

Advancing Referring Expression Segmentation Beyond Single Image

Yixuan Wu^{12*} Zhao Zhang^{1*} Chi Xie¹³ Feng Zhu^{1✉} Rui Zhao¹⁴
¹SenseTime Research ²Zhejiang University ³Tongji University
⁴Qing Yuan Research Institute, Shanghai Jiao Tong University

wyx_chloe@zju.edu.cn z Zhang@mail.nankai.edu.cn zhufeng@sensetime.com

Abstract

Referring Expression Segmentation (RES) is a widely explored multi-modal task, which endeavors to segment the pre-existing object within a single image with a given linguistic expression. However, in broader real-world scenarios, it is not always possible to determine if the described object exists in a specific image. Generally, a collection of images is available, some of which potentially contain the target objects. To this end, we propose a more realistic setting, named Group-wise Referring Expression Segmentation (GRES), which expands RES to a group of related images, allowing the described objects to exist in a subset of the input image group. To support this new setting, we introduce an elaborately compiled dataset named Grouped Referring Dataset (GRD), containing complete group-wise annotations of the target objects described by given expressions. Moreover, we also present a baseline method named Grouped Referring Segmenter (GRSer), which explicitly captures the language-vision and intra-group vision-vision interactions to achieve state-of-the-art results on the proposed GRES setting and related tasks, such as Co-Salient Object Detection and traditional RES. Our dataset and codes are publicly released in <https://github.com/shikras/d-cube>.

1. Introduction

Segmenting target objects described by users in a collection of images is a fundamental but overlooked capability that facilitates various real-world applications (as illustrated in Fig. 1), such as Internet images automatic labeling, multi-monitors joint event discovery, and mobile album retrieval. In recent years, Referring Expression Segmentation (RES) has become a research hotspot with great potentials to solve this demand. Various promising approaches [17, 38, 44, 6, 14] and datasets [21, 46, 36, 42] have contributed to significant advancements in this field.

* Equal Contribution

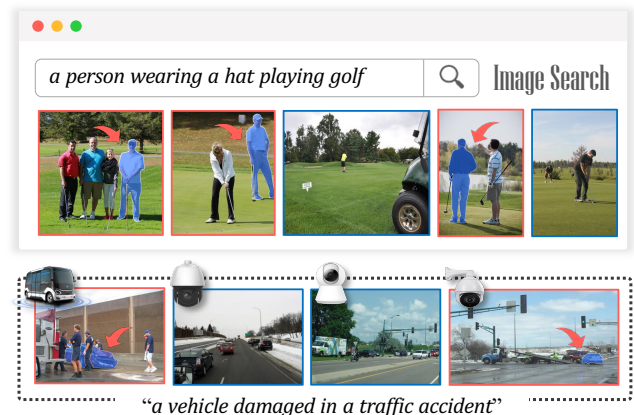


Figure 1: **Real-world applications** of our proposed Group-wise Referring Expression Segmentation (GRES) setting, which facilitates Internet images automatic labeling (upper), multi-monitors joint event discovery (lower), etc.

However, the setting of RES is overly idealistic, which aims to segment what has been known to exist in a single image. This has restricted the practicality of RES in real-world situations, given that it is not always possible to determine if the described object exists in a specific image. Generally, a group of images is available, only some of which may potentially contain the target objects.

To address this limitation, in this paper, we introduce a new realistic setting, namely Group-wise Referring Expression Segmentation (GRES), and define it as segmenting objects described in language expression from a group of related images. Compared with traditional RES task, our GRES setting get extended in two points: (a) tolerance for negative images where no target object exists, and (b) simultaneous processing of a group of images. Note that the application scope of GRES is the **superset** of RES, which means (a) or (b) can be omitted based on the actual scenarios. Moreover, we establish the foundation of GRES in two aspects: firstly, a baseline method named Grouped Referring Segmenter (GRSer) that explicitly leverages language and intra-group vision connections to obtain promis-

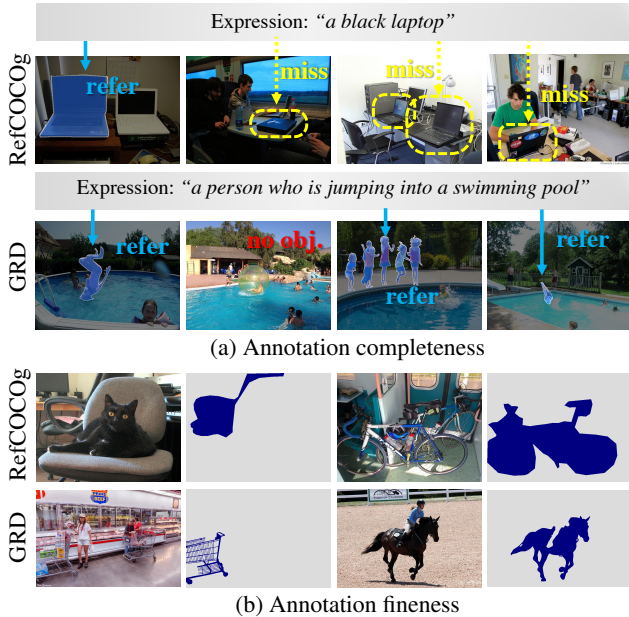


Figure 2: Proposed GRD vs. RefCOCOg on the annotation completeness and annotation fitness.

ing results, and secondly, a meticulously annotated dataset, Group Referring Dataset (GRD), that ensures complete annotations of described objects across all images in a group.

Our proposed baseline method GRSer, illustrated in Fig. 3, facilitates a simultaneous processing of multiple input images with an expression, and generates segmentation masks for all described objects. We devise a Triphasic Query Module (TQM), where the target objects not only queried by linguistic features, but also by intra-group visual features. In contrast to segmenting based solely on linguistic expression, querying target objects with intra-group homo-modal visual features bridges the modal gap and assembles a more precise target concept. In the proposed Heatmap Hierarchizer, heatmaps generated by intra-group visual querying are ranked based on their confidences, and then jointly used to predict segmentation masks in condition of the ranking priorities. Furthermore, we propose a mirror training strategy and triplet loss to learn anti-expression features, which are crucial for the TQM and Heatmap Hierarchizer, and enable GRSer to comprehend the image background and negative samples. The promising performance of GRSer makes it a strong research baseline for our proposed GRES setting.

To facilitate the research in novel GRES setting, the GRD dataset is introduced, which effectively overcomes the incomplete annotation problem in current RES datasets [21, 46, 36]. For example, in Fig. 2 (a), RefCOCOg’s expression of the 1st image also corresponds to objects in images 2, 3, and 4, but the target objects are not annotated. This leads to false positive samples during evaluation even though the

target objects are correctly segmented by the model. In contrast, in our GRD dataset, the described objects are completely annotated in all images across the dataset, including images without the targets or with multiple targets. Additionally, GRD collects images from Internet search engines by a group of keywords, where negative samples inherently exist in each group, making them hard negatives and effectively increasing the dataset’s difficulty. Finally, as shown in Fig. 2 (b), compared with current RES datasets, GRD carefully labels details in segmentation masks, such as blocking and hollowing-out regions, which contributes to a more objective and reliable evaluation efficacy for model performances.

Our contributions can be summarized as:

- We formalize a Group-wise Referring Expression Segmentation (GRES) setting over the RES task, which advances user-specified object segmentation towards more practical applications.
- To support GRES research, we present a meticulously compiled dataset named GRD, possessing complete group-wise annotations of target objects. The dataset will also benefit various other vision-language tasks.
- Extensive experiments show the effectiveness and generalizability of our proposed baseline method GRSer, which achieves SOTA results on the GRES and related tasks, such as Co-Salient Object Detection and RES.

2. Related Work

2.1. Referring Expression Segmentation (RES)

RES aims to ground the target object in the given image referred by the language and generate a corresponding segmentation mask. **Methods.** A common approach to solve RES is to first extract both vision and language features, and then fuse the multi-modal features to predict the mask. Early methods [17, 25, 28] simply concatenate visual features and language features extracted by convolutional neural networks (CNNs) and recurrent neural networks (RNNs), respectively. Due to the breakthrough of Transformer [41, 9, 31], a rich line of works begin to explore its remarkable fusion power for multi-modality. Some [43, 27, 38, 14, 6, 22, 44] conduct cross-model alignment based on Transformer, others [45, 34, 30, 35, 43] adopt various attention mechanisms to achieve better feature weighting and fusing. There are some works to explore how to solve RES working with related tasks, such as visual grounding [35, 24, 55, 29], zero/one-shot segmentation [33], interactive segmentation [5], unified segmentation [57], and referring expression generation [18]. **Datasets.** Several datasets have been introduced to evaluate the performance of RES methods, including RefClef [21], RefCOCO [46], RefCOCO+ [46], RefCOCOg (G-Ref) [36], and PhraseCut [42]. RefClef, RefCOCO,

and RefCOCO+ are collected interactively in a two-player game, named ReferitGame [21], thus the given expressions are more concise and less flowery. Among them, RefCOCO+ bans location words in expressions, making it more challenging. RefCOCOg is collected non-interactively, resulting in more complex expressions, often full sentences instead of phrases. PhraseCut’s phases, consist of attribute, category, and relationship, are automatically generated by predefined templates and existing annotations from Visual Genome [23]. The above datasets fail to serve as reliable evaluation datasets for GRES setting due to their image-text pairs are one-to-one matched, which leads to incomplete annotation for target objects in unmatched images. More datasets comparison can be found in Tab. 1.

2.2. Co-Salient Object Detection (Co-SOD)

Co-SOD is a recent research focus [50, 53, 20, 54, 49, 51, 13, 52, 12, 47, 56], aiming to discover the common semantic objects in a group of related images. In this task, the target object does not need to be specified by language expression, while required to appear commonly in all images. Co-SOD methods need to perceive what the common objects are from the pure visual modality, and then segment them. Historically, researchers refer to Co-SOD as “detection”, but its outputs are actually segmentation maps. **Methods.** Recently, many impressive Co-SOD methods have arisen, focusing primarily on obtaining co-representations of common objects to guide target object segmentation. Co-representations can be obtained through methods like feature concatenation [39], linear addition [54], channel shuffling [53], graph neural networks [19, 51], and iterative purification [56]. There is also a body of research work focused on intra-group information exchange, such as using pair-wise similarity map [20], dynamic convolution [52], group affinity [13], and transformers [16, 37]. Moreover, besides these central lines of exploring, efforts have been made to enhance the Co-SOD model through data enhancement [54, 52], training strategies [13, 47], adversarial attack preventing [15]. **Datasets.** Co-SOD datasets include iCoseg [1], MSRC [40], CoSal2015 [48], CoSOD3k [12], and CoCA [54]. Early datasets such as iCoseg and MSRC contain co-salient objects with similar appearance in similar scenes. CoSal2015 and CoSOD3k are large-scale datasets, featuring target objects with varying appearance in the same category. CoCA, the latest dataset, presents a more challenging setting with at least one extraneous salient object in each image, requiring the model to identify the target object in cluttered scenes. Although the data sets have favorable grouping scenarios, they lack expressions and negative samples, making them unsuitable for direct use as evaluation dataset for GRES.

3. Proposed Method

3.1. Overview

The pipeline of our Grouped Referring Segmenter (GRSer) is demonstrated in Fig. 3. Given an expression that specifies an object, a group of related images are processed simultaneously, and then all corresponding pixel-wise masks of the target object are output. In particular, for the negative sample (*i.e.*, image without target object), its output mask is 0 mask. There are four modules in our GRSer, including a multi-modal encoder, a triphasic query module, a heatmap hierarchizer, and a mask predictor.

Text & Image Encoder. BERT [4] is employed to embed the expression into linguistic features $\mathbf{L} \in \mathbb{R}^{C_l}$, where C_l is the number of channels for the language feature. Meanwhile, we construct an anti-expression by adding a prefix <no> to the given expression, which is embedded as linguistic anti-features $\mathbf{L}^{anti} \in \mathbb{R}^{C_l}$. We follow LAVT [44] to perform visual encoding to obtain visual features $\mathbf{V}_n \in \mathbb{R}^{C_v \times H \times W}$ for each image x_n in the group ($n = 1, \dots, N$), where N is the number of images in one group, and C_v, H , and W denote the channel number, height, and width, respectively.

TQM & Heatmap Hierarchizer. The language-vision and intra-group vision-vision semantic relations are explicitly captured in proposed TQM (Sec. 3.2) to produce heatmaps, which reflect the spatial relation between linguistic and intra-group visual features. And these heatmaps are ranked and rearranged in heatmap hierarchizer (Sec. 3.3) according to their importance with the expression to better activate their locating capability for mask prediction.

Mask Predictor. The well-ranked heatmaps are concatenated with visual features \mathbf{V}_n to obtain the triphasic features \mathbf{z}_n , which integrates the discriminative cues of target object in TQM and heatmap hierarchizer. \mathbf{z}_n is used to distinguish positive or negative samples, and predict the segmentation masks. **In inference**, the positive distance $d^{pos} = d(\mathbf{z}_n, \mathbf{L})$ and negative distance $d^{neg} = d(\mathbf{z}_n, \mathbf{L}^{anti})$ are computed, where the Euclidean Distance $d(\cdot)$ is applied. If $d^{pos} + m < d^{neg}$ (m is the margin value), then image x_n is recognized as a positive sample, and its \mathbf{z}_n is then transmitted to the decoder to output segmentation mask. If not, 0 mask is reassigned as negative output.

3.2. Triphasic Query Module (TQM)

Due to the inherent modality gap, directly querying objects through linguistic features often results in rougher language-activated heatmaps (*e.g.*, the 2nd image in the bottom row of Fig. 4). We resort to intra-group homomodal visual features to act as “experts”, offering suggested heatmaps from their perspectives. To this end, we devise the TQM, where “triphasic” means that the target object not only queried by linguistic features, but also by intra-group

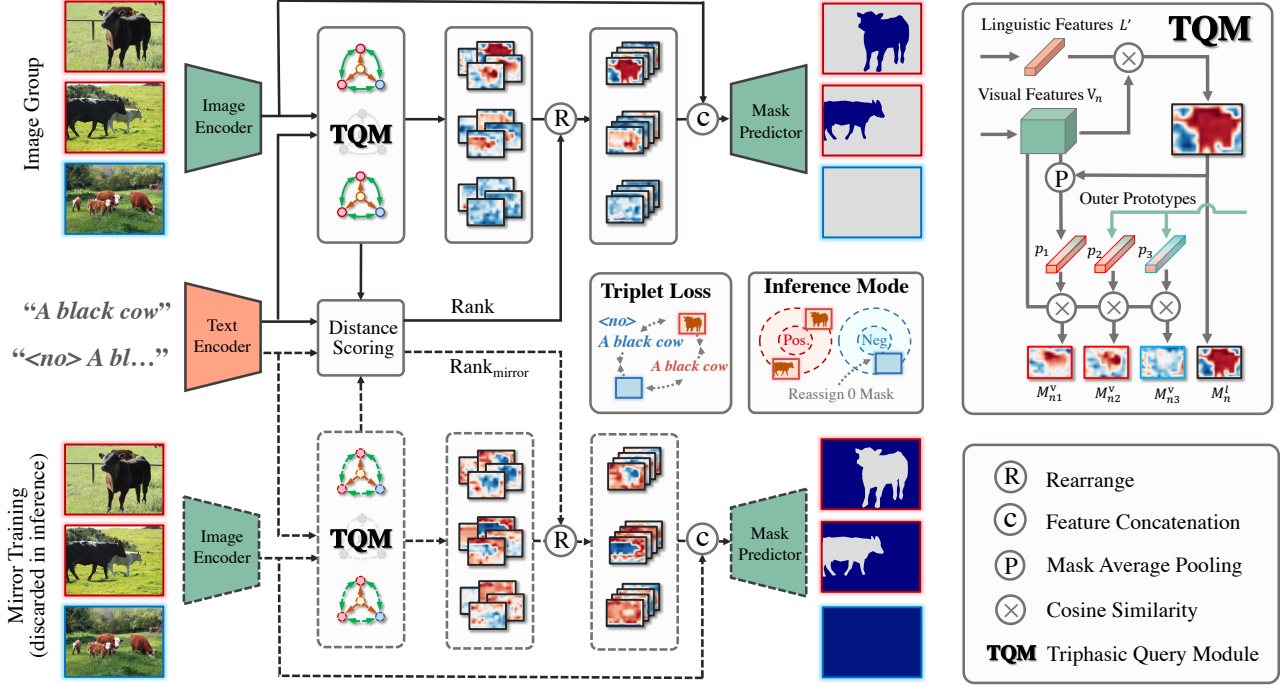


Figure 3: **The pipeline of proposed GRSer.** First, grouped input images, given an expression together with its anti-expression (<no> prefix added), are encoded by image and text encoder, respectively, and fed into a triphasic query module (TQM), to generate a set of heatmaps that indicate the most discriminating region in the visual feature map responding to the target object. Next, these heatmaps are rearranged according to their correlation with the description, and then concatenated with visual features for mask prediction. In training, triplet loss and segmentation loss are both applied, and a mirror training strategy (dotted line) is introduced to better comprehend the anti-expression and image background. In inference, the mirror training will be discarded, and images close to the anti-expression are reassigned 0 masks.

homo-modal visual features.

In the right top of Fig. 3, we take one image x_n as an example to illustrate the detailed process. First, in order to detect the most discriminating region in the visual feature map responded to the referring expression, a language-activated heatmap $\mathbf{M}_n^l \in \mathbb{R}^{H \times W}$ is generated. Specifically, the cosine similarity is computed between the flattened visual features $\mathbf{V}_n \in \mathbb{R}^{C_v \times HW}$ and linguistic features $\mathbf{L}' = \omega_l(\mathbf{L}) \in \mathbb{R}^{C_v}$, where a 1×1 convolution layer ω_l with C_v number of output channels are deployed to align the cross-modal features. This is denoted as

$$\mathbf{M}_n^l = \frac{\mathbf{V}_n^T \cdot \mathbf{L}'}{\|\mathbf{V}_n\| \|\mathbf{L}'\|}. \quad (1)$$

Second, \mathbf{M}_n^l is element-wise multiplied with visual features \mathbf{V}_n , and the output features are averaged along spatial dimension (*i.e.*, $H \times W$) with mask average pooling to generate a prototype $\mathbf{p}_n \in \mathbb{R}^{C_v}$ corresponding to image x_n , as

$$\mathbf{p}_n = \text{avg}(\mathbf{M}_n^l \odot \mathbf{V}_n), \quad (2)$$

where \mathbf{M}_n^l is broadcast to the same size as \mathbf{V}_n , and \odot denotes the element-wise multiplication. In this manner, a

group of prototypes $\{\mathbf{p}_i\}_{i=1}^N$ is generated, with each prototype corresponding to one image from a group. Intuitively, the prototype integrates visual features of the target object.

Next, the intra-group queries are conducted between current image x_n and a group of prototypes, and these prototypes serve as “experts” to provide localization heatmap suggestions from their perspectives.

In details, the cosine similarity is computed between the flattened visual features \mathbf{V}_n and each prototype \mathbf{p}_i from $\{\mathbf{p}_i\}_{i=1}^N$ one-by-one, and then produce N vision-activated heatmaps $\mathbf{M}_n^v = \{\mathbf{M}_{ni}^v\}_{i=1}^N$, as

$$\mathbf{M}_{ni}^v = \frac{\mathbf{V}_n^T \cdot \mathbf{p}_i}{\|\mathbf{V}_n\| \|\mathbf{p}_i\|}, \quad (3)$$

where n denotes the index of image in a group, and i denotes the index of prototype in a group. As shown in Fig. 4, these four \mathbf{M}_{ni}^v (the 3rd - 6th in the bottom row) show stronger locating capability than the \mathbf{M}_n^l (the 2nd in the bottom row), which thus provide more accurate guidance for mask prediction.

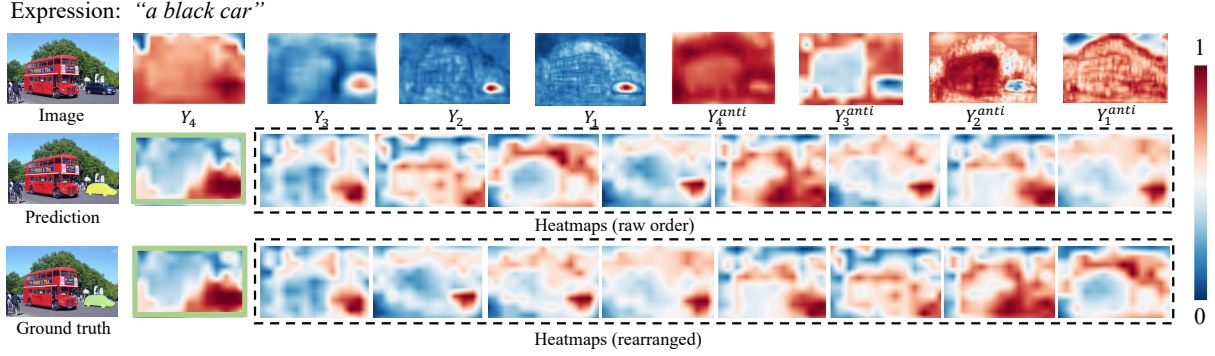


Figure 4: **Visualizations of the prediction mask, feature maps, and heatmaps** on an example from the G-RefCg test set. The leftmost column demonstrates the input image, predicted mask (in yellow) and ground-truth mask (in green). In the first row of other columns, we visualize the feature maps in decoder (*i.e.*, Y_i) activated by linguistic feature L , and feature maps in decoder (*i.e.*, Y_i^{anti}) activated by linguistic anti-feature L^{anti} in our proposed mirror training process (Sec. 3.4). In the second and third row, we demonstrate the language-activated heatmap M^l (the one inside the solid green border) and vision-activated heatmaps M^v (the others inside the dotted border). Note that the raw order of heatmaps are shown in the second row, and the well-ranked heatmaps are demonstrated in the third row. Best viewed in color.

3.3. Heatmap Hierarchizer

Considering that the vision-activated heatmaps suggested by “experts” from TQM can be uneven, especially when there are negative samples. For example, in Fig. 4, prototypes come from negative samples tend to generate counterfactual localization heatmaps (the 7th - 10th in the bottom row). We need experts to give confidence of their suggestions to determine the heatmap priority in following prediction. To this end, we propose a heatmap hierarchizer to rank and rearrange these vision-activated heatmaps based on a confidence evaluation strategy.

To get the rank of different heatmaps, we define a scoring criterion based on the multi-modal representation distance. Specifically, we compute the Euclidean Distance [7] between each prototype from the group $\{p_i\}_{i=1}^N$ and linguistic features L to get the positive score $\{s_i^{pos}\}_{i=1}^N$. Negative score $\{s_i^{neg}\}_{i=1}^N$ is also obtained by computing Euclidean Distance between $\{p_i\}_{i=1}^N$ and linguistic anti-features L^{anti} . In this way, a smaller s_i^{pos} indicates that the prototype p_i gets closer to the target object, which means its corresponding generated heatmap M_{ni}^v is more reliable. Inversely, a smaller s_i^{neg} indicates the prototype p_i fits the background (*i.e.*, outside of the target object in a image) better. Then, we obtain the positive rank R^{pos} and negative rank R^{neg} for N vision-activated heatmaps $M_n^v = \{M_{ni}^v\}_{i=1}^N$, according to corresponding positive score s_i^{pos} (from smallest to largest) and negative score s_i^{neg} (from largest to smallest), respectively. The positive rank R^{pos} and negative rank R^{neg} are summed as the final rank to rearrange M_n^v by

$$\bar{M}_n^v = \text{rearrange}(M_n^v | R^{pos} + R^{neg}), \quad (4)$$

where $\text{rearrange}(\cdot)$ means changing the channel-wise or-

der of these stacked heatmaps. These heatmaps are then concatenated with visual features V_n to get triphasic features z_n for mask prediction. In Fig. 4, it can be seen that heatmaps with lower confidence (generated by negative samples) are relegated to the back after rearrangement.

3.4. Training Objectives

Training with Negative Samples. For training, we set the ratio between positive samples x^{pos} (*i.e.*, image containing target object referred by the expression) and negative samples x^{neg} (*i.e.*, noisy image where no target object exists) in each image group as 1 : 1. The training objectives are twofold: (1) Triplet margin loss to empower model with recognition ability for negative samples; (2) Cross-entropy loss to optimize the model’s segmentation performance.

Triplet Margin Loss. The goal of triplet margin loss [8] is to bring closer together the anchor and the positive example, while pull the anchor from the negative example away, as is illustrated in Eq. 5. The Euclidean Distance $d(\cdot)$ is applied, and m is the margin value. For a positive sample x^{pos} , its triphasic features z_n is regarded as the anchor, and linguistic features L and anti-features L^{anti} are regarded as the positive and negative examples, respectively. And for a negative sample x^{neg} , L^{anti} and L are regarded as its positive and negative examples instead. The triplet margin loss is computed as

$$\mathcal{L}_{tri} = \begin{cases} \max(d(z_n, L) - d(z_n, L^{anti}) + m, 0) & \text{for } x^{pos} \\ \max(d(z_n, L^{anti}) - d(z_n, L) + m, 0) & \text{for } x^{neg} \end{cases} \quad (5)$$

Mirror Training Strategy. To further force our model to comprehend the semantics contained in linguistic anti-features L^{anti} , we design a mirror training strategy. In-

tuitively, linguistic anti-features represent the opposite semantics of the given expression, and thus we explicitly relate the linguistic anti-features to the image background (*i.e.*, outside of the target object in an image). Specifically, during training, on the basis of original pipeline, we add an additional mirror one that swaps the roles of \mathbf{L} and \mathbf{L}^{anti} , and corresponding ground-truth mask is replaced with the background (*i.e.*, $1 - \mathbf{Y}$, where \mathbf{Y} denotes the ground-truth mask for the target object). As shown in the first row of Fig. 4, the feature maps (*i.e.*, \mathbf{Y}_i^{anti}) in decoder activated by \mathbf{L}^{anti} exactly focus on the background outside of the target object. The cross-entropy loss is applied for mirror training, denoted as \mathcal{L}_{ce}^{mirr} .

Objective Function. Note that only positive samples x^{pos} are included for computing cross-entropy loss, while all samples (*i.e.*, x^{pos} and x^{neg}) are used for computing triplet margin loss. We adopt the increasing weighting strategy for triplet margin loss to optimize the training process, by

$$\mathcal{L} = \mathcal{L}_{ce}(\hat{\mathbf{Y}}, \mathbf{Y}) + \lambda \mathcal{L}_{ce}^{mirr}(\hat{\mathbf{Y}}^{anti}, 1 - \mathbf{Y}) + \frac{t}{T} \mathcal{L}_{tri}, \quad (6)$$

where t and T denote the current training epoch and total number of training epochs, respectively; λ is a hyper-parameter to weigh the importance of mirror training strategy; $\hat{\mathbf{Y}}$ denotes predicted mask referred by the expression, and $\hat{\mathbf{Y}}^{anti}$ denotes predicted mask referred by the anti-expression obtained in mirror training strategy.

4. Proposed Dataset

4.1. Dataset Highlights

Within a collection of images, for a given expression, **we label all described objects in all images without any omission.** This constitutes the fundamental attribute that distinguishes GRD from its counterparts. For instance, as shown in Fig. 2 (a), in RefCOCOg’s, for a given expression “*a black laptop*”, the target object is only labeled in a single image, while the same objects are not labeled in all other images. This flaw renders the expression valid only in one image, **making other images in the dataset unsuitable as negative samples.** In addition to complete annotation, there are also some features that make GRD exceptional. One is that the images are collected in a group-wise manner, where images in each group of GRD are related. Thus, even if the described target does not appear on some images in a group, the scenes behind these images are often close to the description, which makes our GRD dataset more challenging. Additionally, as shown in Fig. 2 (b), our meticulous annotations can lead to a more objective and precise evaluation of model performance compared to current RES datasets. More features can be found in Tab. 1. Thanks to these features, GRD can help many other vision-language tasks, such as RES, visual grounding, and grounding caption. *GRD is freely available for non-commercial research purposes.*

Table 1: **Valuable features bring by GRD dataset.** “*scene grouping*” means samples are grouped by similar scenes. “*complete annotation*” means any object satisfying the given description is annotated across dataset. In this case, samples without the label for a specific expression could be reliably considered as “*certified negative samples*” for this expression. If there are “*multiple referred objects*” described in an image, all of them are annotated without omission. “*meticulous masks*” are provided to fits the object perfectly, especially for the hollowed-out and blocking areas. “*object-centric*” means the dataset concentrates on objects rather than broad concepts like grass and sky. “*avg. expression length*” represents the average expression length. RC, RC+, RCg, RCF, and PC denote RefCOCO, RefCOCO+, RefCOCOg, RefClef, and PhaseCut, respectively.

	RC	RC+	RCg	RCF	PC	GRD
scene grouping	✗	✗	✗	✗	✗	✓
complete annotation	✗	✗	✗	✗	✗	✓
certified neg. samples	✗	✗	✗	✗	✗	✓
multi. referred objects	✗	✗	✗	✗	✓	✓
meticulous masks	✗	✗	✗	✗	✗	✓
object-centric	✓	✓	✓	✗	✓	✓
avg. expression length	3.6	3.5	8.4	3.5	2.0	5.9

4.2. Construction Procedures

We collect images searching from Flickr¹. If crawling directly according to the expression, we usually get the iconic images, which appear in profile, unobstructed near the center of a neatly composed photo. In order to meet the real situation and increase the challenge, we employ the combination of target keywords and scene keywords to crawl a group of related images from search engines. Consequently, the images involves intricate scenes, *i.e.*, non-iconic images [26]. Then, for each group of images, we carefully propose several related expressions to be annotated. The announcers will segment the objects in the group according to these expressions, without excluding any referred objects. This completeness allows our dataset to accurately assess the model’s performance on negative samples. Each object annotation takes an average of 3 minutes to precisely define edges and remove hollow areas, guaranteeing accurate evaluation of model segmentation performances.

4.3. Dataset Statistics.

The GRD dataset contains 10,578 images. It includes 106 scenes (groups), such as indoor, outdoor and sports ground. Each group has around 100 images and 3 well-designed expressions referring to various number of positive and negative samples. In total, the dataset is anno-

¹<https://www.flickr.com>

tated with 316 expressions, resulting in 31,524 positive or negative image-text pairs. The expressions have an average length of 5.9 words.

5. Experiments

5.1. Datasets and Metrics

To comprehensively evaluate GRSer’s performance, apart from the proposed GRD, we also introduce RES and Co-SOD datasets as supplements. For RES dataset (*e.g.*, RefCOCO [46], RefCOCO+ [46], and RefCOCOg [36]), given that there exist some repeated sentences in different images, we reconstruct these datasets to the form of “one sentence *vs.* a group of referred images”, named as G-RefC, G-RefC+, and G-RefCg, with randomly sampled negative samples from other groups. These re-built datasets have 8717, 8020, and 2451 image groups, respectively, with positive to negative sample ratio of 1 : 1 for both training and inference. In our experiment, GRD and re-built RES datasets are regarded as RES setting if the negative samples of the dataset are removed, otherwise it is the GRES setting. Besides, we use the CoCA [54] dataset to evaluate our model’s performance in Co-SOD task, where we take category names as expression inputs.

We adopt the metric of mean intersection-over-union (mIoU) for evaluating model’s performance in RES setting with no negative samples included. When negative samples are introduced, their corresponding ground-truth masks are 0 mask, where the originally defined mIoU is not valid (*i.e.*, $\text{IoU} \equiv 0$, for the negative sample). Therefore, we define an adapted metric $\overline{\text{mIoU}}$ to measure model performance on both segmentation accuracy and recognition ability for negative samples. Specifically, the idea of confusion matrix is adopted: for a true positive sample (TP), its $\overline{\text{IoU}}$ is calculated in the same way as the vanilla IoU; for a true negative sample (TN), its $\overline{\text{IoU}}$ is set to 1; for a false positive sample (FP) or false negative sample (FN), its $\overline{\text{IoU}}$ is set to 0. Then, the $\overline{\text{IoU}}$ value of all m test samples are averaged to get the $\overline{\text{mIoU}}$, *i.e.*, $\overline{\text{mIoU}} = \frac{1}{m} \sum_{i=1}^m \overline{\text{IoU}}_i$. Besides, for Co-SOD task, common metrics of mean absolute error (MAE)[3], maximum F-measure [2] (F_{\max}), S-measure [10] (S_α), and mean E-measure [11] (E_ϵ) are adopted.

5.2. Implementation Details

The Transformer layers for visual encoding are initialized with classification weights pre-trained on ImageNet-22K from the Swin Transformer [31]. The language encoder is the base BERT [4] with 12 layers and hidden size of 768 (*i.e.*, C_l), which is implemented from HuggingFace’s Transformer library [41]. C_v is set to 512. Following [31, 44], the AdamW optimizer [32] is adopted with weight decay of 0.01. The initial learning rate is set to 0.00005 with polynomial learning rate decay. The model

is trained for 80 epochs with batch size of 4. Images are resized to 416×416 and no data augmentations are employed. The size of input image group N is set to 8. The margin value m is set to 1 in triplet margin loss.

5.3. Comparison with SOTA Methods

Results on the GRES Setting. In Tab. 2, we compare our GRSer with other RES methods on the re-built G-RefC, G-RefC+, G-RefCg, and our proposed GRD datasets. Specifically, negative samples are introduced to each image group for both training and inference (see Sec. 5.1 for details), where the adapted metric $\overline{\text{mIoU}}$ is used. Compared methods are implemented following their original paradigms to input data in the form of “one image *vs.* one expression”. The ground-truth for the negative sample is set as 0 mask. Note that the proposed GRD dataset is only used for inference, and its corresponding train set is the combination of train sets from G-RefC, G-RefC+ and G-RefCg. It can be seen that our GRSer significantly outperforms other methods, and excels in recognition of negative samples, due to our designed triplet loss and mirror training strategy, which effectively optimize the multi-modal representation space.

Results on the RES Setting. In Tab. 2, we present the results in the conventional RES setting, where mIoU metric is adopted. Here, no negative sample is included and all images in a group do contain target objects. Similarly, our GRSer outperforms all compared methods, particularly on the more difficult G-RefCg and GRD dataset (the given expressions are complex and hard to understand by models). It is the triphasic feature interactions in TQM (Sec. 3.2) that help our model comprehend the complex semantics of the same object from different images in a group.

Results on the Co-SOD Task. In Tab. 3, we further compare our GRSer with methods in the Co-SOD task on the CoCA dataset, where metrics including mean absolute error (MAE), maximum F-measure (F_{\max}), S-measure (S_α), and mean E-measure (E_ϵ) are adopted. Our model is trained on the combination of train sets from G-RefC, G-RefC+ and G-RefCg. The given category names of CoCA are regarded as expressions for grouped images during implementation. It can be seen that our method also achieves remarkable performances on this challenging real-world dataset.

5.4. Ablation Studies

Triphasic Query Module (TQM). We remove the proposed TQM, and only a single language-activated heatmap is concatenated with visual features and then fed to the mask predictor. In Tab. 5, the removal of TQM leads to a $\overline{\text{mIoU}}$ drop of 9.04% and 9.50% in G-RefCg and GRD, respectively, validating the effects of TQM. Besides, we try different image numbers in one group. In Tab. 6, when increasing the group size N , model performances get better.

Heatmap Hierarchizer (HMapHier). To explore the ef-

Table 2: **Quantitative comparisons with RES methods** in terms of mean Intersection-over-Union (mIoU) for the RES setting and our proposed mIoU for the GRES setting on the G-RefC, G-RefC+, G-RefCg, and our proposed GRD datasets. The best results are marked in **bold**.

Method	Pub	GRES (with negative samples)				RES (no negative samples)			
		G-RefC	G-RefC+	G-RefCg	GRD	G-RefC	G-RefC+	G-RefCg	GRD
EFN	CVPR21 [14]	25.42	22.32	20.77	15.29	63.52	55.37	52.88	31.57
VL	PAMI22 [6]	26.87	24.38	22.83	16.58	66.98	59.14	51.73	33.78
CRIS	CVPR22 [38]	29.31	27.27	24.74	19.33	70.62	68.12	58.93	41.23
LAVT	CVPR22 [44]	30.22	27.14	24.38	18.48	75.27	67.93	59.94	39.14
GRSer	Ours	84.77	78.44	75.32	57.12	79.33	70.38	65.47	47.25

Table 3: **Quantitative comparisons with Co-SOD methods** in terms of mean absolute error (MAE)[3], maximum F-measure [2] (F_{max}), S-measure [10] (S_α), and mean E-measure [11] (E_ξ) on the CoCA [54] dataset. “ \uparrow ” means that the higher the numerical value, the better the model performance, and vice versa for “ \downarrow ”. The best results are marked in **bold**.

Metric	CSMG	GCAGC	GICD	ICNet	CoEG	DeepACG	GCoNet	CADC	CoRP	GRSer
	CVPR19 [50]	CVPR20 [51]	ECCV20 [54]	NeurIPS20 [20]	PAMI21 [12]	CVPR21 [49]	CVPR21 [13]	ICCV21 [52]	PAMI2023 [56]	Ours
MAE \downarrow	0.114	0.111	0.126	0.148	0.106	0.102	0.105	0.132	0.121	0.099
F_{max} \uparrow	0.499	0.517	0.513	0.514	0.493	0.552	0.544	0.548	0.551	0.562
S_α \uparrow	0.627	0.666	0.658	0.657	0.612	0.688	0.673	0.681	0.686	0.712
E_ξ \uparrow	0.606	0.668	0.701	0.686	0.679	-	0.739	-	0.715	0.728

Table 4: **Ablation studies of ranking criteria** in heatmap hierarchizer on the G-RefCg and the proposed GRD datasets in the GRES setting. (*) indicates default choices of our model. The best results are marked in **bold**.

Train	Test	G-RefCg			GRD		
		mIoU	E_ξ	R_{neg}	mIoU	E_ξ	R_{neg}
<i>Random</i>	<i>Random</i>	68.92	0.554	89.12	50.73	0.475	75.38
	$R^{pos} + R^{neg}$	68.42	0.550	88.74	50.24	0.470	73.37
$R^{pos} + R^{neg}$ (*)	<i>Random</i>	67.79	0.548	88.23	49.83	0.468	73.28
	$R^{pos} + R^{neg}$ (*)	75.32	0.572	95.25	57.12	0.515	81.09
R^{pos}	R^{pos}	74.57	0.570	94.32	56.28	0.502	80.08
R^{neg}	R^{neg}	73.79	0.563	94.53	56.37	0.507	80.23

Table 5: **Ablation studies of main designs** in our method on G-RefCg and GRD datasets in the GRES setting.

	G-RefCg			GRD		
	mIoU	E_ξ	R_{neg}	mIoU	E_ξ	R_{neg}
w/o. TQM	66.28	0.525	85.33	47.62	0.463	70.29
w/o. HMapHier	68.92	0.554	89.12	50.73	0.475	75.38
w/o. MirrorT	69.38	0.543	90.12	51.47	0.479	75.54
w/o. TriLoss	30.37	0.493	0	23.14	0.435	0
Full model	75.32	0.572	95.25	57.12	0.515	81.09

ffects of the heatmap order in HMapHier, we experiment with different ranking criteria. As shown in Tab. 4, removing HMapHier (*i.e.*, heatmap orders in both training and testing are random) results in the mIoU drops of 6.40% and 6.39% in G-RefCg and GRD, respectively. Besides, inconsistent ranking criteria in training and testing resulted in inferior performance. Also, using the combination of positive rank R^{pos} and negative rank R^{neg} achieves the best results compared to using a single-source criterion.

Table 6: **Ablation studies of group size (N)** in TQM on G-RefCg and GRD datasets in the GRES setting. (*) indicates default choices of our model.

	G-RefCg			GRD		
	mIoU	E_ξ	R_{neg}	mIoU	E_ξ	R_{neg}
$N = 1$	66.28	0.525	85.33	47.62	0.463	70.29
$N = 3$	72.98	0.559	92.38	54.89	0.484	78.23
$N = 5$	74.26	0.567	94.45	56.01	0.502	80.92
$N = 8$ (*)	75.32	0.572	95.25	57.12	0.515	81.09

Mirror Training (MirrorT). In Tab. 5, removing MirrorT leads to a mIoU drop of 5.94% and 5.65% in G-RefCg and GRD, respectively. This is because MirrorT plays a vital role in forcing model to comprehend the semantics contained in anti-expressions, helping our GRSer to be better aware of the image background and negative samples.

Triplet Margin Loss (TriLoss). Tab. 5 shows that TriLoss is critical for GRSer when negative samples are included. Without TriLoss, the recall of negative samples R_{neg} falls

to 0 in both datasets, which means the model fails to recognize negative samples and output non-zero predicted masks for all images. TriLoss optimizes the multi-modal representation distances during training and constructs a well-distributed representation space that helps our model to distinguish between positive and negative samples.

6. Conclusion

In this work, we present a realistic multi-modal setting named Group-wise Referring Expression Segmentation (GRES), which relaxes the limitation of idealized setting of RES and extends it to a group of related images. To facilitate this new setting, we introduce a challenging dataset named GRD, which effectively simulates the real-world scenarios by collecting images in a grouped manner and annotating both positive and negative samples thoroughly. Besides, a novel baseline method GRSer is proposed to explicitly capture the language-vision and vision-vision feature interactions for better comprehension of the target object. Extensive experiments show that our method achieves SOTA performances on GRES, RES, and Co-SOD.

References

- [1] Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, and Tsuhan Chen. iCoseg: Interactive co-segmentation with intelligent scribble guidance. In *CVPR*, pages 3169–3176. IEEE, 2010. 3
- [2] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. *IEEE TIP*, 24(12):5706–5722, 2015. 7, 8
- [3] Ming-Ming Cheng, Jonathan Warrell, Wen-Yan Lin, Shuai Zheng, Vibhav Vineet, and Nigel Crook. Efficient salient region detection with soft image abstraction. In *ICCV*, pages 1529–1536, 2013. 7, 8
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019. 3, 7
- [5] Henghui Ding, Scott Cohen, Brian Price, and Xudong Jiang. Phraseclick: Toward achieving flexible interactive segmentation by phrase and click. In *ECCV*, pages 417–435. Springer, 2020. 2
- [6] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vlt: Vision-language transformer and query generation for referring segmentation. *IEEE TPAMI*, 2022. 1, 2, 8
- [7] Ivan Dokmanic, Reza Parhizkar, Juri Ranieri, and Martin Vetterli. Euclidean distance matrices: essential theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 32(6):12–30, 2015. 5
- [8] Xingping Dong and Jianbing Shen. Triplet loss in siamese network for object tracking. In *ECCV*, pages 459–474, 2018. 5
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [10] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. Structure-measure: A new way to evaluate foreground maps. In *ICCV*, pages 4548–4557, 2017. 7, 8
- [11] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. In *IJCAI*, pages 698–704, 7 2018. 7, 8
- [12] Deng-Ping Fan, Tengpeng Li, Zheng Lin, Ge-Peng Ji, Dingwen Zhang, Ming-Ming Cheng, Huazhu Fu, and Jianbing Shen. Re-thinking co-salient object detection. *IEEE TPAMI*, 2021. 3, 8
- [13] Qi Fan, Deng-Ping Fan, Huazhu Fu, Chi Keung Tang, Ling Shao, and Yu-Wing Tai. Group collaborative learning for co-salient object detection. *CVPR*, 2021. 3, 8
- [14] Guang Feng, Zhiwei Hu, Lihe Zhang, and Huchuan Lu. Encoder fusion network with co-attention embedding for referring image segmentation. In *CVPR*, pages 15506–15515, 2021. 1, 2, 8
- [15] Ruijun Gao, Qing Guo, Felix Juefei-Xu, Hongkai Yu, Huazhu Fu, Wei Feng, Yang Liu, and Song Wang. Can you spot the chameleon? Adversarially camouflaging images from co-salient object detection. In *CVPR*, pages 2150–2159, 2022. 3
- [16] Yanliang Ge, Qiao Zhang, Tian-Zhu Xiang, Cong Zhang, and Hongbo Bi. TCNet: Co-salient object detection via parallel interaction of Transformers and CNNs. *IEEE TCSVT*, 2022. 3
- [17] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In *ECCV*, pages 108–124. Springer, 2016. 1, 2
- [18] Shijia Huang, Feng Li, Hao Zhang, Shilong Liu, Lei Zhang, and Liwei Wang. A unified mutual supervision framework for referring expression segmentation and generation. *arXiv preprint arXiv:2211.07919*, 2022. 2
- [19] Bo Jiang, Xingyue Jiang, Ajian Zhou, Jin Tang, and Bin Luo. A unified multiple graph learning and convolutional network model for co-saliency estimation. In *ACM MM*, pages 1375–1382, 2019. 3
- [20] Wen-Da Jin, Jun Xu, Ming-Ming Cheng, Yi Zhang, and Wei Guo. ICNet: Intra-saliency correlation network for co-saliency detection. *NeurIPS*, 33, 2020. 3, 8
- [21] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014. 1, 2, 3
- [22] Namyup Kim, Dongwon Kim, Cuiling Lan, Wenjun Zeng, and Suha Kwak. ReSTR: Convolution-free referring image segmentation using transformers. In *CVPR*, pages 18145–18154, 2022. 2
- [23] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123:32–73, 2017. 3

- [24] Muchen Li and Leonid Sigal. Referring transformer: A one-step approach to multi-task visual grounding. In *NeurIPS*, volume 34, pages 19652–19664, 2021. [2](#)
- [25] Ruiyu Li, Kaican Li, Yi-Chun Kuo, Michelle Shu, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Referring image segmentation via recurrent refinement networks. In *CVPR*, pages 5745–5753, 2018. [2](#)
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. [6](#)
- [27] Chang Liu, Xudong Jiang, and Henghui Ding. Instance-specific feature propagation for referring segmentation. *IEEE TMM*, 2022. [2](#)
- [28] Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. Recurrent multimodal interaction for referring image segmentation. In *ICCV*, pages 1271–1280, 2017. [2](#)
- [29] Jiang Liu, Hui Ding, Zhaowei Cai, Yuting Zhang, Ravi Kumar Satzoda, Vijay Mahadevan, and R Manmatha. PolyFormer: Referring image segmentation as sequential polygon generation. In *CVPR*, 2023. [2](#)
- [30] Si Liu, Tianrui Hui, Shaofei Huang, Yunchao Wei, Bo Li, and Guanbin Li. Cross-modal progressive comprehension for referring segmentation. *IEEE TPAMI*, 44(9):4761–4775, 2021. [2](#)
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. [2](#), [7](#)
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [7](#)
- [33] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *CVPR*, 2022. [2](#)
- [34] Gen Luo, Yiyi Zhou, Rongrong Ji, Xiaoshuai Sun, Jinsong Su, Chia-Wen Lin, and Qi Tian. Cascade grouped attention network for referring expression segmentation. In *ACM MM*, pages 1274–1282, 2020. [2](#)
- [35] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Liujuan Cao, Chenglin Wu, Cheng Deng, and Rongrong Ji. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *CVPR*, pages 10034–10043, 2020. [2](#)
- [36] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *CVPR*, pages 11–20, 2016. [1](#), [2](#), [7](#)
- [37] Yukun Su, Jingliang Deng, Ruizhou Sun, Guosheng Lin, and Qingyao Wu. A unified transformer framework for group-based segmentation: Co-segmentation, co-saliency detection and video salient object detection. *arXiv preprint arXiv:2203.04708*, 2022. [3](#)
- [38] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. CRIS: CLIP-driven referring image segmentation. In *CVPR*, pages 11686–11695, 2022. [1](#), [2](#), [8](#)
- [39] Lina Wei, Shanshan Zhao, Omar El Farouk Bourahla, Xi Li, Fei Wu, and Yueting Zhuang. Deep group-wise fully convolutional network for co-saliency detection with graph propagation. *IEEE TIP*, 28(10):5052–5063, 2019. [3](#)
- [40] J. Winn, A. Criminisi, and T. Minka. Object categorization by learned universal visual dictionary. In *ICCV*, pages 1800–1807, 2005. [3](#)
- [41] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Transformers: State-of-the-art natural language processing. In *EMNLP*, pages 38–45, 2020. [2](#), [7](#)
- [42] Chenyun Wu, Zhe Lin, Scott Cohen, Trung Bui, and Subhansu Maji. Phrasecut: Language-based image segmentation in the wild. In *CVPR*, pages 10216–10225, 2020. [1](#), [2](#)
- [43] Sibe Yang, Meng Xia, Guanbin Li, Hong-Yu Zhou, and Yizhou Yu. Bottom-up shift and reasoning for referring image segmentation. In *CVPR*, pages 11266–11275, 2021. [2](#)
- [44] Zhao Yang, Jiaqi Wang, Yansong Tang, Kai Chen, Hengshuang Zhao, and Philip HS Torr. LAVT: Language-aware vision transformer for referring image segmentation. In *CVPR*, pages 18155–18165, 2022. [1](#), [2](#), [3](#), [7](#), [8](#)
- [45] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. MattNet: Modular attention network for referring expression comprehension. In *CVPR*, pages 1307–1315, 2018. [2](#)
- [46] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *ECCV*, pages 69–85. Springer, 2016. [1](#), [2](#), [7](#)
- [47] Siyue Yu, Jimin Xiao, Bingfeng Zhang, and Eng Gee Lim. Democracy does matter: Comprehensive feature mining for co-salient object detection. In *CVPR*, pages 979–988, 2022. [3](#)
- [48] Dingwen Zhang, Junwei Han, Chao Li, Jingdong Wang, and Xuelong Li. Detection of co-salient objects by looking deep and wide. *IJCV*, 120(2):215–232, 2016. [3](#)
- [49] Kaihua Zhang, Mingliang Dong, Bo Liu, Xiao-Tong Yuan, and Qingshan Liu. DeepACG: Co-saliency detection via semantic-aware contrast gromov-wasserstein distance. In *CVPR*, pages 13703–13712, 2021. [3](#), [8](#)
- [50] Kaihua Zhang, Tengpeng Li, Bo Liu, and Qingshan Liu. Co-saliency detection via mask-guided fully convolutional networks with multi-scale label smoothing. In *CVPR*, pages 3095–3104, 2019. [3](#), [8](#)
- [51] Kaihua Zhang, Tengpeng Li, Shiwen Shen, Bo Liu, Jin Chen, and Qingshan Liu. Adaptive graph convolutional network with attention graph clustering for co-saliency detection. In *CVPR*, pages 9050–9059, 2020. [3](#), [8](#)
- [52] Ni Zhang, Junwei Han, Nian Liu, and Ling Shao. Summarize and search: Learning consensus-aware dynamic convolution for co-saliency detection. In *ICCV*, pages 4167–4176, 2021. [3](#), [8](#)
- [53] Qijian Zhang, Runmin Cong, Junhui Hou, Chongyi Li, and Yao Zhao. CoADNet: Collaborative aggregation-and-distribution networks for co-salient object detection. In *NeurIPS*, 2020. [3](#)
- [54] Zhao Zhang, Wenda Jin, Jun Xu, and Ming-Ming Cheng. Gradient-induced co-saliency detection. In *ECCV*, pages 455–472. Springer, 2020. [3](#), [7](#), [8](#)

- [55] Chaoyang Zhu, Yiyi Zhou, Yunhang Shen, Gen Luo, Xingjia Pan, Mingbao Lin, Chao Chen, Liujuan Cao, Xiaoshuai Sun, and Rongrong Ji. SeqTR: A simple yet universal network for visual grounding. In *ECCV*, pages 598–615. Springer, 2022. [2](#)
- [56] Ziyue Zhu, Zhao Zhang, Zheng Lin, Xing Sun, and Ming-Ming Cheng. Co-salient object detection with co-representation purification. *IEEE TPAMI*, 2023. [3](#), [8](#)
- [57] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. Generalized decoding for pixel, image, and language. *arXiv preprint arXiv:2212.11270*, 2022. [2](#)