and object categories handled in each dataset. In contrast, our pseudo-supervised model shows consistent performances across different domains, validating a level of robustness superior to that of the fully supervised model trained with manual annotations.

Open-world Results. To validate the generalization ability of our Pseudo-RIS model compared to the fully supervised one, we extend our zero-shot evaluation to an unseen domain, PhraseCut [79], as shown in Table 4. Since PhraseCut has a larger number of instance categories (*i.e.* 3,103 classes) than RefCOCO datasets (80 classes from MS-COCO [38]), we choose PhraseCut to validate the generalization ability in open-world setting. In Table 4, *Unseen* column denotes when evaluated on a subset of instance categories not included in RefCOCO datasets as in [91]. Note that the fully supervised model shows a significant drop in performance on the unseen subset, whereas our pseudo-supervised models are robust and even improve performance on this subset. These results highlight the



Fig. 4: mIoU results of two models: our Pseudo-RIS in a semi-supervised setting and a fully supervised model, across different volumes of human-labeled data. Our method uses both our pseudo-supervisions with varying proportions of human-labeled data, while the fully supervised method only relies on human labeled data. Both models are based on the ETRIS [85]. We also provide performance comparisons with other semi-supervised RIS methods [50, 86, 92] in our supplementary.

superior generalization ability of our pseudo-supervised model in light of the open-world nature of RIS.

Results on Semi-supervised Learning. Fig. 4 illustrates a comparison between our method and a supervised approach in a semi-supervised setting, both using the ETRIS model. In this comparison, our method is first trained with our pseudo supervisions and then fine-tuned with labeled data, while the supervised method is trained solely on proportions of human-labeled data (0.1%, 1%, 5%, 10%, and up to 100%). Remarkably, our approach achieves promising results even when just 0.1% labeled data. As the proportion of labeled data increases to 1%, 5% and 10%, we observe a consistent improvement in our performance, demonstrating the effectiveness of our method even with minimal labeled data. These results highlight the potential of our method, especially in real-world scenarios where obtaining large amounts of labeled data is expensive and challenging. Additionally, we provide a performance comparison with other semi-supervised RIS methods [50, 86, 92] in our supplementary materials.

5.3 Ablation Study

We conduct ablation studies on the val set of each RefCOCO dataset to validate the effectiveness of our method. We adopt ETRIS as our base model trained with our pseudo supervisions in all experiments.

Impact of Proposed Components. In this ablation study, we analyze the impact of each component in our method, as illustrated in Table 5. The results show that the integration of our *distinctive caption sampling* into the naïve method significantly improves the performances across all datasets in terms of mIoU scores. We also demonstrate that our *distinctiveness-based text filtering* improves mIoU scores. Furthermore, combining *distinctiveness-based text filtering* with *distinctive caption sampling* leads to further improvements in mIoU scores. These results indicate that our two strategies are more effective than the naïve method in generating distinctive captions for each mask.

Table 5: mIoU results with different components of our approach. DCS and filtering denote our distinctive caption sampling and distictiveness-based text filtering methods, respectively.

DCS	Filtering	RefCOCO	RefCOCO+	RefCOCOg
		37.96	39.87	42.08
\checkmark		40.01 + 2.05	43.27 +3.40	44.46 + 2.38
	\checkmark	40.40 + 2.44	43.40 + 3.53	45.09 + 3.01
\checkmark	\checkmark	41.05 + 3.09	44.33 + 4.46	45.99 + 3.91

Table 6: mIoU results of ETRIS model with full supervisions, GT segmentation masks with our generated pseudo text supervisions, and solely our pseudo-supervisions. S and T indicate GT segmentation masks and GT texts used for training.

Supervision	Human	RefCOCO+			RefCOCOg			
Types	Label	val	testA	testB	val(U)	test(U)	val(G)	
Fully	S+T	62.23	68.51	52.79	60.28	60.42	57.86	
Pseudo	S	52.54	60.24	42.43	54.66	54.78	55.82	
(Ours)	-	44.33	51.42	35.08	45.99	46.67	46.80	

Effects on GT segmentation masks. In Table 6, we show the impact of quality on segmentation masks using ETRIS as a base model. By using GT masks to generate our distinctive text supervisions, we achieve remarkable performances that are close to a fully supervised model. These results highlight the powerful ability of our method to generate distinctive expressions for RIS.

Ablation on threshold τ . We conduct an ablation study on the threshold τ , which is crucial for filtering out misleading captions through our *distinctiveness metric*. The mIoU scores under various threshold τ values are reported in our supplementary material. We choose $\tau = 1.3$ for all datasets.

Qualitative Analysis. We visualize the expressions generated by our method and those generated by the naïve method in Fig. 5a, as well as compare ours with GT expressions in RefCOCO datasets in Fig. 5b. A naïve method generates nondistinctive captions, which make it hard to distinguish the target object well from other objects. In contrast, our method produces distinctive captions describing the details of the target mask. Compared to GT expressions, ours provide more variety of visual concepts in the target mask, making the RIS models robust and general to different domains.

Limitations and Discussions. The proposed method is not explicitly designed to generate phrases describing object positions (*e.g. left, right*) or relations between objects (*e.g. on top of*). However, we find that it can still produce such phrases to some extent as showcased in Fig 6a and 6b. This capability is due to two factors: (1) we crop patches with margins that include contextual information around the target, and (2) we employ a captioning model trained on web datasets that include positional and relational words. We plan to explore this matter further in future work.



(b) Comparison of our generated expressions with GT expressions.

Fig. 5: Qualitative analysis of our generated expressions compared to the naïve methods and GT captions.



(b) Examples of our captions with relations between objects.

Fig. 6: Qualitative analysis of our generated expressions with positions and relations.

Social Impacts. Since our method relies on foundation models that are trained on large-scale datasets, it also inherits their limitations, such as dataset bias and unfairness issues. Specifically, we utilize the image captioning model trained with non-curated web-crawled large-scale datasets, and it may generate captions that include discriminatory or biased languages present in the dataset. This highlights the need to carefully monitor the generated expressions for identifying and correcting biased or harmful content.

6 Conclusion

In this paper, we propose a novel framework, Pseudo-RIS, to generate pseudosupervisions for RIS, where two techniques are introduced: (1) generating multiple candidates with distinctive words, and (2) filtering out the misleading captions, resulting in distinctive captions, like manual supervisions. Our approach outperforms SoTA weakly and zero-shot approaches, even fully supervised methods on an unseen domain. Furthermore, we show effectiveness of our method in a semi-supervised setting with a few labeled data. Acknowledgements. This work was supported by the IITP grants (No.2019-0-01842 (3%), No.2021-0-02068 (2%), No.2022-0-00926 (50%), No.2024-2020-01819 (5%)) funded by MSIT, Culture, Sports and Tourism R&D Program of the KOCCA grant (RS-2024-00345025 (5%)) funded by MCST, and the GIST-MIT Research Collaboration grant (35%) funded by GIST, Korea.

References

- Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. In: CVPR (2018) 4
- Basu, S., Ramachandran, G.S., Keskar, N.S., Varshney, L.R.: Mirostat: A perplexity-controlled neural text decoding algorithm. In: ICLR (2021) 4
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., others.: Language models are fewshot learners. In: NeurIPS (2020) 4
- Chng, Y.X., Zheng, H., Han, Y., Qiu, X., Huang, G.: Mask grounding for referring image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 26573–26583 (2024) 4
- 5. Cho, Y., Yu, H., Kang, S.J.: Cross-aware early fusion with stage-divided vision and language transformer encoders for referring image segmentation. TMM (2023) 4
- 6. Choi, M.: Referring object manipulation of natural images with conditional classifier-free guidance. In: ECCV (2022) 1
- 7. Dai, Q., Yang, S.: Curriculum point prompting for weakly-supervised referring image segmentation. In: CVPR (2024) 4
- Dang, R., Feng, J., Zhang, H., Chongjian, G., Song, L., Lijun, G., Liu, C., Chen, Q., Zhu, F., Zhao, R., et al.: Instructdet: Diversifying referring object detection with generalized instructions. In: ICLR (2024) 3
- 9. Ding, H., Liu, C., Wang, S., Jiang, X.: Vision-language transformer and query generation for referring segmentation. In: ICCV (2021) 1, 4
- Fan, A., Lewis, M., Dauphin, Y.: Hierarchical neural story generation. In: ACL (2018) 6
- 11. Feng, G., Hu, Z., Zhang, L., Lu, H.: Encoder fusion network with co-attention embedding for referring image segmentation. In: CVPR (2021) 4
- 12. Feng, G., Hu, Z., Zhang, L., Lu, H.: Encoder fusion network with co-attention embedding for referring image segmentation. In: CVPR (2021) 4
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) 9
- 14. Holla, M., Lourentzou, I.: Commonsense for zero-shot natural language video localization. In: AAAI (2024)3
- Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. In: ICLR (2019) 4, 6, 9
- Hu, Y., Wang, Q., Shao, W., Xie, E., Li, Z., Han, J., Luo, P.: Beyond one-to-one: Rethinking the referring image segmentation. In: ICCV (2023) 4
- 17. Hu, Z., Feng, G., Sun, J., Zhang, L., Lu, H.: Bi-directional relationship inferring network for referring image segmentation. In: CVPR (2020) 4
- 18. 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 (2020) 4

- 16 Yu et al.
- Huang, Y., Wu, Z., Gao, C., Peng, J., Yang, X.: Exploring the distinctiveness and fidelity of the descriptions generated by large vision-language models. arXiv preprint arXiv:2404.17534 (2024) 5
- 20. Huang, Z., Satoh, S.: Referring image segmentation via joint mask contextual embedding learning and progressive alignment network. In: EMNLP (2023) 4
- 21. 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 (2020) 4
- Jiang, H., Lin, Y., Han, D., Song, S., Huang, G.: Pseudo-q: Generating pseudo language queries for visual grounding. In: CVPR (2022) 4
- 23. Jing, Y., Kong, T., Wang, W., Wang, L., Li, L., Tan, T.: Locate then segment: A strong pipeline for referring image segmentation. In: CVPR (2021) 4
- Jung, M., Choi, S., Kim, J., Kim, J.H., Zhang, B.T.: Modal-specific pseudo query generation for video corpus moment retrieval. In: EMNLP (2022) 4
- Kazemzadeh, S., Ordonez, V., Matten, M., Berg, T.: Referitgame: Referring to objects in photographs of natural scenes. In: EMNLP (2014) 9
- Kim, D., Park, J., Lee, J., Park, S., Sohn, K.: Language-free training for zero-shot video grounding. In: WACV (2023) 3
- Kim, D., Kim, N., Lan, C., Kwak, S.: Shatter and gather: Learning referring image segmentation with text supervision. In: ICCV (2023) 2, 4, 10
- Kim, N., Kim, D., Lan, C., Zeng, W., Kwak, S.: Restr: Convolution-free referring image segmentation using transformers. In: CVPR (2022) 1, 4
- Kim, S., Kang, M., Kim, D., Park, J., Kwak, S.: Extending clip's image-text alignment to referring image segmentation. In: NAACL (2024) 4
- Kim, T., Kim, J., Lee, G., Yun, S.Y.: Instructive decoding: Instruction-tuned large language models are self-refiner from noisy instructions. In: ICLR (2024) 4
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollar, P., Girshick, R.: Segment anything. In: ICCV (2023) 2, 4, 5, 9
- 32. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. IJCV **123**, 32–73 (2017) 4
- 33. Lai, X., Tian, Z., Chen, Y., Li, Y., Yuan, Y., Liu, S., Jia, J.: Lisa: Reasoning segmentation via large language model. In: CVPR (2024) 4
- 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: ICCV (2023) 2, 4, 10
- Leng, S., Zhang, H., Chen, G., Li, X., Lu, S., Miao, C., Bing, L.: Mitigating object hallucinations in large vision-language models through visual contrastive decoding. In: CVPR (2024) 4
- 36. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: ICML (2022) 4
- Li, X.L., Holtzman, A., Fried, D., Liang, P., Eisner, J., Hashimoto, T., Zettlemoyer, L., Lewis, M.: Contrastive decoding: Open-ended text generation as optimization. In: ACL (2023) 4
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014) 11
- Liu, C., Ding, H., Jiang, X.: Gres: Generalized referring expression segmentation. In: CVPR (2023) 1, 4
- 40. Liu, F., Liu, Y., Kong, Y., Xu, K., Zhang, L., Yin, B., Hancke, G., Lau, R.: Referring image segmentation using text supervision. In: ICCV (2023) 2, 4, 10

- 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 (2023) 4
- Liu, S., Ma, Y., Zhang, X., Wang, H., Ji, J., Sun, X., Ji, R.: Rotated multi-scale interaction network for referring remote sensing image segmentation. In: CVPR (2024) 4
- Liu, Y., Zhang, J., Chen, Q., Peng, Y.: Confidence-aware pseudo-label learning for weakly supervised visual grounding. In: ICCV (2023) 3
- 44. Locatello, F., Weissenborn, D., Unterthiner, T., Mahendran, A., Heigold, G., Uszkoreit, J., Dosovitskiy, A., Kipf, T.: Object-centric learning with slot attention. In: NeurIPS (2020) 4
- Lu, Y., Quan, R., Zhu, L., Yang, Y.: Zero-shot video grounding with pseudo query lookup and verification. TIP 33, 1643–1654 (2024) 3
- Luo, G., Zhou, Y., Sun, X., Cao, L., Wu, C., Deng, C., Ji, R.: Multi-task collaborative network for joint referring expression comprehension and segmentation. In: CVPR (2020) 4
- 47. Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A.L., Murphy, K.: Generation and comprehension of unambiguous object descriptions. In: CVPR (2016) 9
- Meister, C., Pimentel, T., Malagutti, L., Wilcox, E., Cotterell, R.: On the efficacy of sampling adapters. In: ACL (2023) 4
- Meister, C., Pimentel, T., Wiher, G., Cotterell, R.: Locally typical sampling. TACL 11, 102–121 (2023)
- Nag, S., Goswami, K., Karanam, S.: Safari: Adaptive sequence transformer for weakly supervised referring expression segmentation. arXiv preprint arXiv:2407.02389 (2024) 12
- Nagaraja, V.K., Morariu, V.I., Davis, L.S.: Modeling context between objects for referring expression understanding. In: ECCV (2016) 9
- 52. Nam, J., Ahn, D., Kang, D., Ha, S.J., Choi, J.: Zero-shot natural language video localization. In: ICCV (2021) 3, 4
- Ni, M., Zhang, Y., Feng, K., Li, X., Guo, Y., Zuo, W.: Ref-diff: Zero-shot referring image segmentation with generative models. arXiv preprint arXiv:2308.16777 (2023) 2, 4, 10
- 54. O'Brien, S., Lewis, M.: Contrastive decoding improves reasoning in large language models. arXiv preprint arXiv:2309.09117 (2023) 4
- Padmakumar, A., Thomason, J., Shrivastava, A., Lange, P., Narayan-Chen, A., Gella, S., Piramuthu, R., Tur, G., Hakkani-Tur, D.: Teach: Task-driven embodied agents that chat. In: AAAI (2022) 1
- Qiu, Z., Ou, Z., Wu, B., Li, J., Liu, A., King, I.: Entropy-based decoding for retrieval-augmented large language models. arXiv preprint arXiv:2406.17519 (2024) 4
- 57. Qu, M., Wu, Y., Wei, Y., Liu, W., Liang, X., Zhao, Y.: Learning to segment every referring object point by point. In: CVPR (2023) 4
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021) 2, 4, 5, 7, 9
- 59. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., others.: Language models are unsupervised multitask learners. OpenAI blog (2019) 4
- 60. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) 2, 4

- 18 Yu et al.
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open largescale dataset for training next generation image-text models. In: NeurIPS (2022) 9
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Gradcam: Visual explanations from deep networks via gradient-based localization. In: ICCV (2017) 4
- Sennrich, R., Vamvas, J., Mohammadshahi, A.: Mitigating hallucinations and offtarget machine translation with source-contrastive and language-contrastive decoding. In: ACL (2024) 4
- 64. Shah, N.A., VS, V., Patel, V.M.: Lqmformer: Language-aware query mask transformer for referring image segmentation. In: CVPR (2024) 4
- Shang, C., Song, Z., Qiu, H., Wang, L., Meng, F., Li, H.: Prompt-driven referring image segmentation with instance contrasting. In: CVPR (2024) 4
- 66. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: ACL (2018) 9
- Shi, H., Pan, W., Zhao, Z., Zhang, M., Wu, F.: Unsupervised domain adaptation for referring semantic segmentation. In: ACM (2023) 11
- Shi, W., Han, X., Lewis, M., Tsvetkov, Y., Zettlemoyer, L., Yih, S.W.t.: Trusting your evidence: Hallucinate less with context-aware decoding. arXiv preprint arXiv:2305.14739 (2023) 4
- Strudel, R., Laptev, I., Schmid, C.: Weakly-supervised segmentation of referring expressions. arXiv preprint arXiv:2205.04725 (2022) 2, 4
- 70. Su, Y., Lan, T., Wang, Y., Yogatama, D., Kong, L., Collier, N.: A contrastive framework for neural text generation. In: NeurIPS (2022) 4
- Suo, Y., Zhu, L., Yang, Y.: Text augmented spatial aware zero-shot referring image segmentation. In: EMNLP Findings (2023) 4
- 72. Tang, J., Zheng, G., Shi, C., Yang, S.: Contrastive grouping with transformer for referring image segmentation. In: CVPR (2023) 4
- Wang, W., He, X., Zhang, Y., Guo, L., Shen, J., Li, J., Liu, J.: Cm-masksd: Crossmodality masked self-distillation for referring image segmentation. TMM (2024) 4
- Wang, W., Yue, T., Zhang, Y., Guo, L., He, X., Wang, X., Liu, J.: Unveiling parts beyond objects: Towards finer-granularity referring expression segmentation. In: CVPR (2024) 4
- 75. Wang, X., Yu, Z., De Mello, S., Kautz, J., Anandkumar, A., Shen, C., Alvarez, J.M.: Freesolo: Learning to segment objects without annotations. In: CVPR (2022) 4
- Wang, X., Girdhar, R., Yu, S.X., Misra, I.: Cut and learn for unsupervised object detection and instance segmentation. In: CVPR (2023) 9
- 77. Wang, Y., Li, J., Zhang, X., Shi, B., Li, C., Dai, W., Xiong, H., Tian, Q.: Barleria: An efficient tuning framework for referring image segmentation. In: ICLR (2024) 4
- 78. Wang, Z., Lu, Y., Li, Q., Tao, X., Guo, Y., Gong, M., Liu, T.: Cris: Clip-driven referring image segmentation. In: CVPR (2022) 1, 4, 9, 10
- Wu, C., Lin, Z., Cohen, S., Bui, T., Maji, S.: Phrasecut: Language-based image segmentation in the wild. In: CVPR (2020) 9, 11
- Wu, J., Li, X., Li, X., Ding, H., Tong, Y., Tao, D.: Towards robust referring image segmentation. TIP (2024) 4

- 81. Wu, Y., Zhang, Z., Xie, C., Zhu, F., Zhao, R.: Advancing referring expression segmentation beyond single image. In: ICCV (October 2023) 1, 4
- 82. Xia, Z., Han, D., Han, Y., Pan, X., Song, S., Huang, G.: Gsva: Generalized segmentation via multimodal large language models. In: CVPR (2024) 4
- Xu, L., Huang, M.H., Shang, X., Yuan, Z., Sun, Y., Liu, J.: Meta compositional referring expression segmentation. In: CVPR (2023) 4
- Xu, Z., Jiang, F., Niu, L., Jia, J., Lin, B.Y., Poovendran, R.: Safedecoding: Defending against jailbreak attacks via safety-aware decoding. arXiv preprint arXiv:2402.08983 (2024) 4
- Xu, Z., Chen, Z., Zhang, Y., Song, Y., Wan, X., Li, G.: Bridging vision and language encoders: Parameter-efficient tuning for referring image segmentation. In: ICCV (2023) 1, 4, 9, 10, 12
- Yang, D., Ji, J., Ma, Y., Guo, T., Wang, H., Sun, X., Ji, R.: Sam as the guide: Mastering pseudo-label refinement in semi-supervised referring expression segmentation. In: ICML 12
- 87. Yang, S., Xia, M., Li, G., Zhou, H.Y., Yu, Y.: Bottom-up shift and reasoning for referring image segmentation. In: CVPR (2021) 4
- 88. Yang, Z., Wang, J., Tang, Y., Chen, K., Zhao, H., Torr, P.H.: Lavt: Language-aware vision transformer for referring image segmentation. In: CVPR (2022) 4
- Ye, L., Rochan, M., Liu, Z., Wang, Y.: Cross-modal self-attention network for referring image segmentation. In: CVPR (2019) 4
- Yu, J., Wang, Z., Vasudevan, V., Yeung, L., Seyedhosseini, M., Wu, Y.: Coca: Contrastive captioners are image-text foundation models. TMLR (2022) 2, 4, 5, 9
- Yu, S., Seo, P.H., Son, J.: Zero-shot referring image segmentation with global-local context features. In: CVPR (2023) 2, 4, 8, 10, 11
- 92. Zang, Y., Fu, C., Cao, R., Zhu, D., Zhang, M., Hu, W., Zhu, L., Chen, T.: Resmatch: Referring expression segmentation in a semi-supervised manner. arXiv preprint arXiv:2402.05589 (2024) 12
- Zheng, M., Gong, S., Jin, H., Peng, Y., Liu, Y.: Generating structured pseudo labels for noise-resistant zero-shot video sentence localization. In: ACL (2023) 3, 4
- 94. Zou, X., Dou, Z.Y., Yang, J., Gan, Z., Li, L., Li, C., Dai, X., Behl, H., Wang, J., Yuan, L., et al.: Generalized decoding for pixel, image, and language. In: CVPR (2023) 4