

Rotated Multi-Scale Interaction Network for Referring Remote Sensing Image Segmentation

Sihan Liu*, Yiwei Ma*, Xiaoqing Zhang*, Haowei Wang, Jiayi Ji†, Xiaoshuai Sun, Rongrong Ji
 Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
 Ministry of Education of China, Xiamen University, 361005, P.R. China

{liusihan, yiweima, 36920221153149, wanghaowei}@stu.xmu.edu.cn

jyixmu@gmail.com

{xssun, rrji}@xmu.edu.cn

Abstract

Referring Remote Sensing Image Segmentation (RRSIS) is a new challenge that combines computer vision and natural language processing. Traditional Referring Image Segmentation (RIS) approaches have been impeded by the complex spatial scales and orientations found in aerial imagery, leading to suboptimal segmentation results. To address these challenges, we introduce the Rotated Multi-Scale Interaction Network (RMSIN), an innovative approach designed for the unique demands of RRSIS. RMSIN incorporates an Intra-scale Interaction Module (IIM) to effectively address the fine-grained detail required at multiple scales and a Cross-scale Interaction Module (CIM) for integrating these details coherently across the network. Furthermore, RMSIN employs an Adaptive Rotated Convolution (ARC) to account for the diverse orientations of objects, a novel contribution that significantly enhances segmentation accuracy. To assess the efficacy of RMSIN, we have curated an expansive dataset comprising 17,402 image-caption-mask triplets, which is unparalleled in terms of scale and variety. This dataset not only presents the model with a wide range of spatial and rotational scenarios but also establishes a stringent benchmark for the RRSIS task, ensuring a rigorous evaluation of performance. Experimental evaluations demonstrate the exceptional performance of RMSIN, surpassing existing state-of-the-art models by a significant margin. Datasets and code are available at <https://github.com/Lsan2401/RMSIN>.

1. Introduction

Referring Remote Sensing Image Segmentation (RRSIS) stands at the forefront of integrating computer vision with natural language processing [18, 36–38], aiming to segment

*These authors contributed equally to this work.

†The corresponding author.

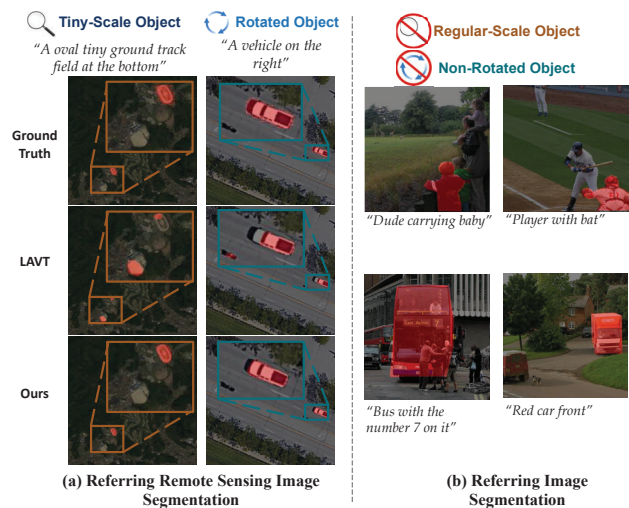


Figure 1. Comparison between the newly constructed RRSIS-D and conventional RIS datasets [62], highlighting the complex spatial scales and orientations prevalent in aerial imagery. (a) Examples from our RRSIS-D, demonstrating the limitations of traditional RIS methods (e.g., LAVT [58]) in handling such complexities. (b) Examples from a standard RIS dataset [62].

specific areas from aerial images based on textual descriptions. This sophisticated task goes beyond the capabilities of traditional RIS [13, 25, 28], requiring an intricate understanding of the spatial and geographic nuances conveyed from aerial perspectives. RRSIS plays a crucial role in a wide range of applications, including land use categorization [9], climate impact studies [44], and urban infrastructure management [8]. By pushing the boundaries of semantic understanding in remote sensing data, RRSIS is advancing the possibilities in these domains. Despite this, the field has been constrained by the limited scale and scope of existing datasets, which are insufficient for training models to the level of accuracy required for these critical tasks.

In light of these requirements, our research introduces an

expansive new benchmark dataset, namely RRSIS-D, designed to propel the development of RRSIS. This dataset surpasses its predecessors* [63] in threefold size, encompassing not only higher resolution images but also a significantly broader range of geographic diversity. The development of this dataset is guided by the Segment Anything Model (SAM) [23], which facilitates a semi-automated annotation process, thereby mitigating the labor-intensive nature of generating accurate pixel-level masks traditionally. This process involves deriving initial segmentation masks from bounding box prompts and refining them to ensure high fidelity to the complex reality of aerial imagery. The result is a comprehensive corpus of 17,402 remote sensing image-caption-mask triplets, an invaluable resource aimed at advancing the precision and utility of RRSIS.

Furthermore, although the existing RIS methodologies [17, 25, 29, 48] have demonstrated effectiveness in specific domains [39, 41, 62], they face limitations when applied to the diverse and intricate nature of remote sensing imagery. As illustrated in Fig. 1, aerial images pose distinct challenges that are not encountered in conventional image datasets, including vast and diverse spatial scales, as well as objects captured from multiple orientations. Current RIS approaches typically excel in aligning visual and linguistic elements in well-bounded contexts [16, 61] but falter when faced with the chaotic and unstructured nature of aerial images. The inability of these methods to grapple with high levels of spatial variation and rotational diversity results in a notable performance gap in RRSIS tasks, highlighting the need for a more robust and versatile approach.

To overcome the inherent limitations in existing approaches, we present the Rotated Multi-Scale Interaction Network (RMSIN), a pioneering architectural solution meticulously designed to tackle the complexities of RRSIS. Our approach introduces a sophisticated Intra-scale Interaction Module (IIM) that excels at extracting detailed features within individual layers, as well as a Cross-scale Interaction Module (CIM) that facilitates comprehensive feature fusion across the entire network. Furthermore, we integrate an Adaptive Rotated Convolution (ARC) into the decoder, empowering the model to effectively handle the intricate rotational variations exhibited by objects. By seamlessly integrating these modules, RMSIN proficiently extracts and aligns features across diverse scales and orientations, resulting in remarkable performance enhancements for RRSIS.

To sum up, our key contributions are as follows:

- We introduce RRSIS-D, a novel benchmark dataset tailored for Referring Remote Sensing Image Segmentation (RRSIS). This dataset accommodates substantial variations in both spatial resolution and object orientation.
- We propose the Rotated Multi-Scale Interaction Network

(RMSIN) to address the challenges posed by the multiple spatial scales and orientations prevalent in aerial imagery.

- We propose IIM and CIM to handle fine-grained information within and across different scales. Meanwhile, We design ARC to enhance the model’s robustness against the ubiquitous rotational phenomena in RRSIS.

2. Related work

Referring Image Detection and Segmentation. Referring Image Detection aims to predict a bounding box corresponding to a given referring expression. Existing works can be classified into two-staged methods [12, 14, 65, 67] which are based on region proposal ranking, and one-stage methods [1, 24, 26, 33–35, 57, 66] which directly predict the target bounding box. Referring Image Segmentation aims to achieve pixel-level localization of target objects within images based on associated referring expressions, presenting a more complex task [13]. Early works [25, 28, 41] leverage convolution networks and recurrent neural networks to extract vision and language features, respectively. These features are then fused by simple concatenation to generate final predictions. Subsequent methods [2, 3, 6, 19, 20, 28, 29, 40, 47, 48, 51, 54] mainly focus on vision-language alignment to enhance predictions. Some employ recurrent refinement [2, 28], while others utilize dynamic filters [3, 19, 40] to fuse visual and linguistic features. Recently, leveraging the Transformer’s outstanding performance, methods can be divided into two categories: those performing cross-modal decoder fusion based on Transformer [6, 20, 29, 48] and those incorporating language-aware visual encoding instead of post-feature fusion [22, 52, 58]. However, due to the specific characteristics of aerial images, these methods exhibit limited performance in the Remote Sensing field. Some approaches [21, 42, 60] have introduced extra scale interaction modules to enhance feature extraction. However, the extreme semantic gap between natural images and aerial images still results in suboptimal performance.

Remote Sensing Referring Image Detection and Segmentation. Referring Image Detection in Remote Sensing field is a novel task with limited research. It was first introduced by [45], where a new dataset and a baseline model were proposed. Recently, the transformer-based method RSVG [64] has been proposed. RSVG utilized the Vision Transformer [7] and BERT [5] as backbones, incorporating the Multi-Level Cross-Modal feature learning module to address multi-scale variations in aerial images. Remote Sensing Referring Image Segmentation (RRSIS) is also a nascent field owing to the aforementioned challenges. Yuan et al. [63] constructed the first RRSIS dataset and proposed a model that utilizes the deep and shallow feature interactions to enhance the multi-scale feature extraction. However, this model encounters limitations in handling more complex datasets. In an effort to address the existing gaps in

*Known as RefSegRS, as of November 17, 2023, this dataset is not yet publicly available.

| Dataset | Number of images | Image size | Spatial resolution | Attributes of expression | Mask generation |
|---------------|------------------|------------|--------------------|--------------------------|--------------------|
| RefSegRS [63] | 4420 | 512 × 512 | 0.13m | 3 | Manually |
| RRSIS-D | 17402 | 800 × 800 | 0.5m ~ 30m | 7 | Semi automatically |

Table 1. Compare our dataset with the previous dataset.

4. RMSIN

4.1. Overview

The pipeline of our proposed model is depicted in Fig. 5. Initially, given the input image $I \in \mathbb{R}^{H \times W \times 3}$ and the language expression $E = \{\omega_i\}, i \in \{0, \dots, N\}$, where H and W represent the height and width of the input image, and N is the length of the expression, the input language expression E is transformed into the feature space $F_\ell \in \mathbb{R}^{N \times C}$ via the backbone f_ℓ . The following Compounded Scale Interaction Encoder (CSIE), which is composed of an Intra-scale Interaction Module (IIM) at each stage, and a Cross-scale Interaction Module (CIM), is applied to generate the fused features with sufficient semantics across multiple scales. Finally, we propose an Adaptive Rotated Convolution (ARC) based Oriented-Aware Decoder (OAD) to generate the segmentation mask by the parallel inference on the features from the multiple stages of the CSIE.

4.2. Compounded Scale Interaction Encoder

To effectively locate diverse targets with the guidance of the referred texts, the information for multi-scale is just as important as the referring expressions. Given the language features F_ℓ and the input image $I \in \mathbb{R}^{H \times W \times 3}$, the Compounded Scale Interaction Encoder (CSIE) brings about the fusion across vision-language modality in a multi-stage way with both intra- and inter-perspective.

Specifically, the CSIE is constructed with two components, Intra-scale Interaction Module (IIM) and Cross-scale Interaction Module (CIM).

4.2.1 Intra-scale Interaction Module

The first part of each stage in CSIE, the Intra-scale Interaction Module (IIM) is designed to further excavate the rich information within each scale and facilitate interaction between the vision and language modalities. Based on a hierarchy of four stages, IIM could be denoted as $\{\phi_i\}_{i \in \{1,2,3,4\}}$. After obtaining the language features $F_\ell \in \mathbb{R}^{N \times C}$ through the text backbone, where C denotes the number of channels, the output features F_e^i of IIM at stage i could be described as:

$$F_e^i = \phi_i(F_e^{i-1}, F_\ell), \quad (1)$$

where F_e^0 is extracted from the vision backbone f_v with the input I . More detailed, during the stage i , the input

features F_e^{i-1} undergo a combination of downsampling and MLP [31] to reduce the scale and unify the dimension of features, resulting in \hat{F}_e^{i-1} . The downsampled input \hat{F}_e^{i-1} is fed into two branches for enhancing visual priors and fusing cross-modal information individually.

Various Receptive Branch is the first branch. The feature \hat{F}_e^{i-1} is transformed through multiple branches with different settings of convolution kernels to yield features with various receptive fields, which could be formulated as:

$$\omega^i = \sigma \left(\sum_{j=0}^J \left(\frac{1}{C} \sum k_j^i * \hat{F}_e^{i-1} \right) \right), \quad (2)$$

where the k_j^i means the j -th branch of convolution and the σ is the Sigmoid Function. The above formulation indicates that the different convolution setting is utilized to balance the weight $\omega^i \in (0, 1)^{H \times W}$ between all the pixels. The weight is taken to enhance the features by:

$$\hat{F}_{e1}^{i-1} = \omega^i \otimes \hat{F}_e^{i-1}. \quad (3)$$

In addition, the output is regulated by a Visual Gate, adding to the raw image features as a complement to local detail information. The specific implementation of the gate is:

$$\alpha = \text{Tanh}(\text{Linear}(\text{ReLU}(\text{Linear}(\hat{F}_{e1}^{i-1}))))), \quad (4)$$

where $\text{Linear}(\cdot)$ denotes the linear projection, and $\text{Tanh}(\cdot)$ and $\text{ReLU}(\cdot)$ represent the activation functions.

Cross-modal Alignment Branch is designed for multi-modal alignment, which is the key to enabling the model to comprehend natural language.

Concretely, taking the input \hat{F}_e^{i-1} and language features F_ℓ , the module first implements scaled dot-product attention [46] using \hat{F}_e^{i-1} as the query and F_ℓ as the key and value to obtain the multi-modal features:

$$A^i = \text{attention}(\hat{F}_e^{i-1} W_q^i, F_\ell W_k^i, F_\ell W_v^i), \quad (5)$$

where W_q^i , W_k^i and W_v^i are the linear projection matrices. Subsequently, the attention A^i is combined with \hat{F}_e^{i-1} to obtain language-guided image features:

$$\hat{F}_{e2}^{i-1} = \text{Proj}(A^i W_w^i \otimes \hat{F}_e^{i-1} W_m^i), \quad (6)$$

where W_w^i and W_m^i are the projection matrices, and \otimes denotes element-wise multiplication. The obtained result is passed through a final 1×1 convolution, denoted as $\text{Proj}(\cdot)$, to produce the final output.

Similar to the operation performed on the output of \hat{F}_{e1}^{i-1} , the result is regulated by β from the Linguistic Gate and is added to the raw image features, serving as supplementary linguistic features. Consequently, the overall output features of IIM at stage i can be illustrated as:

$$F_e^i = \hat{F}_e^{i-1} + \alpha \hat{F}_{e1}^{i-1} + \beta \hat{F}_{e2}^{i-1}. \quad (7)$$

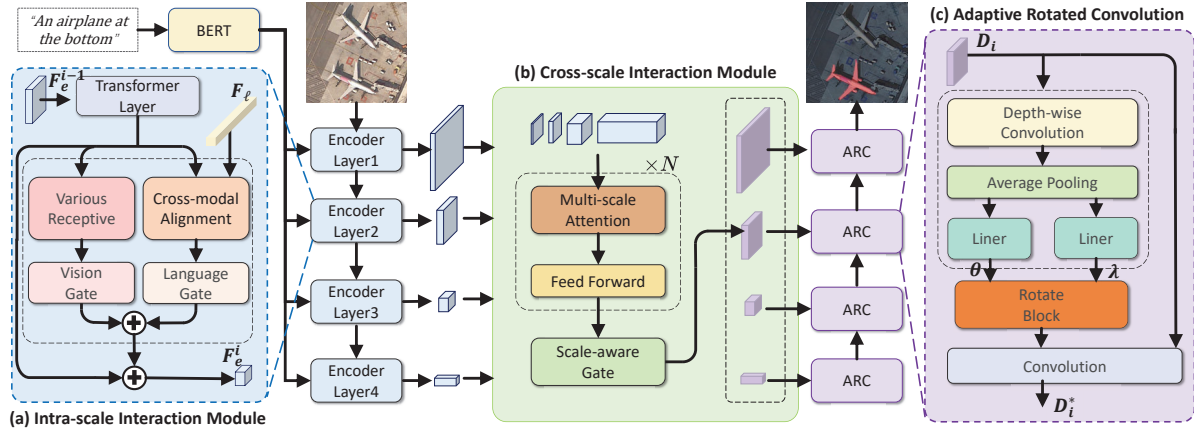


Figure 5. Overview of the proposed RMSIN.

4.2.2 Cross-scale Interaction Module

While the IIM adequately extracts localized multi-scale information guided by linguistic features, we additionally design a Cross-scale Interaction Module (CIM) to further enhance the interaction between the coarse and fine stages, particularly in response to the scale variation challenge observed in aerial images. Specifically, the module takes features collected from each stage of the IIM, *i.e.*, the previously mentioned F_e^i , $i \in \{1, 2, 3, 4\}$ as input and performs multi-stage interaction. The structure is depicted schematically in Fig. 5.

Multi-stage Feature Combination is first performed, where the features F_e^i are downsampled to the same size and concatenated along the channel dimension. The formula expression is as follows:

$$\begin{aligned} F_d^i &= \text{downsample}(F_e^i), \quad i \in \{1, 2, 3, 4\}, \\ F_c^* &= \text{concat}_c(F_d^1, F_d^2, F_d^3, F_d^4), \end{aligned} \quad (8)$$

where F_d^i represents the downsampled features, and F_c^* represents the multi-stage feature concatenated along the channel dimension. $\text{downsample}(\cdot)$ is typically implemented through average pooling.

Multi-scale Attention Layer is subsequently implemented. Specifically, we design various perceptive fields for the concatenated feature F_c^* to achieve deep multi-scale interaction. F_c^* is resized to different scales through the depth-wise convolutions with diverse kernel sizes and strides, defined as follows:

$$\begin{aligned} F_c^m &= k^m * F_c^*, \\ h^m &= \lfloor \frac{h-1}{m} + 1 \rfloor, w^m = \lfloor \frac{w-1}{m} + 1 \rfloor, \end{aligned} \quad (9)$$

where $m \in \{1, \dots, M\}$, M is the number of resized scales, k_m is the m -th depth-wise convolution and h_m and w_m are the corresponding height and weight of the F_c^m . With the

set $\{F_c^m | m \in \{1, \dots, M\}\}$, we flatten all the elements on the size dimension and concatenate them as a sequence features $\hat{F}_c^* \in \mathbb{R}^{(\sum_1^M h^m \times w^m) \times C}$. Similar to vanilla attention [46], we take the origin feature F_c^* as the query, and the multi-scale-aware feature \hat{F}_c^* as the key and value to perform cross-scale interaction:

$$\tilde{F}_c^* = \text{softmax}\left(\frac{F_c^* W_q \cdot \hat{F}_c^* W_k^T}{\sqrt{C}}\right) \cdot \hat{F}_c^* W_v. \quad (10)$$

For better preservation of local details, following inspiration from HRViT [10], a local relationship compensation called LRC is incorporated to regulate the output of the multi-scale attention. Consequently, the final output of the Multi-scale Attention Layer is expressed as:

$$F_c = \tilde{F}_c^* + \text{DWConv}(\text{Hardswish}(F_c^*)), \quad (11)$$

where $\text{DWConv}(\cdot)$ represents depth-wise convolution, and $\text{Hardswish}(\cdot)$ is the activity function, implemented in accordance with [10] to enhance the extraction of multi-scale local information.

The Feed Forward Layer follows the Multi-scale Attention layer which is identical to the standard attention block [46]. The feature F_c is divided into four parts to revert to the size of F_e^i by upsampling and subsequently fed into the Scale-aware Gate to obtain the final output.

Scale-aware Gate is employed to alleviate the semantic gap before and after multi-scale attention. Specifically, for each part from F_c , we implement the corresponding part from F_e to measure the weight of the cross-scale interaction. This weight is considered as the assistance residual for the features from IIM. The formulation is as follows:

$$F_o^i = \text{sigmoid}(F_e^i W_1) \otimes F_c^i W_2 + F_e^i W_3, \quad (12)$$

where $i \in \{1, 2, 3, 4\}$. The output of the Scale-aware Gate is utilized in the decoder for final mask prediction.

4.3. Oriented-aware Decoder

The set of features $\{F_o^i | i \in \{1, 2, 3, 4\}\}$ from the CSIE are used to generate the segmentation. Considering that object instances in aerial images often exhibit various orientations, using static horizontal convolution kernels for mask generation may result in a loss of precision. Inspired by oriented object detection, where the problem has been researched for decades and achieved considerable progress [43, 53, 55, 56], we incorporate the Adaptive Rotated Convolution (ARC) into the segmentation decoder tailored for the specific needs of RRSIS task to achieve better mask prediction.

4.3.1 Adaptive Rotated Convolution

The ARC captures angle information from input features and dynamically re-parameterizes the kernel weights to filter out redundant features. Specifically, it extracts orientation features and predicts n angles $\theta \in \{1, \dots, n\}$ and corresponding weights $\lambda \in \{1, \dots, n\}$ based on the input. For the input X , the θ and λ are predicted as:

$$\theta, \lambda = \text{Routing}(X), \quad (13)$$

where the concrete structure of the Routing Block is illustrated in Fig. 5. The static convolution kernel weights can be viewed as specific sampling points from the two-dimensional kernel space. Thus, the rotation of the convolution kernel is the process of rotary resampling. Specifically, the convolution kernel weights W_i are re-parameterized according to the predicted angles as follows:

$$\begin{aligned} Y_i' &= M^{-1}(\theta_i)Y_i, \\ W_i' &= \text{interpolation}(W_i, Y_i'), \end{aligned} \quad (14)$$

where Y_i is the coordinates of original sampling points, $M^{-1}(\theta_i)$ is the inverse matrix of the rotation matrix for affine transformation by angle θ around the origin, and $\text{interpolation}(\cdot)$ is implemented as bilinear interpolation. Finally, the features are filtered by the obtained convolution kernel and subjected to a weighted sum operation to produce orientation-aware features:

$$X^* = X * \sum_{i=1}^n \lambda_i W_i'. \quad (15)$$

The overall top-down process of mask prediction can be concluded as follows:

$$\begin{aligned} D_4 &= F_o^4, \\ D_i &= \text{Seg}(\text{ARC}([D_{i+1}; F_o^i])), \quad i \in \{1, 2, 3\}, \\ D_0 &= \text{Proj}(D_1), \end{aligned} \quad (16)$$

where $\text{Seg}(\cdot)$ refers to a nonlinear block comprising a 3×3 convolution layer, a batch normalization layer, and a ReLU

activity function to enhance the nonlinearity of the segmentation feature space. And $\text{Proj}(\cdot)$ is implemented as a linear projection function to map the final feature D_1 into two class scores. It is notable that half of the convolution layers are replaced by the ARC to leverage orientation information in the feature space, thereby eliminating redundancy for enhanced accuracy in boundary details.

5. Experiments

5.1. Implementation Details

Experiment Settings. In our experiments, the visual backbone utilizes Swin Transformer [31], pre-trained on ImageNet22K [4], while the language backbone employs the base BERT model from HuggingFace’s library [49]. The model is trained for 40 epochs using AdamW [32] with a weight decay of 0.01 and a starting learning rate of $3e-5$, reducing according to polynomial decay. The setup ran on four RTX 2080 GPUs with a batch size of 8.

Metrics. We utilize Overall Intersection-over-Union (oIoU), Mean Intersection-over-Union (mIoU), and Precision@X (P@X) as evaluation metrics, similar to prior studies [50, 64].

5.2. Comparison with state-of-the-art RIS methods

In our experiments, we compared RMSIN’s performance with existing state-of-the-art referring image segmentation methods on the validation and test subsets of our RRSIS-D dataset (see Tab. 2). For a fair comparison, we adopted the original implementation details of these competing methods. Notably, RMSIN outperforms its counterparts across almost all metrics on both subsets, marking a significant improvement with a 3.64% and 3.16% increase in mIoU on the validation and test subsets respectively over the closest competitor, LAVT. This leap in performance is particularly evident in complex scenarios, such as detecting small or rotated objects, where it secured over 3.0% gains in Precision at IoU thresholds of 0.5, 0.6, and 0.7.

5.3. Ablation study

We have performed various ablation experiments on the validation subset of RRSIS-D to assess the efficacy of the pivotal components within our proposed network.

Effectiveness of IIM and CIM. To validate the efficacy of our proposed two-scale interaction modules in CSIE, we conduct ablation studies on all the combinations of IIM and CIM. As illustrated in Tab. 3, The introduction of the IIM brings about discernible improvements in precision at lower IoU thresholds, while the incorporation of the CIM further refines predictions across various IoU levels. The combined effect of both modules demonstrates a synergistic enhancement, yielding the highest performance across all evaluated metrics, particularly in P@0.5, P@0.7, and mIoU, with

| Method | Visual Encoder | Text Encoder | P@0.5 | | P@0.6 | | P@0.7 | | P@0.8 | | P@0.9 | | oIoU | | mIoU | |
|---------------------|----------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | | | Val | Test | Val | Test | Val | Test | Val | Test | Val | Test | Val | Test | Val | Test |
| RRN [25] | R-101 | LSTM | 51.09 | 51.07 | 42.47 | 42.11 | 33.04 | 32.77 | 20.80 | 21.57 | 6.14 | 6.37 | 66.53 | 66.43 | 46.06 | 45.64 |
| CSMA [61] | R-101 | None | 55.68 | 55.32 | 48.04 | 46.45 | 38.27 | 37.43 | 26.55 | 25.39 | 9.02 | 8.15 | 69.68 | 69.39 | 48.85 | 48.54 |
| LSCM [17] | R-101 | LSTM | 57.12 | 56.02 | 48.04 | 46.25 | 37.87 | 37.70 | 26.37 | 25.28 | 7.93 | 8.27 | 69.28 | 69.05 | 50.36 | 49.92 |
| CMPC [16] | R-101 | LSTM | 57.93 | 55.83 | 48.85 | 47.40 | 38.50 | 36.94 | 25.28 | 25.45 | 9.31 | 9.19 | 70.15 | 69.22 | 50.41 | 49.24 |
| BRINet [15] | R-101 | LSTM | 58.79 | 56.90 | 49.54 | 48.77 | 39.65 | 39.12 | 28.21 | 27.03 | 9.19 | 8.73 | 70.73 | 69.88 | 51.14 | 49.65 |
| CMPC+ [30] | R-101 | LSTM | 59.19 | 57.65 | 49.36 | 47.51 | 38.67 | 36.97 | 25.91 | 24.33 | 8.16 | 7.78 | 70.14 | 68.64 | 51.41 | 50.24 |
| LGCE [63] | Swin-B | BERT | 68.10 | 67.65 | 60.52 | 61.53 | 52.24 | 51.45 | 42.24 | 39.62 | 23.85 | 23.33 | 76.68 | 76.34 | 60.16 | 59.37 |
| LAVT [58] | Swin-B | BERT | 69.54 | 69.52 | 63.51 | 63.63 | 53.16 | 53.29 | 43.97 | 41.60 | 24.25 | 24.94 | 77.59 | 77.19 | 61.46 | 61.04 |
| RMSIN (Ours) | Swin-B | BERT | 74.66 | 74.26 | 68.22 | 67.25 | 57.41 | 55.93 | 45.29 | 42.55 | 24.43 | 24.53 | 78.27 | 77.79 | 65.10 | 64.20 |

Table 2. Comparison with state-of-the-art methods on the proposed RRSIS-D dataset. R-101 and Swin-B represent ResNet-101 [11] and base Swin Transformer [31] models, respectively. The best result is bold.

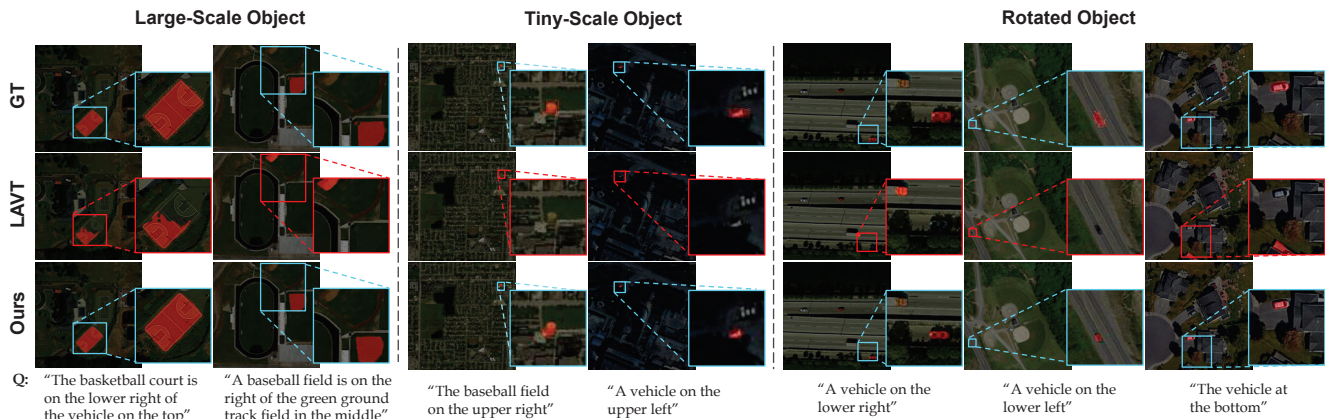


Figure 6. Qualitative comparisons between RMSIN and the previous SOTA LAVT. The left part illustrates the predictions of large-scale objects, while the middle part offers exceedingly diminutive objects amidst a highly noisy background. The right part exhibits the predictions for scenarios wherein objects are situated at diverse angles.

| IIM | CIM | P@0.5 | P@0.7 | P@0.9 | oIoU | mIoU |
|-----|-----|--------------|--------------|--------------|--------------|--------------|
| ✗ | ✗ | 69.54 | 53.16 | 24.25 | 77.59 | 61.46 |
| ✓ | ✗ | 71.09 | 53.45 | 24.71 | 77.68 | 62.27 |
| ✗ | ✓ | 73.68 | 56.67 | 25.69 | 77.40 | 64.25 |
| ✓ | ✓ | 74.14 | 57.59 | 25.69 | 77.91 | 64.91 |

Table 3. Ablation on the scale interaction modules IIM and CIM.

margins ranging from 3.5% to 4.5%. These findings affirm the pivotal role played by the IIM and CIM in capturing multi-scale features from images, thus substantiating their efficacy in advancing the overall segmentation capabilities. **Design options of CIM.** To further substantiate the effectiveness of CIM, we conduct a detailed analysis of its main components, as outlined in Tab. 4. The most substantial enhancement in results is observed upon the inclusion of the complete module, showcasing the highest metric enhancement of over 4.14%. This confirms the role of CIM in preserving local details and extracting multi-scale information. **Design options of Decoder.** We explore the design of the segmentation decoder structure as demonstrated in Tab. 5. The CIM yields output features with robust semantics and

| Options | P@0.5 | P@0.7 | P@0.9 | oIoU | mIoU |
|-------------------------|--------------|--------------|--------------|--------------|--------------|
| Default | 69.54 | 53.16 | 24.25 | 77.59 | 61.46 |
| + Multi-scale Attention | 68.91 | 53.68 | 25.11 | 77.61 | 61.46 |
| + Feed Forward | 69.83 | 52.70 | 25.57 | 77.85 | 61.46 |
| + LRC | 73.68 | 56.67 | 25.69 | 77.40 | 64.25 |

Table 4. Ablation on options design of CIM. Default means the vanilla self attention and we reintroduce all the designs cumulatively to demonstrate the effectiveness of each major component.

intricate spatial details. Thus our proposed Oriented-aware Decoder straightforwardly concatenates the features and extracts angular information through ARC to obtain more accurate predictions better suited to RS tasks. We also experiment with two alternative decoder structures. The exceptional results of our proposed decoder, surpassing others across all metrics, underscore the significance of incorporating angle information in the decoding process. This outcome firmly reaffirms the efficacy of our approach in customizing mask predictions for remote sensing applications, where the inclusion of precise angular information emerges as a critical factor for optimizing segmentation accuracy.

| Decoder Design | P@0.5 | P@0.7 | P@0.9 | oIoU | mIoU |
|----------------|--------------|--------------|--------------|--------------|--------------|
| Sum | 72.36 | 55.23 | 24.89 | 76.11 | 62.76 |
| Concat | 70.34 | 53.33 | 24.66 | 77.74 | 61.56 |
| Oriented-aware | 73.85 | 56.84 | 25.40 | 77.76 | 64.15 |

Table 5. Ablation study examining decoder designs. ‘‘Sum’’ decoder employs summation instead of concatenation for cross-stage feature connection, while ‘‘Concat’’ decoder substitutes ARC in ‘‘Oriented-aware’’ decoder with static convolutions.

| Options | P@0.5 | P@0.7 | P@0.9 | oIoU | mIoU |
|------------------------------------------------------|--------------|--------------|--------------|--------------|--------------|
| (a) Replacement of ARC in the Oriented-aware Decoder | | | | | |
| L = 1 | 71.72 | 54.94 | 24.43 | 75.93 | 62.53 |
| L = 2 | 72.18 | 55.29 | 24.66 | 76.37 | 62.48 |
| L = 3 | 72.36 | 55.98 | 25.00 | 77.06 | 63.81 |
| (b) Predicted number of angles in ARC | | | | | |
| n = 1 | 72.36 | 55.98 | 25.00 | 77.06 | 63.81 |
| n = 2 | 72.24 | 54.48 | 24.02 | 77.39 | 63.14 |
| n = 4 | 73.85 | 56.84 | 25.40 | 77.76 | 64.15 |

Table 6. Ablation studies of ARC. L=1 indicates the replacement of the first layer in the decoder with Adaptive Rotated Convolution. Experiments on the predicted number of angles are performed under L=3.

Design options of ARC. We further investigate the impact of the Adaptive Rotated Convolution (ARC) replacement strategy on the results, as demonstrated in Tab. 6 (a). We progressively replace the convolution layers in each stage of the decoder, and the result exhibits a consistent upward trend. Consequently, we opt to replace all three layers of the decoder. Additionally, we explore the influence of varying the number of prediction angles for ARC on the prediction results illustrated in Tab. 6 (b). The decoder showcases a consistent improvement in performance with an increase in the predicted number of angles, resulting in a performance boost of approximately 1% when seated to 4 compared to 1.

5.4. Visualization

5.4.1 Quantitative Results

We qualitatively compare our model with the baseline to provide a comprehensive understanding. As shown in Fig. 6, our model excels at identifying targets across various scales accurately, even within noisy backgrounds and at different angles. In contrast, the baseline model exhibits shortcomings such as missing parts and noticeable shifts in predicted masks.

5.4.2 Visualization of Features from Encoder

In Fig. 7, we visualize the feature maps from the RMSIN during training under the ablation of ARC and CSIE. It’s obvious that RMSIN can accurately capture boundary infor-




Figure 7. Visualization of predictions and feature maps across stages, where F_0^i denotes the feature map for stage i . Each row shows the outcomes from progressively adding modules.

mation with the assistance of scale interaction and rotated convolution. With the scale interaction performed by CSIE and the orientation extraction performed by ARC, RMSIN can focus more keenly on the referred targets. Compared with the first row, CSIE provides more accurate semantics in the deeper layer, and ARC supplies the space prior, which is important for rotated object segmentation.

These qualitative comparisons underscore the efficacy of our approach in addressing challenges related to scale variations and orientation robustness, affirming its capabilities in diverse scenarios.

6. Conclusion

In this paper, we propose RMSIN, a novel method adept at navigating the complex scales and orientations prevalent in RRSIS. By integrating the IIM and CIM, RMSIN capably addresses the wide range of spatial scales encountered in aerial imagery. Additionally, the implementation of the ARC offers a solid strategy for tackling various orientations. The construction of our expansive RRSIS-D dataset, featuring 17,402 image-caption-mask triplets provides an unparalleled resource in terms of scale and variety for rigorous evaluation. The comprehensive validation on the RRSIS-D dataset not only underscores RMSIN’s superior performance but also establishes a new benchmark for future research in this domain.

Acknowledgement

This work was supported by National Key R&D Program of China (No.2023YFB4502804), the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No. U21B2037, No. U22B2051, No. 62072389), the National Natural Science Fund for Young Scholars of China (No. 62302411), China Postdoctoral Science Foundation (No. 2023M732948), the Natural Science Foundation of Fujian Province of China (No.2021J01002, No.2022J06001), and partially sponsored by CCF-NetEase ThunderFire Innovation Research Funding (NO. CCF-Netease 202301).

References

- [1] Zhaowei Cai, Gukyeong Kwon, Avinash Ravichandran, Erhan Bas, Zhuowen Tu, Rahul Bhotika, and Stefan O Soatto. X-detr: A versatile architecture for instance-wise vision-language tasks. *ArXiv*, abs/2204.05626, 2022. 2
- [2] Ding-Jie Chen, Songhao Jia, Yi-Chen Lo, Hwann-Tzong Chen, and Tyng-Luh Liu. See-through-text grouping for referring image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 2
- [3] Yi-Wen Chen, Yi-Hsuan Tsai, Tiantian Wang, Yen-Yu Lin, and Ming-Hsuan Yang. Referring expression object segmentation with caption-aware consistency. In *British Machine Vision Conference (BMVC)*, 2019. 2
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 6
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. 2
- [6] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 2
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 2
- [8] Liuyun Duan and Florent Lafarge. Towards large-scale city reconstruction from satellites. In *European Conference on Computer Vision*, 2016. 1
- [9] G.M. Foody. Remote sensing of tropical forest environments: Towards the monitoring of environmental resources for sustainable development. *International Journal of Remote Sensing*, 24:4035–4046, 2003. 1
- [10] Jiaqi Gu, Hyoukjun Kwon, Dilin Wang, Wei Ye, Meng Li, Yu-Hsin Chen, Liangzhen Lai, Vikas Chandra, and David Z. Pan. Multi-scale high-resolution vision transformer for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12094–12103, 2022. 5
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [12] Richang Hong, Daqing Liu, Xiaoyu Mo, Xiangnan He, and Hanwang Zhang. Learning to compose and reason with language tree structures for visual grounding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44:684–696, 2019. 2
- [13] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 1, 2
- [14] Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. Modeling relationships in referential expressions with compositional modular networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [15] Zhiwei Hu, Guang Feng, Jiayu Sun, Lihe Zhang, and Huchuan Lu. Bi-directional relationship inferring network for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7
- [16] Shaofei Huang, Tianrui Hui, Si Liu, Guanbin Li, Yunchao Wei, Jizhong Han, Luoqi Liu, and Bo Li. Referring image segmentation via cross-modal progressive comprehension. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 7
- [17] Tianrui Hui, Si Liu, Shaofei Huang, Guanbin Li, Sansi Yu, Faxi Zhang, and Jizhong Han. Linguistic structure guided context modeling for referring image segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 59–75, 2020. 2, 7
- [18] Jiayi Ji, Yiwei Ma, Xiaoshuai Sun, Yiyi Zhou, Yongjian Wu, and Rongrong Ji. Knowing what to learn: a metric-oriented focal mechanism for image captioning. *IEEE Transactions on Image Processing*, 31:4321–4335, 2022. 1
- [19] Ya Jing, Tao Kong, Wei Wang, Liang Wang, Lei Li, and Tieniu Tan. Locate then segment: A strong pipeline for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9858–9867, 2021. 2
- [20] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr - modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1780–1790, 2021. 2
- [21] Dongwon Kim, Namyup Kim, Cuiling Lan, and Suha Kwak. Shatter and gather: Learning referring image segmentation with text supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15547–15557, 2023. 2
- [22] Namyup Kim, Dongwon Kim, Cuiling Lan, Wenjun Zeng, and Suha Kwak. Restr: Convolution-free referring image segmentation using transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18145–18154, 2022. 2
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything, 2023. 2, 3
- [24] Muchen Li and Leonid Sigal. Referring transformer: A one-step approach to multi-task visual grounding. *ArXiv*, abs/2106.03089, 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 *Proceed-*

- ings of the *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 7
- [26] Yue Liao, Si Liu, Guanbin Li, Fei Wang, Yanjie Chen, Chen Qian, and Bo Li. A real-time cross-modality correlation filtering method for referring expression comprehension. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [27] 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 3
- [28] Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. Recurrent multimodal interaction for referring image segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2
- [29] Chang Liu, Henghui Ding, Yulun Zhang, and Xudong Jiang. Multi-modal mutual attention and iterative interaction for referring image segmentation. *IEEE Transactions on Image Processing*, 32:3054–3065, 2023. 2
- [30] Si Liu, Tianrui Hui, Shaofei Huang, Yunchao Wei, Bo Li, and Guanbin Li. Cross-modal progressive comprehension for referring segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44:4761–4775, 2021. 7
- [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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 4, 6, 7
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2017. 6
- [33] 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 *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 10034–10043, 2020. 2
- [34] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Yan Wang, Liujuan Cao, Yongjian Wu, Feiyue Huang, and Rongrong Ji. Towards lightweight transformer via group-wise transformation for vision-and-language tasks. *IEEE Transactions on Image Processing*, 31:3386–3398, 2022.
- [35] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Yongjian Wu, Yue Gao, and Rongrong Ji. Towards language-guided visual recognition via dynamic convolutions. *International Journal of Computer Vision*, pages 1–19, 2023. 2
- [36] Yunpeng Luo, Jiayi Ji, Xiaoshuai Sun, Liujuan Cao, Yongjian Wu, Feiyue Huang, Chia-Wen Lin, and Rongrong Ji. Dual-level collaborative transformer for image captioning. In *Proceedings of the AAAI conference on artificial intelligence*, pages 2286–2293, 2021. 1
- [37] Yiwei Ma, Guohai Xu, Xiaoshuai Sun, Ming Yan, Ji Zhang, and Rongrong Ji. X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 638–647, 2022.
- [38] Yiwei Ma, Jiayi Ji, Xiaoshuai Sun, Yiyi Zhou, and Rongrong Ji. Towards local visual modeling for image captioning. *Pattern Recognition*, 138:109420, 2023. 1
- [39] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Computer Vision and Pattern Recognition*, 2016. 2
- [40] Edgar Margffoy-Tuay, Juan C. Perez, Emilio Botero, and Pablo Arbelaez. Dynamic multimodal instance segmentation guided by natural language queries. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [41] Varun K. Nagaraja, Vlad I. Morariu, and Larry S. Davis. Modeling context between objects for referring expression understanding. In *ECCV*, 2016. 2
- [42] Shuyi Ouyang, Hongyi Wang, Shiao Xie, Ziwei Niu, Ruofeng Tong, Yen-Wei Chen, and Lanfen Lin. Slvit: Scale-wise language-guided vision transformer for referring image segmentation. pages 1294–1302, 2023. 2
- [43] Yifan Pu, Yiru Wang, Zhuofan Xia, Yizeng Han, Yulin Wang, Weihao Gan, Zidong Wang, Shiji Song, and Gao Huang. Adaptive rotated convolution for rotated object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6589–6600, 2023. 6
- [44] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Sasha Luccioni, Tegan Maharaj, Evan D. Sherwin, Surya Karthik Mukkavilli, Konrad Paul Kording, Carla P. Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer T. Chayes, and Yoshua Bengio. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55:1 – 96, 2019. 1
- [45] Yuxi Sun, Shanshan Feng, Xutao Li, Yunming Ye, Jian Kang, and Xu Huang. Visual grounding in remote sensing images. *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. 2
- [46] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Information Processing Systems*, 2017. 4, 5
- [47] Haowei Wang, Jiayi Ji, Yiyi Zhou, Yongjian Wu, and Xiaoshuai Sun. Towards real-time panoptic narrative grounding by an end-to-end grounding network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2528–2536, 2023. 2
- [48] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. 2
- [49] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. Transformers: State-of-the-art natural language processing. In *Conference on Empirical Methods in Natural Language Processing*, 2019. 6

- [50] Chenyun Wu, Zhe Lin, Scott Cohen, Trung Bui, and Subhransu Maji. Phrasedcut: Language-based image segmentation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [6](#)
- [51] Changli Wu, Yiwei Ma, Qi Chen, Haowei Wang, Gen Luo, Jiayi Ji, and Xiaoshuai Sun. 3d-stmn: Dependency-driven superpoint-text matching network for end-to-end 3d referring expression segmentation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. [2](#)
- [52] Jianzong Wu, Xiangtai Li, Xia Li, Henghui Ding, Yu Tong, and Dacheng Tao. Towards robust referring image segmentation. *ArXiv*, abs/2209.09554, 2022. [2](#)
- [53] Yongchao Xu, Mingtao Fu, Qimeng Wang, Yukang Wang, Kai Chen, Gui-Song Xia, and Xiang Bai. Gliding vertex on the horizontal bounding box for multi-oriented object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(4):1452–1459, 2021. [6](#)
- [54] Danni Yang, Jiayi Ji, Xiaoshuai Sun, Haowei Wang, Yinan Li, Yiwei Ma, and Rongrong Ji. Semi-supervised panoptic narrative grounding. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7164–7174, 2023. [2](#)
- [55] Xue Yang, Xiaojiang Yang, Jirui Yang, Qi Ming, Wentao Wang, Qi Tian, and Junchi Yan. Learning high-precision bounding box for rotated object detection via kullback-leibler divergence. *ArXiv*, abs/2106.01883, 2021. [6](#)
- [56] Xue Yang, Yue Zhou, Gefan Zhang, Jitui Yang, Wentao Wang, Junchi Yan, Xiaopeng Zhang, and Qi Tian. The kfiou loss for rotated object detection. *ArXiv*, abs/2201.12558, 2022. [6](#)
- [57] Zhengyuan Yang, Boqing Gong, Liwei Wang, Wenbing Huang, Dong Yu, and Jiebo Luo. A fast and accurate one-stage approach to visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. [2](#)
- [58] Zhao Yang, Jiaqi Wang, Yansong Tang, Kai Chen, Hengshuang Zhao, and Philip H.S. Torr. Lavt: Language-aware vision transformer for referring image segmentation. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18134–18144, 2022. [1](#), [2](#), [7](#)
- [59] Alias-z yatengLG and horffmanwang. Isat with segment anything: Image segmentation annotation tool with segment anything, 2023. [3](#)
- [60] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10494–10503, 2019. [2](#)
- [61] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10502–10511, 2019. [2](#), [7](#)
- [62] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *ECCV*, pages 69–85, 2016. [1](#), [2](#)
- [63] Zhenghang Yuan, Lichao Mou, Yuansheng Hua, and Xiao Xiang Zhu. Rrsis: Referring remote sensing image segmentation, 2023. [2](#), [3](#), [4](#), [7](#)
- [64] Yang Zhan, Zhitong Xiong, and Yuan Yuan. Rsvg: Exploring data and models for visual grounding on remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13, 2023. [2](#), [3](#), [6](#)
- [65] Hanwang Zhang, Yulei Niu, and Shih-Fu Chang. Grounding referring expressions in images by variational context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [66] 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 *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXV*, pages 598–615. Springer, 2022. [2](#)
- [67] Bohan Zhuang, Qi Wu, Chunhua Shen, Ian Reid, and Anton van den Hengel. Parallel attention: A unified framework for visual object discovery through dialogs and queries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)