

RS2-SAM2: Customized SAM2 for Referring Remote Sensing Image Segmentation

Fu Rong¹, Meng Lan², Qian Zhang³, Lefei Zhang^{1,*}

¹National Engineering Research Center for Multimedia Software, School of Computer Science, Wuhan University

²Hong Kong University of Science and Technology

³Horizon Robotics

{furong, zhanglefei}@whu.edu.cn, eemenglan@ust.hk, qian01.zhang@horizon.auto

Abstract

Referring Remote Sensing Image Segmentation (RRSIS) aims to segment target objects in remote sensing (RS) images based on textual descriptions. Although Segment Anything Model 2 (SAM2) has shown remarkable performance in various segmentation tasks, its application to RRSIS presents several challenges, including understanding the text-described RS scenes and generating effective prompts from text. To address these issues, we propose **RS2-SAM2**, a novel framework that adapts SAM2 to RRSIS by aligning the adapted RS features and textual features while providing pseudo-mask-based dense prompts. Specifically, we employ a union encoder to jointly encode the visual and textual inputs, generating aligned visual and text embeddings as well as multimodal class tokens. A bidirectional hierarchical fusion module is introduced to adapt SAM2 to RS scenes and align adapted visual features with the visually enhanced text embeddings, improving the model’s interpretation of text-described RS scenes. To provide precise target cues for SAM2, we design a mask prompt generator, which takes the visual embeddings and class tokens as input and produces a pseudo-mask as the dense prompt of SAM2. Experimental results on several RRSIS benchmarks demonstrate that RS2-SAM2 achieves state-of-the-art performance.

Introduction

Referring Remote Sensing Image Segmentation (RRSIS) (Liu et al. 2024b; Yuan et al. 2024) aims to segment specified targets from aerial images based on textual descriptions. This task extends the capabilities of traditional Referring Image Segmentation (RIS) (Wang et al. 2022; Yang et al. 2022; Liu et al. 2023a) by addressing the unique challenges inherent to remote sensing images. These challenges include handling diverse spatial scales, interpreting complex scene contexts, and resolving ambiguous object boundaries, which are particularly prevalent in remote sensing scenes.

Recent advances in the Segment Anything Model (SAM) (Kirillov et al. 2023) and its variants (Zhang et al. 2023; Xiong et al. 2024; Ke et al. 2024) have demonstrated significant improvements in efficiency and accuracy for promptable segmentation tasks in natural images. These models

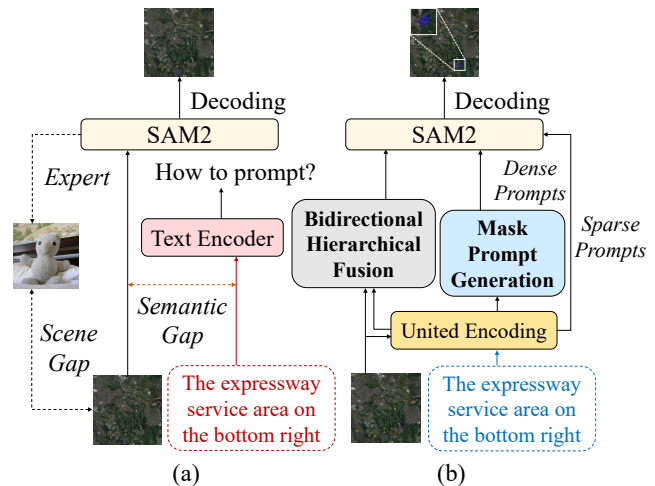


Figure 1: Comparison of two SAM2 adaptations for RRSIS. (a) vanilla SAM2, (b) our RS2-SAM2.

exhibit powerful segmentation capabilities and robust interactive prompting mechanisms. Notably, SAM2 (Ravi et al. 2024) incorporates the hierarchical encoder Hiera (Ryali et al. 2023), enhancing its ability to process images with diverse spatial scales. However, despite these advancements, applying SAM2 to RRSIS remains challenging due to the unique complexities of remote sensing images.

First, while SAM2 performs well on natural images, its effectiveness declines in remote sensing scenes due to the low target distinguishability and limited foreground-background contrast. Moreover, SAM2 struggles to exploit spatial information from textual descriptions, as shown in Fig. 1 (a). Although existing studies have attempted to address these issues, significant room for improvement remains. For example, SAM2-Adapter (Chen et al. 2024) modifies SAM2’s encoding process to better adapt to complex scenes. However, this unimodal approach lacks hierarchical information interaction, resulting in an insufficient fine-grained understanding of textual information. Consequently, achieving effective alignment of vision-language information and improving generalization to complex remote sensing scenes remain critical challenges in adapting SAM2 for RRSIS. **Second**, SAM2’s lack of tex-

*Corresponding author.

tual prompt integration limits its ability to generate prompts aligned with textual descriptions. Previous RRSIS methods, such as RMSIN (Liu et al. 2024b), rely on independent encoding for visual and textual inputs, but such traditional encoder-decoder structures fail to effectively integrate textual information into SAM2. EVF-SAM (Zhang et al. 2024) mitigates this by jointly encoding both modalities and generating sparse prompts via an MLP, achieving promising results on natural images. However, in remote sensing, sparse prompts alone cannot ensure fine-grained understanding, especially for subtle or indistinct objects. Designing effective text-based prompts to guide decoding therefore remains a key challenge for adapting SAM2 to RRSIS.

To address these challenges, this paper proposes RS2-SAM2, a novel RRSIS framework adapted from SAM2. As illustrated in Fig. 1 (b), our approach focuses on two key aspects: (1) *adapting SAM2 image features to remote sensing scenes and aligning them with textual features during the image encoding process*, and (2) *generating dense prompts for precise segmentation*. Specifically, to adapt SAM2 to remote sensing scenes and align remote sensing visual features with textual features, we design a bidirectional hierarchical fusion module, which is embedded both within and after the frozen SAM2 image encoder. Initially, we utilize a union encoder to jointly encode visual and textual inputs, producing semantically aligned visual and textual embeddings as well as multimodal class tokens. Subsequently, our proposed fusion module hierarchically aligns the visually enhanced textual embeddings with adapted remote sensing visual features in the SAM2 encoder.

To equip SAM2 with more precise prompts, we introduce a mask prompt generator. This module integrates jointly encoded visual embeddings and multimodal class tokens to generate a pseudo-mask representing remote sensing target objects. The pseudo-mask is then fed into the SAM2 prompt encoder as a dense prompt, providing pixel-level positional information to enhance the model’s segmentation capability.

Experimental results across several RRSIS benchmarks demonstrate the superior performance of our model and the effectiveness of the proposed modules. The primary contributions of this work are outlined as follows:

- We propose the RS2-SAM2 framework, which adapts the SAM2 model to the RRSIS task by hierarchically aligning the adapted remote sensing visual and textual features, providing pseudo-mask-based dense prompts, and enforcing boundary constraints.
- We design a bidirectional hierarchical fusion module that adapts the SAM2 encoder to the remote sensing images and aligns the adapted visual features with textual features during the encoding process.
- We develop a mask prompt generator that leverages multimodal features to produce a pseudo-mask as the dense prompt, offering precise target cues for SAM2.

Related Work

Referring Remote Sensing Image Segmentation. RRSIS emerges as an important research direction in remote sensing (Wang et al. 2025a,b), allowing users to extract geospa-

tial information through natural language queries (Yingrui Ji 2025). Compared to the remote sensing visual grounding (RSVG) (Sun et al. 2022; Zhan, Xiong, and Yuan 2023; Lan et al. 2024) task, RRSIS focuses on fine-grained pixel-level analysis rather than region-level identification. However, research in this area remains in its early stages with limited exploration. Yuan et al. (Yuan et al. 2024) pioneer this task by introducing the RefsegRS dataset and adapting the LAVT (Yang et al. 2022) framework for remote sensing. To address challenges like small and fragmented targets easily confused with complex backgrounds, they develop a multi-level feature fusion mechanism incorporating linguistic guidance, significantly enhancing small object detection. Subsequently, Liu et al. (Liu et al. 2024b) build upon the RSVG dataset DIOR-RSVG (Zhan, Xiong, and Yuan 2023) to establish the first large-scale RRSIS dataset, RRSIS-D. They further introduce the RMSIN architecture, which leverages rotational convolution operations to better handle spatial-scale variations and directional complexities in remote sensing images, providing a more robust modeling framework for RRSIS. Building upon RMSIN (Liu et al. 2024b), Lei et al. propose FIANet (Lei et al. 2024), which focuses more on fine-grained vision-language interaction and adaptive understanding of objects at different scales.

Segment Anything Model. SAM (Kirillov et al. 2023) is an interactive model designed to generate non-semantic segmentation masks based on a variety of input prompts. Leveraging large-scale training data, it demonstrates strong generalization across diverse segmentation tasks. SAM has also been widely utilized in multiple domains, including remote sensing image analysis (Wang et al. 2024b, 2025c, 2024a), video object tracking (Cheng et al. 2023; Yan et al. 2024), and medical image segmentation (Yue et al. 2024; Cheng et al. 2024). To improve its accuracy and efficiency, several optimized versions (Zhang et al. 2023; Xiong et al. 2024; Zhong et al. 2024) have been developed. More recently, SAM2 (Ravi et al. 2024) has been introduced for video segmentation (Rong et al. 2025; Lan et al. 2023), incorporating multi-scale feature encoding to enhance segmentation robustness. Despite these advances, SAM remains limited by its lack of linguistic understanding, preventing it from directly handling referring segmentation tasks that require textual guidance. To address this limitation, recent works (Lai et al. 2024; Xu et al. 2023; Rasheed et al. 2024) have explored integrating multimodal large language models (MLLMs) to enhance SAM’s ability to process language-based instructions. LISA (Lai et al. 2024), for example, fine-tunes LLaVA (Liu et al. 2024a) to derive hidden embeddings, thus producing multimodal features for enhanced segmentation. u-LLaVA (Xu et al. 2023) extends this paradigm by enabling simultaneous region- and pixel-level segmentation through multi-task learning. With its lightweight pre-fusion strategy, EVF-SAM (Zhang et al. 2024) leverages joint visual-language encoding to strengthen text-driven segmentation prompts, achieving superior segmentation accuracy. However, while these models demonstrate strong performance on natural images, their effectiveness in complex remote sensing scenes (Zhang et al. 2025) remains limited, underscoring the need for further refinement and adaptation.

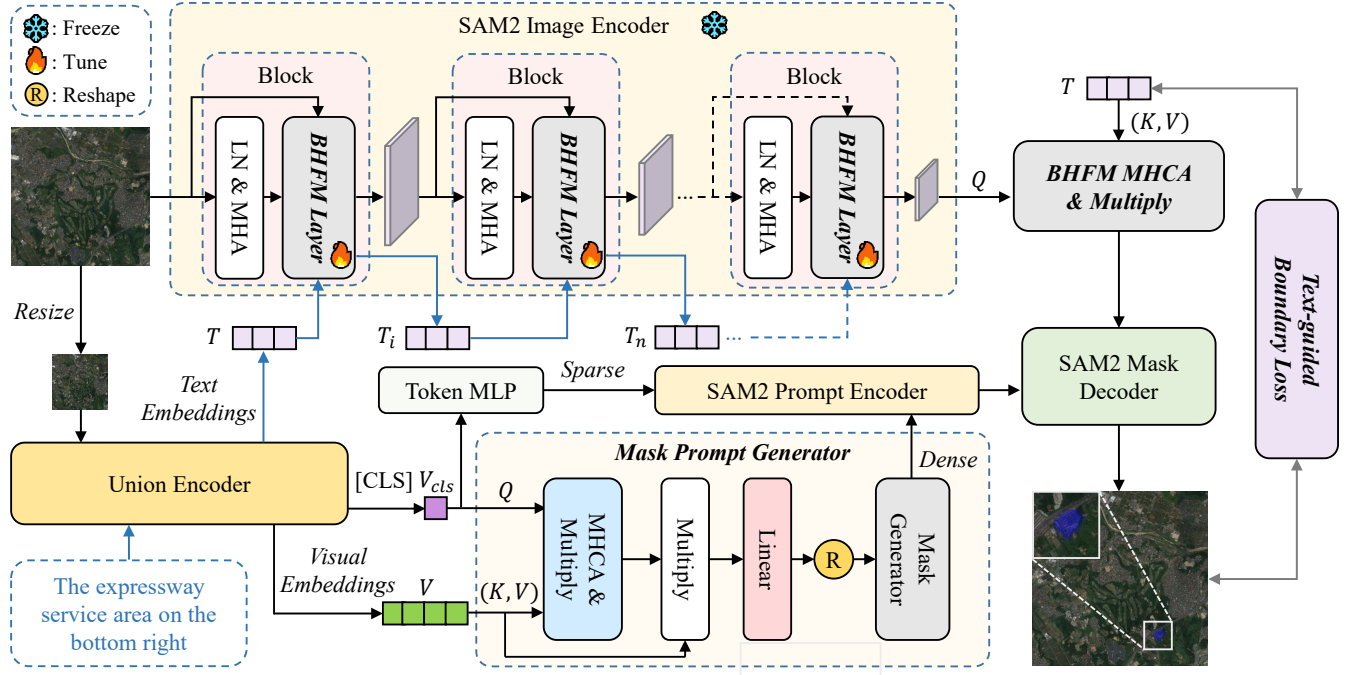


Figure 2: The overview of the proposed RS2-SAM2 framework. It consists of four key components: the union encoder, the bidirectional hierarchical fusion module, the mask prompt generator, and SAM2. The union encoder extracts multimodal representations from the input image and text. The bidirectional hierarchical fusion module enhances image features with textual embeddings. The mask prompt generator produces a prior mask as the dense prompt for SAM2. Finally, SAM2 generates precise masks, while the text-guided boundary loss constrains their boundary accuracy.

Method

Overview

The architecture of the proposed RS2-SAM2 framework is depicted in Fig. 2. RS2-SAM2 comprises four essential parts: the union encoder, the bidirectional hierarchical fusion module, the mask prompt generator, and the SAM2 model. Given an input remote sensing image \mathcal{I} and its associated textual description $\mathcal{E} = \{e_i\}_{i=1}^L$, where L is the number of words, the union encoder processes both modalities to extract a multimodal [CLS] token, visual patch embeddings, and textual embeddings. The SAM2 image encoder, incorporating the bidirectional hierarchical fusion module, further refines the extracted image features by leveraging textual information. Next, the mask prompt generator utilizes the visual patch embeddings and the multimodal [CLS] token to produce a pseudo-mask estimate for the target object. This prior mask, along with the multimodal [CLS] token, serves as a guiding signal for SAM2. Ultimately, the SAM2 decoder synthesizes the extracted image features and generated prompts to produce high-precision segmentation masks.

Feature Extraction

To align visual and textual modalities in RRSIS, this work uses the union BEiT-3 encoder (Wang et al. 2023), following (Zhang et al. 2024; Yu et al. 2024). Each input image $I_u \in \mathbb{R}^{H_u \times W_u \times 3}$ is divided into non-overlapping patches $P_v \in \mathbb{R}^{N_p \times (p^2 \times 3)}$ and then projected to $P_v \in \mathbb{R}^{N_p \times D}$,

where $N_p = \frac{H_u \times W_u}{p^2}$ and D is the embedding dimension. A visual class token $V_{cls} \in \mathbb{R}^{1 \times D}$ and positional embedding $V_{pos} \in \mathbb{R}^{(N_p+1) \times D}$ are prepended to obtain:

$$V_0 = [V_{cls}, P_v] + V_{pos}$$

Text of length L is tokenized using XLM-Roberta (Conneau 2019) as T_{seq} . A class token and end-of-sequence marker are added, followed by positional embedding: $T_0 = [T_{cls}, T_{seq}, T_{eos}] + T_{pos}$, where $T_0 \in \mathbb{R}^{N_t \times D}$ and $N_t = L + 2$.

The concatenated multimodal representation U_0 is constructed as:

$$U_0 = [V_0; T_0] \in \mathbb{R}^{(N_p+N_t+1) \times D}$$

After multimodal fusion and FFNs, the final representation U is obtained, which is decomposed into $V_{cls} \in \mathbb{R}^{1 \times D}$, $V \in \mathbb{R}^{N_p \times D}$, and $T \in \mathbb{R}^{N_t \times D}$.

The SAM2 image encoder with the proposed bidirectional hierarchical fusion module encodes the input image $I_s \in \mathbb{R}^{H_s \times W_s \times 3}$. It extracts multi-scale features, and the final-layer output $F_n \in \mathbb{R}^{\frac{H_s}{16} \times \frac{W_s}{16} \times C}$ is used for decoding, where H_s , W_s , and C denote the height, width, and channel number of the SAM2 image encoder, respectively.

Bidirectional Hierarchical Fusion Module

Although SAM2 demonstrates powerful segmentation capabilities for natural images, it struggles to achieve accurate segmentation for remote sensing images with more complex

scenes. Inspired by MSA (Wu et al. 2023), a straightforward approach is to incorporate linear layers into the SAM2 image encoder to enhance its adaptability to remote sensing images. However, for the RRSIS task, integrating textual information during SAM2’s image encoding process could make the model more sensitive to referred objects. To address this, we designed a bidirectional hierarchical fusion module and embedded it into SAM2 image encoder, enabling the visual feature of SAM2 to better adapt to remote sensing scenes and the referring text, thereby achieving more precise segmentation.

As illustrated in Fig. 3, our bidirectional hierarchical fusion module begins with the preprocessed SAM2 image feature F_i of the current layer being projected to a lower dimensionality through a linear layer and an activation function. Simultaneously, the text feature T_i of the current layer is also projected to match the dimensionality of the image feature using a linear layer. Subsequently, the image and text features are utilized as query embeddings and key-value pairs, respectively, and undergo cross-attention interaction to capture modality-specific dependencies. The resulting features are then integrated with the pre-interaction representations through element-wise addition. Following this, linear layers are used to restore the dimensionality of both the image and text features. The above process can be expressed by the following equation:

$$\begin{aligned} F'_i &= \sigma(\text{Linear}(F_i)), T'_i = \text{Linear}(T_i), \\ F''_i &= \text{MHCA}(F'_i, T'_i) + F'_i, \\ T''_i &= \text{MHCA}(T'_i, F'_i) + T'_i, \end{aligned} \quad (1)$$

where σ denotes the activation function GeLU and MHCA represents the multi-head cross-attention layer.

To preserve textual integrity during visual enhancement, the text feature is weighted and summed with its pre-interaction representation. Meanwhile, after skip-connection, the visual feature is summed with the original image feature F_{in} of the current layer, then separately processed by the MLP branch and the linear branch, followed by weighted fusion. The entire process is depicted as follows:

$$\begin{aligned} T_{i+1} &= (1 - \alpha_t)T_i + \alpha_t \text{Linear}(T''_i), \\ F'''_i &= F_{in} + \text{Linear}(F''_i) + F_i, \\ F_{out} &= F'''_i + \text{MLP}(\text{LN}(F'''_i)) + \\ &\quad \alpha_i \text{Linear}(\sigma(\text{Linear}(F'''_i))), \end{aligned} \quad (2)$$

here, α_t represents the text weighting coefficient, F_{out} denotes the image feature output to the next layer, and α_i represents the image weighting coefficient.

After feature encoding, the original text feature T encoded by the union encoder is used to further guide the visual feature F at the high-level. Specifically, the visual feature acts as query and interacts with the text feature as key-value pairs through multi-head cross-attention. The resulting feature is then element-wise multiplied with the visual feature before interaction, yielding the text-guided hierarchical feature F_{en} , which is subsequently fed into the SAM2 decoder for accurate decoding.

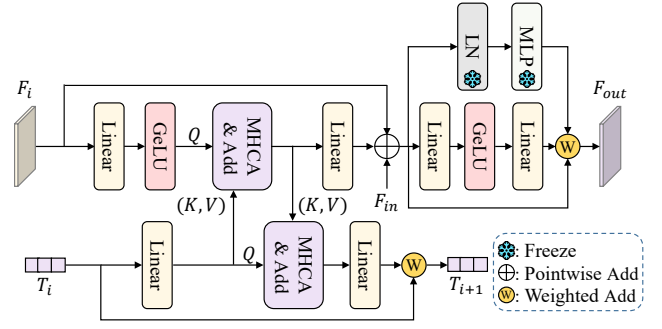


Figure 3: The structure of the bidirectional hierarchical fusion module.

Mask Prompt Generator

Although the SAM2 image feature has been hierarchically guided by text features using the bidirectional hierarchical fusion module, further providing pixel-level semantic guidance could offer finer-grained textual information. Given the excellent semantic alignment between the visual embeddings V and the textual embeddings T in the joint encoding process, we propose combining the multimodal [CLS] token with visually aligned embeddings that are well-aligned with text. This combination generates a multimodal pseudo-mask prior, which provides pixel-level guidance for the SAM2 decoding process.

As illustrated in Fig. 2, we first employ the multimodal [CLS] token V_{cls} as query and the visual embeddings V as key-value pairs for cross-attention computation. The interaction result is then element-wise multiplied with the multimodal [CLS] token to further align the multimodal token with visual information. Subsequently, the visual embeddings are reshaped into feature maps with dimensions $\frac{H_u}{p} \times \frac{W_u}{p}$, and the multimodal token is passed through a linear layer and broadcasted to match the shape of the visual embeddings. After element-wise multiplication, the result is input into the mask generator, composed of MLP layers, to produce the pseudo-mask $M_p \in \mathbb{R}^{\frac{H_u}{p} \times \frac{W_u}{p}}$ that is well-aligned with multimodal information.

The pseudo-mask M_p is subsequently upsampled to dimensions H_s and W_s using linear interpolation, to match the feature size of the SAM2 decoder. This upsampled pseudo-mask is subsequently fed as the dense prompt into the SAM2 prompt encoder.

SAM2 Prompt Encoder and Mask Decoder

The prompt encoder in RS2-SAM2 receives two inputs: a dense prompt M_p from the mask prompt generator and a sparse prompt V_{cls} from the union encoder. Following EVF-SAM (Zhang et al. 2024), V_{cls} is processed via a token MLP and combined with zero-initialized sparse embeddings. To ensure compatibility, the pixel-level prompt M_p is adjusted to the spatial dimensions of SAM2 features before being input as the dense spatial information into the mask decoder.

The SAM2 mask decoder leverages both sparse object-level and dense pixel-level prompts: the sparse forms queries

with object tokens, while the latter, as a mask prior, is added to visual features to enable direct high-resolution decoding.

Training Loss

RS2-SAM2 utilizes a comprehensive loss function akin to that proposed in (Liu et al. 2024b) to regulate the predicted mask, formulated as follows:

$$\mathcal{L} = \lambda_{ce}\mathcal{L}_{ce} + \lambda_{dice}\mathcal{L}_{dice} + \lambda_{tbl}\mathcal{L}_{tbl}, \quad (3)$$

where \mathcal{L}_{ce} corresponds to the cross-entropy loss, \mathcal{L}_{dice} denotes the DICE loss (Milletari, Navab, and Ahmadi 2016), and \mathcal{L}_{tbl} signifies the text-guided boundary loss, a loss function we designed that uses text-weighted constraints to predict mask boundaries.

Text-guided Boundary Loss. In the RRSIS task, unlike natural images, the target objects often exhibit low visual distinguishability from the background, making the boundaries of remote sensing targets less distinct. To address this, we introduce a text-guided boundary loss function L_{tbl} to constrain the predicted mask boundaries. Specifically, we first compute the difference between each pixel value and its neighboring pixel value in both horizontal and vertical directions to obtain the gradient, which serves as an indicator for boundary detection. Then, we abstract the text embeddings T into sentence embeddings T_s , reduce their dimensionality to a single scalar through a linear layer, and expand it to match the size of the mask as a text-guided weighting factor for the boundary gradient. Finally, we measure the boundary similarity between the predicted mask and the ground truth mask under the guidance of the text weights by MSE loss:

$$\begin{aligned} \nabla_{pre} &= Absd_h(M_{pre}) + Absd_v(M_{pre}) \\ \nabla_{gt} &= Absd_h(M_{gt}) + Absd_v(M_{gt}) \\ \mathcal{L}_{tbl} &= \frac{1}{N} \sum_{i=1}^N (\text{Linear}(T_s)(\nabla_{pre_i} - \nabla_{gt_i}))^2 \end{aligned} \quad (4)$$

where $Absd_h$ and $Absd_v$ represent the absolute differences between adjacent pixels in the horizontal and vertical directions, respectively, while ∇ denotes the gradient and N is the number of all pixels.

Experiments

Datasets and Metrics

Datasets. The experiments are conducted on two key RRSIS datasets: RefSegRS (Yuan et al. 2024) and RRSIS-D (Liu et al. 2024b). RefSegRS is the first remote sensing referring segmentation dataset, containing 2172 images for training, 413 for validation, and 1817 for testing, each with a resolution of 512×512. RRSIS-D is a widely used large-scale dataset in the RRSIS field, including 12181 training samples, 1740 validation samples, and 3481 test samples, with each image sized at 800×800.

Evaluation Metrics. In accordance with the evaluation protocol established in (Liu et al. 2024b), we assess our model using several metrics, including precision at various IoU thresholds (Pr@0.5 to Pr@0.9), mean Intersection-over-Union (mIoU), and overall Intersection-over-Union (oIoU). These metrics are computed on the validation and test sets of both RefSegRS and RRSIS-D datasets.

Implementation Details

Model Settings. We initialize the key modules of SAM2 and the union encoder using pre-trained weights from SAM2-Hiera-Large (Ravi et al. 2024) and BEiT-3-Large (Wang et al. 2023). For feature extraction, each image is resized to 1024 × 1024 and 224 × 224, which are then fed into the SAM2 image encoder with an output dimension of $C = 256$ and the union encoder with an output dimension of $D = 1024$. Unlike SAM2’s full capabilities, we do not employ its memory mechanism, focusing exclusively on image processing. In the bidirectional hierarchical fusion module, we set the text weight coefficient α_t to 0.2 and the image weight coefficient α_i to 0.5.

Training Details. Experiments are conducted on 8 NVIDIA GeForce RTX 4090 GPUs. The experiment adopts a setup similar to (Lei et al. 2024), with training conducted for 60 epochs on the RefSegRS (Yuan et al. 2024) dataset and 40 epochs on the RRSIS-D (Liu et al. 2024b) dataset. We use the AdamW optimizer (Loshchilov and Hutter 2018) with a unified batch size of 1. The learning rates for models on RefSegRS and RRSIS-D dataset are set to 5e-5 and 1e-5, respectively. To facilitate the adaptation of our designed modules to remote sensing data, the learning rates for bidirectional hierarchical fusion module and mask prompt generator are assigned as 1e-4 on RefSegRS dataset and 5e-5 on RRSIS-D dataset. All learning rates are reduced to 0.1 times their original values during the last 10 epochs. The weighting coefficients for various loss functions are defined as follows: $\lambda_{dice} = 0.1$, $\lambda_{ce} = 1$, and $\lambda_{tbl} = 0.2$.

Comparison with State-of-the-Art Methods

RefSegRS dataset. We conduct a comprehensive comparison between our method, RS2-SAM2, and several state-of-the-art approaches on the RefSegRS (Yuan et al. 2024) dataset. The results of the comparison are reported in Tab. 1. It can be observed that our RS2-SAM2 achieves 88.03% oIoU and 85.21% mIoU on the validation set, surpassing the previous best method, RMSIN, by 5.62% in oIoU and 11.37% in mIoU. On the test set, it reaches 80.87% oIoU and 73.90% mIoU, outperforming FIANet (Lei et al. 2024) by 2.55% and 5.23%, respectively. Moreover, the Pr metric exhibits significant improvements across all thresholds, particularly at Pr@0.7, Pr@0.8, and Pr@0.9, demonstrating the model’s strong multimodal segmentation capabilities.

RRSIS-D dataset. We also conduct comparative experiments between our RS2-SAM2 and existing methods such as LGCE (Yuan et al. 2024), LAVT (Yang et al. 2022), EVF-SAM (Zhang et al. 2024) and RMSIN (Liu et al. 2024b), etc., on the RRSIS-D (Liu et al. 2024b) dataset, with results documented in Tab. 2. On this dataset, our method achieves the best performance, surpassing the current state-of-the-art method RMSIN on the validation set by 4.59% in Pr@0.5, 5.86% in Pr@0.6, 6.44% in Pr@0.7, 5.28% in Pr@0.8, 5.97% in Pr@0.9, 1.89% in oIoU, and 3.71% in mIoU. It also significantly outperformed existing methods on the test set, demonstrating its strong capability in remote sensing object segmentation.

Fig. 4 illustrates the visual comparison between our model and RMSIN (Liu et al. 2024b) on the RRSIS-D dataset. The

Method	Pr@0.5		Pr@0.6		Pr@0.7		Pr@0.8		Pr@0.9		oIoU		mIoU	
	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test
BRINet (Hu et al. 2020)	36.86	20.72	35.53	14.26	19.93	9.87	10.66	2.98	2.84	1.14	61.59	58.22	38.73	31.51
LSCM (Hui et al. 2020)	56.82	31.54	41.24	20.41	21.85	9.51	12.11	5.29	2.51	0.84	62.82	61.27	40.59	35.54
CMPC (Huang et al. 2020)	46.09	32.36	26.45	14.14	12.76	6.55	7.42	1.76	1.39	0.22	63.55	55.39	42.08	40.63
CMSA (Ye et al. 2019)	39.24	28.07	38.44	20.25	20.39	12.71	11.79	5.61	1.52	0.83	65.84	64.53	43.62	41.47
RRN (Li et al. 2018)	55.43	30.26	42.98	23.01	23.11	14.87	13.72	7.17	2.64	0.98	69.24	65.06	50.81	41.88
EVF-SAM (Zhang et al. 2024)	57.77	35.17	37.59	22.34	16.24	9.36	4.87	2.86	1.86	0.39	59.61	55.51	46.98	36.64
CMPC+ (Liu et al. 2021)	56.84	49.19	37.59	28.31	20.42	15.31	10.67	8.12	2.78	2.55	70.62	66.53	47.13	43.65
CARIS (Liu et al. 2023b)	68.45	45.40	47.10	27.19	25.52	15.08	14.62	8.87	3.71	1.98	75.79	69.74	54.30	42.66
CRIS (Wang et al. 2022)	53.13	35.77	36.19	24.11	24.36	14.36	11.83	6.38	2.55	1.21	72.14	65.87	53.74	43.26
LAVT (Yang et al. 2022)	80.97	51.84	58.70	30.27	31.09	17.34	15.55	9.52	4.64	2.09	78.50	71.86	61.53	47.40
RIS-DMMI (Hu et al. 2023)	86.17	63.89	74.71	44.30	38.05	19.81	18.10	6.49	3.25	1.00	74.02	68.58	65.72	52.15
LGCE (Yuan et al. 2024)	90.72	73.75	86.31	61.14	71.93	39.46	<u>32.95</u>	16.02	<u>10.21</u>	5.45	<u>83.56</u>	76.81	72.51	59.96
RMSIN (Liu et al. 2024b)	<u>93.97</u>	79.20	<u>89.33</u>	65.99	<u>74.25</u>	42.98	29.70	16.51	7.89	3.25	82.41	75.72	<u>73.84</u>	62.58
FIANet (Lei et al. 2024)	-	<u>84.09</u>	-	<u>77.05</u>	-	<u>61.86</u>	-	<u>33.41</u>	-	<u>7.10</u>	-	<u>78.32</u>	-	<u>68.67</u>
RS2-SAM2	95.36	84.31	94.90	79.42	92.58	70.89	83.76	55.70	36.66	21.19	88.03	80.87	85.21	73.90

Table 1: Comparison with state-of-the-art methods on the RefSegRS dataset. The top-performing results are presented in bold, while the second-best results are underlined. Our model achieves the best performance across all metrics.

Method	Pr@0.5		Pr@0.6		Pr@0.7		Pr@0.8		Pr@0.9		oIoU		mIoU	
	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test
RRN (Li et al. 2018)	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
CMSA (Ye et al. 2019)	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 (Hui et al. 2020)	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 (Huang et al. 2020)	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 (Hu et al. 2020)	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+ (Liu et al. 2021)	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 (Yuan et al. 2024)	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
RIS-DMMI (Hu et al. 2023)	70.40	68.74	63.05	60.96	54.14	50.33	41.95	38.38	23.85	21.63	77.01	76.20	60.72	60.12
LAVT (Yang et al. 2022)	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
EVF-SAM (Zhang et al. 2024)	73.51	72.16	67.87	66.50	<u>58.33</u>	<u>56.59</u>	<u>46.15</u>	<u>43.92</u>	<u>25.92</u>	25.48	76.32	76.77	64.03	62.75
FIANet (Lei et al. 2024)	-	<u>74.46</u>	-	66.96	-	56.31	-	42.83	-	24.13	-	76.91	-	64.01
RMSIN (Liu et al. 2024b)	<u>74.66</u>	74.26	<u>68.22</u>	<u>67.25</u>	57.41	55.93	45.29	42.55	24.43	24.53	<u>78.27</u>	<u>77.79</u>	<u>65.10</u>	<u>64.20</u>
RS2-SAM2	79.25	77.56	74.08	72.34	63.85	61.76	50.57	47.92	30.40	29.73	80.16	78.99	68.81	66.72

Table 2: Comparison with state-of-the-art methods on the RRSIS-D dataset. The top-performing results are presented in bold, while the second-best results are underlined. Our model achieves the best performance across all metrics.

results clearly indicate that RS2-SAM2 consistently outperforms RMSIN, particularly in terms of accurate target localization, reliable mask prediction, and boundary precision.

Model Analysis

In this section, we perform extensive ablation studies to analyze the impact of the essential components in our RS2-SAM2 framework, as well as the effects of different model configurations. The experiments are conducted on the test dataset of RefSegRS (Yuan et al. 2024).

Components Analysis. To investigate the impact of key components in our model, we first construct a baseline model consisting solely of SAM2 and the union encoder. As shown in Tab. 3, when the text-guided boundary loss (\mathcal{L}_{tbl}) is incorporated during training, the model achieves 38.63% mIoU and 57.36% oIoU, surpassing the baseline by 1.99% and 1.85%, respectively. Subsequently, on top of the

model with \mathcal{L}_{tbl} , we add the mask prompt generator (MPG), which results in the RS2-SAM2 achieving 60.20% mIoU and 70.89% oIoU, an improvement of 21.57% and 13.53% over the previous model. When the bidirectional hierarchical fusion module (BHF) is added to the baseline model with \mathcal{L}_{tbl} , the model shows a further increase of 30.08% in mIoU and 21.00% in oIoU, further validating the superiority of this module. Finally, when all components are integrated, our RS2-SAM2 achieves the best performance, with a mIoU of 73.90% and an oIoU of 80.87%.

Mask Prompt Generator. In this section, we explore different interaction forms within the mask prompt generator (MPG), with results presented in Tab. 4. When the multi-head cross-attention (MHCA) between the multimodal token V_{cls} and visual embeddings V is omitted, the performance of RS2-SAM2 drops by 2.31% in mIoU and 0.98% in oIoU. These findings underscore the importance of further

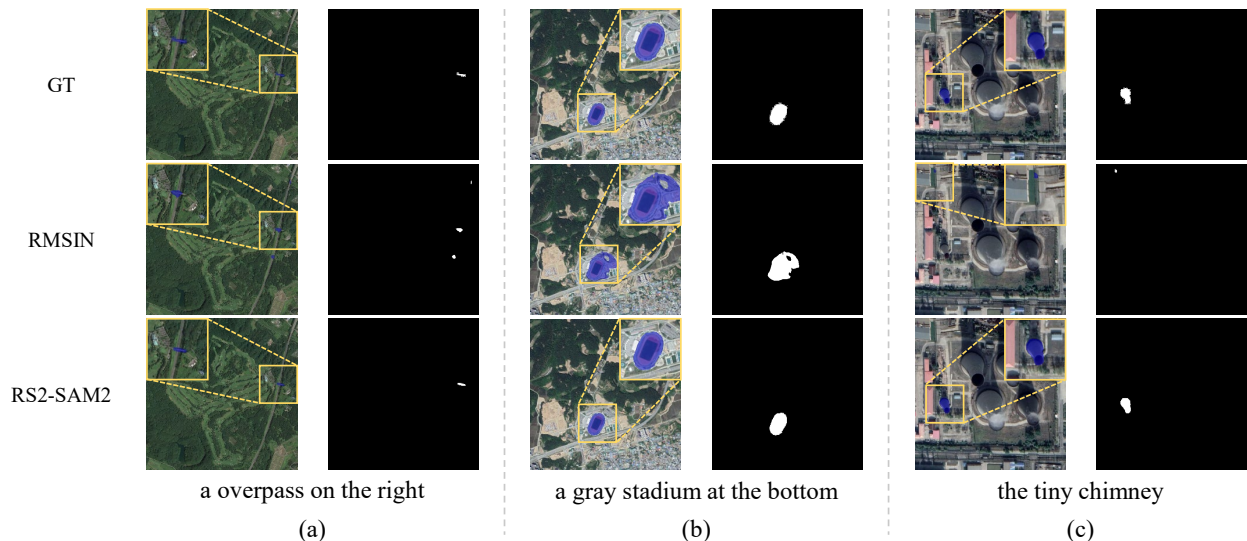


Figure 4: Visualization result on RRSIS-D. Compared to RMSIN (Liu et al. 2024b), RS2-SAM2 demonstrates superior capability in handling local details and boundary regions.

Method	\mathcal{L}_{tbl}	MPG	BHFM	Pr@0.5	Pr@0.7	Pr@0.9	mIoU	oIoU
Baseline				35.17	9.36	0.39	36.64	55.51
RS2-SAM2	✓			39.79	12.49	0.39	38.63	57.36
RS2-SAM2	✓	✓		71.00	43.42	3.96	60.20	70.89
RS2-SAM2	✓		✓	81.89	61.86	7.10	68.71	78.36
RS2-SAM2	✓	✓	✓	84.31	70.89	21.19	73.90	80.87

Table 3: Ablation study of different components of RS2-SAM2 on RefSegRS test set.

strengthening the semantic connections between the multi-modal token and visual embeddings during the multimodal mask prompt generation process.

Bidirectional Hierarchical Fusion Module. The effect of different BHFM configurations is also investigated. We first assess how the BHFM structure in the encoder influences performance. A simple variant uses a linear adapter-like layer without text interaction (“Linear”). Another uses text to enhance visual features without feedback (“Uni”). The third performs bidirectional enhancement between text and vision (“Bi”). As shown in Tab. 4, the layer-by-layer bidirectional enhancement proves more effective in infusing linguistic cues into visual features, enabling progressive textual refinement from low to high levels and resulting in more precise remote sensing segmentation.

Additionally, we investigate the impact of BHFM components on segmentation performance by setting up experiments where the BHFM cross-attention (BC) after encoding and the BHFM layer (BL) without the encoder are removed. The results in Tab. 4 indicate that both interactions during and after encoding are essential. The hierarchical interaction combining both can help the model understand text features from a global to a local perspective, enhancing its pixel-level

Method	Settings	Pr@0.5	Pr@0.7	Pr@0.9	mIoU	oIoU
Interaction Form of MPG						
RS2-SAM2	w/o MHCA	83.71	67.64	16.79	71.59	79.89
RS2-SAM2	w MHCA	84.31	70.89	21.19	73.90	80.87
Structure of BHFM layer						
RS2-SAM2	Linear	81.01	61.42	6.77	68.19	77.39
RS2-SAM2	Uni	81.23	63.84	14.14	70.10	78.93
RS2-SAM2	Bi	84.31	70.89	21.19	73.90	80.87
Components of BHFM						
RS2-SAM2	w/o BC	79.97	58.17	11.61	67.61	77.54
RS2-SAM2	w/o BL	72.21	43.48	4.79	61.08	72.33
RS2-SAM2	w BC&BL	84.31	70.89	21.19	73.90	80.87

Table 4: Model analysis of different settings in RS2-SAM2.

understanding of text.

Conclusion

In this paper, we present RS2-SAM2, an advanced end-to-end framework designed to enhance SAM2 for RRSIS. Our approach leverages a union encoder to jointly encode visual and textual features, producing semantically aligned visual-text embeddings and multimodal class tokens. To effectively integrate spatial and textual information, we introduce a bidirectional hierarchical fusion module, which incorporates textual semantics and spatial context both during and after encoding, enabling a hierarchical refinement from global to local levels. Additionally, a mask prompt generator generates multimodal mask as dense prompt, improving the segmentation of visually indistinct objects by providing stronger pixel-level guidance. Extensive experiments on multiple RRSIS benchmarks demonstrate the superiority of RS2-SAM2 over state-of-the-art methods, validating the effectiveness of our proposed modules.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62431020, the National Key Research and Development Program of China under Grant 2024YFE0111800, and the Fundamental Research Funds for the Central Universities under Grant 2042025kf0030.

References

- Chen, T.; Lu, A.; Zhu, L.; Ding, C.; Yu, C.; Ji, D.; Li, Z.; Sun, L.; Mao, P.; and Zang, Y. 2024. Sam2-adapter: Evaluating & adapting segment anything 2 in downstream tasks: Camouflage, shadow, medical image segmentation, and more. *arXiv preprint arXiv:2408.04579*.
- Cheng, Y.; Li, L.; Xu, Y.; Li, X.; Yang, Z.; Wang, W.; and Yang, Y. 2023. Segment and track anything. *arXiv preprint arXiv:2305.06558*.
- Cheng, Z.; Wei, Q.; Zhu, H.; Wang, Y.; Qu, L.; Shao, W.; and Zhou, Y. 2024. Unleashing the potential of SAM for medical adaptation via hierarchical decoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3511–3522.
- Conneau, A. 2019. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Hu, Y.; Wang, Q.; Shao, W.; Xie, E.; Li, Z.; Han, J.; and Luo, P. 2023. Beyond one-to-one: Rethinking the referring image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4067–4077.
- Hu, Z.; Feng, G.; Sun, J.; Zhang, L.; and Lu, H. 2020. Bi-directional relationship inferring network for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4424–4433.
- Huang, S.; Hui, T.; Liu, S.; Li, G.; Wei, Y.; Han, J.; Liu, L.; and Li, B. 2020. Referring image segmentation via cross-modal progressive comprehension. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10488–10497.
- Hui, T.; Liu, S.; Huang, S.; Li, G.; Yu, S.; Zhang, F.; and Han, J. 2020. Linguistic structure guided context modeling for referring image segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, 59–75. Springer.
- Ke, L.; Ye, M.; Danelljan, M.; Tai, Y.-W.; Tang, C.-K.; Yu, F.; et al. 2024. Segment anything in high quality. *Advances in Neural Information Processing Systems*, 36.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Lai, X.; Tian, Z.; Chen, Y.; Li, Y.; Yuan, Y.; Liu, S.; and Jia, J. 2024. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9579–9589.
- Lan, M.; Rong, F.; Jiao, H.; Gao, Z.; and Zhang, L. 2024. Language query based transformer with multi-scale cross-modal alignment for visual grounding on remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*.
- Lan, M.; Zhang, J.; Zhang, L.; and Tao, D. 2023. Learning to learn better for video object segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1205–1212.
- Lei, S.; Xiao, X.; Zhang, T.; Li, H.-C.; Shi, Z.; and Zhu, Q. 2024. Exploring fine-grained image-text alignment for referring remote sensing image segmentation. *IEEE Transactions on Geoscience and Remote Sensing*.
- Li, R.; Li, K.; Kuo, Y.-C.; Shu, M.; Qi, X.; Shen, X.; and Jia, J. 2018. Referring image segmentation via recurrent refinement networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5745–5753.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024a. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liu, J.; Ding, H.; Cai, Z.; Zhang, Y.; Satzoda, R. K.; Mahadevan, V.; and Manmatha, R. 2023a. Polyformer: Referring image segmentation as sequential polygon generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 18653–18663.
- Liu, S.; Hui, T.; Huang, S.; Wei, Y.; Li, B.; and Li, G. 2021. Cross-modal progressive comprehension for referring segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9): 4761–4775.
- Liu, S.; Ma, Y.; Zhang, X.; Wang, H.; Ji, J.; Sun, X.; and Ji, R. 2024b. Rotated multi-scale interaction network for referring remote sensing image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26658–26668.
- Liu, S.-A.; Zhang, Y.; Qiu, Z.; Xie, H.; Zhang, Y.; and Yao, T. 2023b. CARIS: Context-aware referring image segmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, 779–788.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *ICLR*.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Proceedings of the International Conference on 3D Vision (3DV)*, 565–571.
- Rasheed, H.; Maaz, M.; Shaji, S.; Shaker, A.; Khan, S.; Cholakkal, H.; Anwer, R. M.; King, E.; Yang, M.-H.; and Khan, F. S. 2024. Glamm: Pixel grounding large multi-modal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13009–13018.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Rong, F.; Lan, M.; Zhang, Q.; and Zhang, L. 2025. MPG-SAM 2: Adapting SAM 2 with Mask Priors and Global Context for Referring Video Object Segmentation. *arXiv preprint arXiv:2501.13667*.

- Ryali, C.; Hu, Y.-T.; Bolya, D.; Wei, C.; Fan, H.; Huang, P.-Y.; Aggarwal, V.; Chowdhury, A.; Poursaeed, O.; Hoffman, J.; et al. 2023. Hiera: A hierarchical vision transformer without the bells-and-whistles. In *International conference on machine learning*, 29441–29454. PMLR.
- Sun, Y.; Feng, S.; Li, X.; Ye, Y.; Kang, J.; and Huang, X. 2022. Visual grounding in remote sensing images. In *Proceedings of the 30th ACM International conference on Multimedia*, 404–412.
- Wang, C.; Chen, J.; Meng, Y.; Deng, Y.; Li, K.; and Kong, Y. 2024a. SAMPolyBuild: Adapting the Segment Anything Model for polygonal building extraction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 218: 707–720.
- Wang, C.; Ji, Y.; Meng, Y.; Zhang, Y.; and Zhu, Y. 2025a. SOPSeg: Prompt-based Small Object Instance Segmentation in Remote Sensing Imagery. [arXiv:2509.03002](https://arxiv.org/abs/2509.03002).
- Wang, C.; Xi, Z.; Liu, D.; Feng, Y.; Deng, Y.; Li, K.; Chen, J.; Chen, J.; and Meng, Y. 2025b. PCP: A Prompt-Based Cartographic-Level Polygonal Vector Extraction Framework for Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing*, 63: TGRS–2025.
- Wang, D.; Zhang, J.; Du, B.; Xu, M.; Liu, L.; Tao, D.; and Zhang, L. 2024b. Samrs: Scaling-up remote sensing segmentation dataset with segment anything model. *Advances in Neural Information Processing Systems*, 36.
- Wang, T.; Xiao, X.; Chen, G.; Chi, H.; Zhang, Q.; Cheng, G.; and Ji, Y. 2025c. TASAM: Terrain-and-Aware Segment Anything Model for Temporal-Scale Remote Sensing Segmentation. [arXiv:2509.15795](https://arxiv.org/abs/2509.15795).
- Wang, W.; Bao, H.; Dong, L.; Bjorck, J.; Peng, Z.; Liu, Q.; Aggarwal, K.; Mohammed, O. K.; Singhal, S.; Som, S.; et al. 2023. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19175–19186.
- Wang, Z.; Lu, Y.; Li, Q.; Tao, X.; Guo, Y.; Gong, M.; and Liu, T. 2022. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11686–11695.
- Wu, J.; Ji, W.; Liu, Y.; Fu, H.; Xu, M.; Xu, Y.; and Jin, Y. 2023. Medical sam adapter: Adapting segment anything model for medical image segmentation. [arXiv preprint arXiv:2304.12620](https://arxiv.org/abs/2304.12620).
- Xiong, Y.; Varadarajan, B.; Wu, L.; Xiang, X.; Xiao, F.; Zhu, C.; Dai, X.; Wang, D.; Sun, F.; Iandola, F.; et al. 2024. Efficientsam: Leveraged masked image pretraining for efficient segment anything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16111–16121.
- Xu, J.; Xu, L.; Yang, Y.; Li, X.; Wang, F.; Xie, Y.; Huang, Y.-J.; and Li, Y. 2023. u-llava: Unifying multi-modal tasks via large language model. [arXiv preprint arXiv:2311.05348](https://arxiv.org/abs/2311.05348).
- Yan, C.; Wang, H.; Yan, S.; Jiang, X.; Hu, Y.; Kang, G.; Xie, W.; and Gavves, E. 2024. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*, 98–115. Springer.
- Yang, Z.; Wang, J.; Tang, Y.; Chen, K.; Zhao, H.; and Torr, P. H. 2022. Lavt: Language-aware vision transformer for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18155–18165.
- Ye, L.; Rochan, M.; Liu, Z.; and Wang, Y. 2019. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10502–10511.
- Yingrui Ji, J. C. A. Y. C. W. K. L. Y. Z., Jiansheng Chen. 2025. RS-OOD: A Vision-Language Augmented Framework for Out-of-Distribution Detection in Remote Sensing. [arXiv:2509.02273](https://arxiv.org/abs/2509.02273).
- Yu, S.; Jung, I.; Han, B.; Kim, T.; Kim, Y.; Wee, D.; and Son, J. 2024. A Simple Baseline with Single-encoder for Referring Image Segmentation. [arXiv preprint arXiv:2408.15521](https://arxiv.org/abs/2408.15521).
- Yuan, Z.; Mou, L.; Hua, Y.; and Zhu, X. X. 2024. Rrsis: Referring remote sensing image segmentation. *IEEE Transactions on Geoscience and Remote Sensing*.
- Yue, W.; Zhang, J.; Hu, K.; Xia, Y.; Luo, J.; and Wang, Z. 2024. Surgicalsam: Efficient class promptable surgical instrument segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6890–6898.
- Zhan, Y.; Xiong, Z.; and Yuan, Y. 2023. Rsvg: Exploring data and models for visual grounding on remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13.
- Zhang, C.; Han, D.; Qiao, Y.; Kim, J. U.; Bae, S.-H.; Lee, S.; and Hong, C. S. 2023. Faster segment anything: Towards lightweight sam for mobile applications. [arXiv preprint arXiv:2306.14289](https://arxiv.org/abs/2306.14289).
- Zhang, X.; Zhang, H.; Wang, G.; Zhang, Q.; Zhang, L.; and Du, B. 2025. UniUIR: Considering Underwater Image Restoration as An All-in-One Learner. [arXiv preprint arXiv:2501.12981](https://arxiv.org/abs/2501.12981).
- Zhang, Y.; Cheng, T.; Hu, R.; Liu, L.; Liu, H.; Ran, L.; Chen, X.; Liu, W.; and Wang, X. 2024. Evf-sam: Early vision-language fusion for text-prompted segment anything model. [arXiv preprint arXiv:2406.20076](https://arxiv.org/abs/2406.20076).
- Zhong, Z.; Tang, Z.; He, T.; Fang, H.; and Yuan, C. 2024. Convolution meets lora: Parameter efficient finetuning for segment anything model. [arXiv preprint arXiv:2401.17868](https://arxiv.org/abs/2401.17868).