

RSVG-ZeroOV: Exploring a Training-Free Framework for Zero-Shot Open-Vocabulary Visual Grounding in Remote Sensing Images

Ke Li¹, Di Wang^{1*}, Ting Wang¹, Fuyu Dong¹, Yiming Zhang², Luyao Zhang¹, Xiangyu Wang¹, Shaofeng Li¹, Quan Wang¹

¹School of Computer Science and Technology, Xidian University, 710126, China.

²Department of Mathematics, University of California, San Diego, USA.

Abstract

Remote sensing visual grounding (RSVG) aims to localize objects in remote sensing images based on free-form natural language expressions. Existing approaches are typically constrained to closed-set vocabularies, limiting their applicability in open-world scenarios. While recent attempts to leverage generic foundation models for open-vocabulary RSVG, they overly rely on expensive high-quality datasets and time-consuming fine-tuning. To address these limitations, we propose **RSVG-ZeroOV**, a training-free framework that aims to explore the potential of frozen generic foundation models for zero-shot open-vocabulary RSVG. Specifically, RSVG-ZeroOV comprises three key stages: (i) *Overview*: We utilize a vision-language model (VLM) to obtain cross-attention maps that capture semantic correlations between text queries and visual regions. (ii) *Focus*: By leveraging the fine-grained modeling priors of a diffusion model (DM), we fill in gaps in structural and shape information of objects, which are often overlooked by VLM. (iii) *Evolve*: A simple yet effective attention evolution module is introduced to suppress irrelevant activations, yielding purified segmentation masks over the referred objects. Without cumbersome task-specific training, RSVG-ZeroOV offers an efficient and scalable solution. Extensive experiments demonstrate that the proposed framework consistently outperforms existing weakly-supervised and zero-shot methods.

Introduction

Deep learning has been widely applied in remote sensing tasks, significantly improving the performance of applications such as image classification, object detection, and semantic segmentation (Chen et al. 2024a; Li et al. 2024a; Shan et al. 2025; Pan et al. 2025). However, these advances are heavily dependent on large-scale annotated data to learn how to recognize and localize objects. Given the broad geographic coverage and the diverse types of objects present in remote sensing imagery, existing datasets cannot fully capture the diversity of real-world objects. As a result, when encountering out-of-distribution objects, existing models may struggle to recognize and localize these objects, limiting their reliability in practical remote sensing scenarios.

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

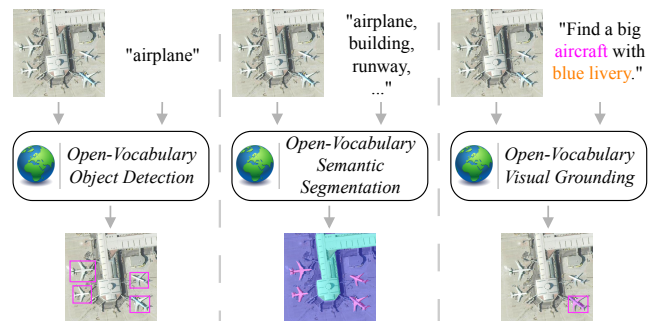


Figure 1: Illustration of different open-vocabulary remote sensing tasks. Left and middle panels depict *category-driven* tasks, where models rely on predefined category names to distinguish different image regions. Right panel represents *intention-driven* paradigm, enabling users to flexibly specify objects using natural language expressions.

Recently, generic foundation models have demonstrated remarkable open-world perception capabilities (Liu et al. 2023; Kirillov et al. 2023; Liu et al. 2024d). Benefiting from pre-training on extensive image-text pairs, these models are able to effectively understand the complex semantic relationships between textual descriptions and visual elements, thereby achieving impressive vision-language alignment.

Some researchers (Yao et al. 2025; Huang, He, and Wen 2025) have begun exploring the potential of foundation models in open-vocabulary category-driven tasks, such as remote sensing object detection and segmentation, relying on predefined object categories (e.g., “road”, “farmland”) to calculate the similarity with image regions, as illustrated in Fig. 1 (left and middle). Furthermore, in many real-world scenarios, users are more interested in locating specific targets. For instance, in urban analysis, planners may seek to identify “the tallest building by the river” or “the factory next to the playground.” Moreover, some objects cannot be represented by a simple category name. For example, temporary roadside parking zones may visually resemble roads but functionally serve as informal parking lots. Such expressions require models to reason about visual attributes, spatial relationships, and functional roles. Conventional open-vocabulary detection or segmentation methods are insuffi-

cient for handling such intention-driven tasks. In contrast, zero-shot open-vocabulary remote sensing visual grounding (RSVG) offers a more practical and flexible solution.

In this work, we focus on a central question: *how to extend the capabilities of frozen foundation models to zero-shot open-vocabulary RSVG?* To answer this, we conduct a series of exploratory experiments to progressively analyze the impact of different foundation models and architectural components on performance. Through these explorations, we summarize three empirical guidelines: **1)** Generic VLMs exhibit strong generalization ability, which is critical for object localization in remote sensing. **2)** Compared to other visual models, DMs demonstrate superior perception of the inherent structural concepts of objects. **3)** Attention maps from VLMs and DMs are complementary, and their integration consistently improves grounding performance.

Based on the above guidelines, we propose a training-free framework named **RSVG-ZeroOV**, a straightforward yet effective approach to solve zero-shot open-vocabulary RSVG. Our framework follows an overview-focus-evolve strategy that leverages the distinct yet complementary attention patterns of frozen VLM and DM to form the basis of our grounding results. Specifically, we first aggregate cross-attention maps¹ from all transformer heads in VLM to highlight regions corresponding to the referred objects. However, we observe two major limitations in these attention maps: one is that attention tends to concentrate on object boundaries and corners rather than the full extent, while another is that attention is rather scattered, often including irrelevant regions. To address the first issue, we devise an attention interaction module to facilitate interaction between the cross-attention maps and the self-attention maps, thereby enhancing the activation of pixels corresponding to the referred objects. For the latter, we propose a simple yet effective attention evolution module to suppress irrelevant activations, resulting in accurate and purified segmentation masks. Moreover, we incorporate the segment anything model (SAM) as an optional post-processing step to further improve mask quality. Extensive experiments conducted on two RSVG benchmarks show that RSVG-ZeroOV achieves exceptional performance in open-vocabulary RSVG under zero-shot setting.

- We propose RSVG-ZeroOV, a training-free framework for zero-shot open-vocabulary RSVG. We take the pioneering step by integrating attention maps from VLM and DM to enhance the perception of referred objects.
- We introduce an overview-focus-evolve strategy to progressively refine the attention maps, enabling the model to effectively capture the referred objects in remote sensing images.
- Extensive experiments on two RSVG datasets demonstrate that RSVG-ZeroOV achieves superior zero-shot performance, confirming the effectiveness of foundation models in open-world remote sensing perception.

¹In this paper, although decoder-only VLMs use self-attention over all tokens, we refer to the image-text interaction part as cross-attention to distinguish it from pure visual self-attention.

Related Work

Generic Foundation Model. The emergence of large-scale pretraining has fueled the rapid advancement of multimodal foundation models. Among them, visual-language models (VLMs) are designed for cross-modal understanding by aligning visual and textual representations. Early attempts like CLIP (Radford et al. 2021) demonstrated impressive zero-shot generalization by contrastively pretraining on massive image-text pairs. LLaVA (Liu et al. 2023) builds upon CLIP by adding a language decoder Vicuna (Chiang et al. 2023) to support multimodal instruction-following. More recent approaches, such as CogVLM (Wang et al. 2024b), InternVL (Chen et al. 2024b), and Qwen-VL (Wang et al. 2024a), further enhance this paradigm and demonstrate outstanding performance in open-world visual understanding. However, due to their emphasis on high-level semantic reasoning and lack of pixel-level supervision, these models often struggle with fine-grained spatial perception. In contrast, another branch represented by diffusion models (DMs) (Rombach et al. 2022), focuses on visual content generation from text prompts. Recent studies (Pnvr et al. 2023; Xu et al. 2023) reveal that DMs’ hierarchical attention encodes strong structural priors, making their attention maps applicable to downstream tasks like classification, detection, and segmentation. In this study, we conduct systematic exploratory experiments to investigate the potential of integrating VLM and DM models for open-vocabulary RSVG, which will be elaborated in the next section.

Remote Sensing Visual Grounding. RSVG aims to localize objects in remote sensing imagery based on natural language descriptions. Compared to natural scene images, remote sensing data is acquired via satellites and often encounters unique challenges such as scale variation and cluttered backgrounds. Remote sensing referring expression comprehension (RSREC) and segmentation (RSRES) are two typical RSVG tasks. RSREC aims to predict a bounding box corresponding to the given referring expression. The pioneering work (Sun et al. 2022) explores RSREC through a one-stage method based on dense (Ding et al. 2024). Recently, transformer-based methods (Zhan, Xiong, and Yuan 2023; Li et al. 2024b; Lan et al. 2024; Ding et al. 2024) can effectively capture cross-modal context and achieve better performance, benefiting from the attention mechanism. For RSRES, the goal is to predict a precise pixel-wise binary mask specified by the given referring expression. Recent works (Yuan et al. 2023; Liu et al. 2024c; Dong et al. 2024; Li et al. 2025b) improve segmentation by enhancing multi-scale visual-text feature interactions.

Open-Vocabulary Visual Perception. With the rapid development of multimodal foundation models, open-vocabulary learning has achieved remarkable success in natural image domains. Inspired by this progress, recent studies have extended these paradigms to the remote sensing field. For instance, RemoteCLIP (Liu et al. 2024a) improves transferability in downstream tasks by pretraining on extensive remote sensing image-expression pairs. Similarly, VisGT (Pan et al. 2025) and GNet (Ye, Zhuge, and Zhang 2025) construct large-scale open-vocabulary detection and segmentation datasets, respectively, enabling models to recognize a

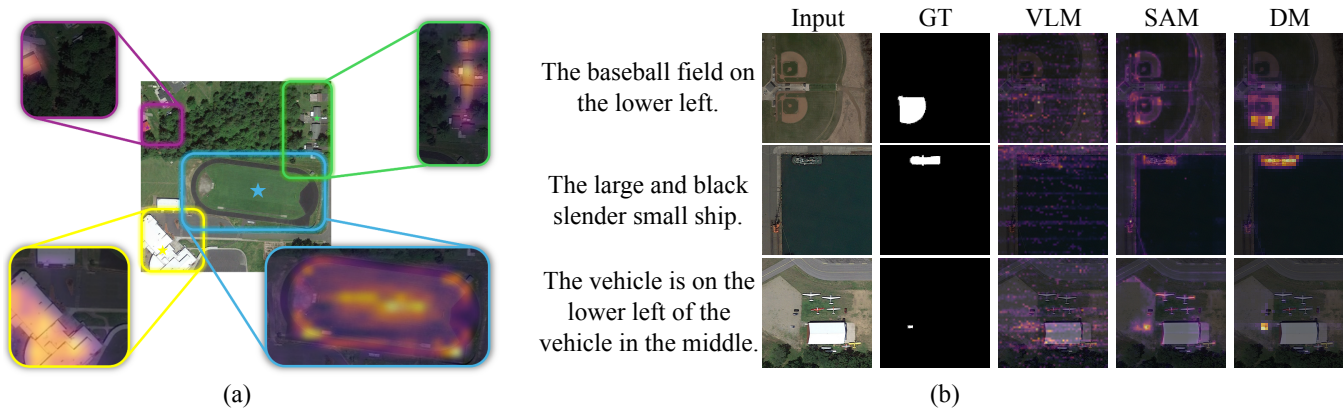


Figure 2: (a) Visualization of receptive fields derived from self-attention maps in DM. (b) Comparison of attention embedding results using self-attention maps from different models.

broader range of object categories. However, curating high-quality annotated datasets for remote sensing remains costly and labor-intensive. To alleviate reliance on manual annotations, several efforts (Zheng et al. 2025) have explored training-free approaches. For example, SegEarth-OV (Li et al. 2025a) adapts existing generic models by modifying or removing task-specific components, allowing seamless application to high-resolution remote sensing segmentation. Nevertheless, these methods still rely on predefined category labels during inference and lack the flexibility to understand user-intended objects described in free-form language. As discussed in the introduction, users in practical applications often prefer to locate objects by describing visual attributes, spatial relationships, or functional roles using natural language expressions. To address this limitation, we introduce the task of zero-shot open-vocabulary RSVG, which requires models to flexibly interpret free-form expressions and localize the referred objects. To this end, we propose a training-free framework that integrates attention maps from frozen VLM and DM to facilitate this task.

Guidelines of Exploring Foundation Models

Naively transferring generic foundation models to open-vocabulary remote sensing visual perception tasks does not yield satisfactory results. In this section, we summarize three guidelines for the effective use of foundation models.

Guideline 1: generic VLMs exhibit strong generalization ability on unseen remote sensing data. To demonstrate this, we conduct zero-shot experiments on RRSIS-D test set. Specifically, we compare three 7B-scale models, including LLaVA-1.5 (Liu et al. 2024b), Qwen2.5-VL (Bai et al. 2025), and GeoChat (Kuckreja et al. 2024). For RSREC task, we generate bounding box predictions of referred objects using a unified prompt format: [Locate the object referred to by ‘{referring expression}’ and return its box coordinates (x1, y1, x2, y2)]. Notably, GeoChat has been finetuned on 318k image-instruction pairs collected from remote sensing datasets, enabling it to learn more domain-specific priors. However, as shown in Tab. 1, GeoChat achieves 23.93% Pr@0.5 on RSREC task, noticeably lower

Method (Zero-Shot)	RSREC		RSRES	
	Pr@0.5	mIoU	Pr@0.5	mIoU
GeoChat (CVPR’24)	23.93	30.12	-	-
LLaVA-1.5 (CVPR’24)	13.21	20.59	-	-
Qwen2.5-VL (arXiv’25)	28.66	30.90	-	-
GeoChat + SAM	27.61	32.53	18.85	24.92
LLaVA-1.5 + SAM	16.26	21.97	10.86	16.86
Qwen2.5-VL + SAM	30.35	31.93	24.68	25.72

Table 1: Comparison of generic VLMs and remote sensing VLM on RRSIS-D test set, all using 7B-scale LLMs.

than Qwen2.5-VL’s **28.66%**. This result suggests that VLMs pretrained on large-scale generic datasets can exhibit strong cross-domain generalization even without access to remote sensing data. For RSRES task, we utilize outputs of VLMs as box prompts for SAM to get pixel-level masks. As we can see, Qwen2.5-VL still outperforms GeoChat by **0.8%** in mIoU. These findings highlight the potential of generic VLMs as a solid foundation for zero-shot open-vocabulary RSVG, even without any domain-specific training.

Guideline 2: self-attention maps in DM encode superior structural priors of objects. Theoretically, self-attention mechanisms are inherently suited for modeling object structure, as they compute pairwise similarities across visual elements, enabling the network to capture spatial dependencies and shape layouts. As shown in Fig. 2(a), we visualize the receptive fields of self-attention maps anchored at different query points. The results demonstrate that self-attention can effectively distinguish foreground objects from the background, reflecting a strong capacity for structural perception. To further investigate this, we compare the spatial distributions of self-attention extracted from three representative architectures: the visual encoder of VLM, the image backbone of SAM, and the U-Net of DM, as illustrated in Fig. 2(b). Specifically, the VLM attention maps tend to be broadly distributed and spatially scattered, consistent with its objective of capturing global semantic context. While SAM produces

Self-Attention Maps	RSREC		RSRES	
	Pr@0.5	mIoU	Pr@0.5	mIoU
w/o Cross-Attention Map	21.49	26.26	1.18	6.15
ViT in VLM	12.53	21.02	2.13	7.40
ViT in SAM	18.85	23.56	4.77	10.23
U-Net in DM (Ours)	30.15	32.92	12.84	21.85

Table 2: Evaluation of different self-attention maps on RRSIS-D test set.

sharper boundaries, its purely visual design lacks semantic understanding, which often leads to over-attention on surrounding background areas. For instance, in the second row, SAM simultaneously highlights both the ship and port. In contrast, the self-attention maps from DM display the most coherent structural representations, *i.e.*, attention is evenly and densely distributed across the entire object extent, revealing a more precise and holistic understanding of object shapes. To quantitatively validate these observations, we integrate the different types of self-attention maps into our visual grounding pipeline and evaluate their performance. As reported in Tab. 2, the use of self-attention maps from DM (the last row) achieves the highest Pr@0.5 and mIoU on both RSREC and RSRES benchmarks. These results confirm the superiority of DM’s self-attention in capturing object-centric structural cues for visual grounding tasks.

Guideline 3: combining cross- and self-attention maps is vital especially for pixel-level perception tasks. To validate this, we first evaluate the grounding performance using only the cross-attention maps from the VLM, as shown in Tab. 2 (the first row). Removing the self-attention maps from DM leads to significant performance drops of **8.66%** Pr@0.5 and **6.66%** mIoU on RSREC, and **11.66%** and **15.70%** on RSRES, respectively. Such a phenomenon indicates the necessity of combining cross- and self-attention maps, especially for pixel-level perception tasks like RSRES. These findings naturally lead to a question: *which interaction strategy best integrates cross- and self-attention maps for optimal grounding performance?* We try four strategies: 1) Anchor-based: selecting high-response pixels from the cross-attention map to represent the object, followed by aggregating their corresponding self-attention regions. 2) Multiplication: performing element-wise multiplication on cross- and self-attention maps, denoted as $\mathcal{A}_S \cdot \mathcal{A}_C$. 3) Exponentiation: enhancing the self-attention map by raising it to a power γ , *i.e.*, $\mathcal{A}_S^\gamma \cdot \mathcal{A}_C$. 4) Similarity: computing cosine similarity between the cross- and self-attention maps at each spatial location (details in the next section). As summarized in Tab. 3, although the anchor-based and exponentiation strategies achieve reasonable performance on RSREC, they suffer on RSRES. The former tends to oversimplify by omitting many regions, while the latter tends to overamplify high activations. In contrast, the proposed similarity strategy yields more semantically consistent initial masks, which allows the evolve stage to further improve segmentation quality, ultimately achieving the best performance on both benchmarks.

Interaction Strategy	RSREC		RSRES	
	Pr@0.5	mIoU	Pr@0.5	mIoU
Anchor-based	23.18	26.94	6.58	16.69
+ Evolve stage	28.73	31.48	7.61	16.58
Multiplication	21.94	26.32	10.03	19.78
+ Evolve stage	29.26	31.54	11.52	20.75
Exponentiation	24.07	28.61	5.37	15.30
+ Evolve stage	27.38	31.02	5.31	14.00
Similarity (Ours)	22.63	26.65	10.26	20.56
+ Evolve stage	30.15	32.92	12.84	21.85

Table 3: Comparison of different interaction strategies between cross- and self-attention maps on RRSIS-D test set.

Proposed Framework

Following our proposed guidelines, we present a straightforward training-free framework named **RSVG-ZeroOV** for zero-shot open-vocabulary RSVG. As illustrated in Fig. 3, the framework comprises three sequential stages: (i) *Overview*: we prompt a VLM with a referring expression to identify the referred object and extract its corresponding cross-attention map, which highlights regions most relevant to the given textual query. (ii) *Focus*: we cache the self-attention maps from the U-Net of a DM to obtain structural priors of objects, and then propose an effective interaction strategy to embed these priors into the cross-attention distribution. (iii) *Evolve*: to generate a high-quality attention map, we introduce an attention evolution module to filter out irrelevant activations, thus improving the quality of the segmentation mask.

Overview Stage. Referring to **Guideline 1**, we utilize the attention map from a frozen VLM to capture the coarse-level localization cues of the referred object. Given a remote sensing image \mathcal{I} and a referring expression \mathcal{T} , we first hook into the attention weights $\mathcal{W}^{(t)} \in \mathbb{R}^{H \times 1 \times N}$ from VLM, where N is the token length, H is the number of attention heads, and $t \in \{1, \dots, T\}$ indexes the T autoregressive forward passes. To obtain cross-attention parts, we extract the image-text-related segment $\mathcal{W}_{p:p'}^{(t)}$, where p and p' indicate the span of visual tokens in the input sequence. After that, we aggregate information from the transformer heads with mean weights, and form a sentence-level cross-attention map by averaging all t as:

$$\mathcal{A}_C = 1/T \sum_{t=1}^T (1/H \sum_{h=1}^H \mathcal{W}_{p:p'}^{(t)}). \quad (1)$$

As shown in Fig. 4, cross-attention maps establish the relationship between textual queries and visual pixels. However, it exhibits two key limitations: (i) *Attention tends to focus on object boundaries or corners rather than the full extent of the object.* This behavior stems from the high-level semantic concentration of VLMs, which encourages attention to concentrate on key features of objects. (ii) *Attention is often scattered and includes irrelevant regions.* This dispersion arises because VLMs need to aggregate contextual cues from multiple visual regions to understand complex referring expressions.

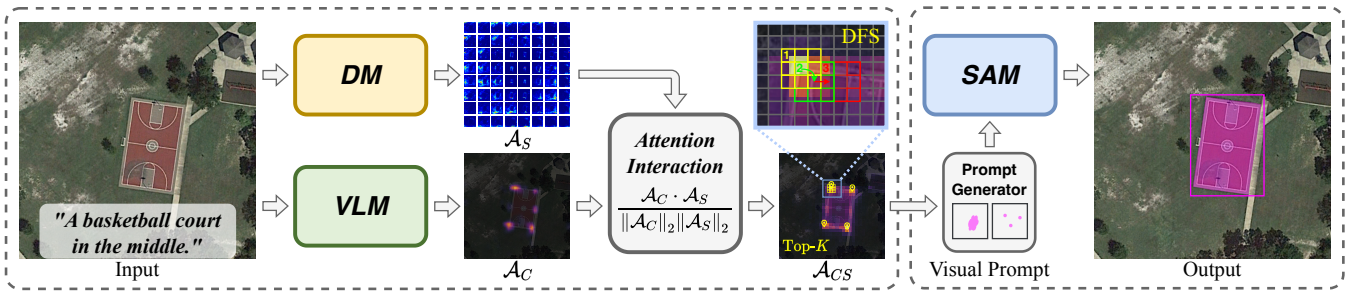


Figure 3: The framework of the proposed RSVG-ZeroOV.

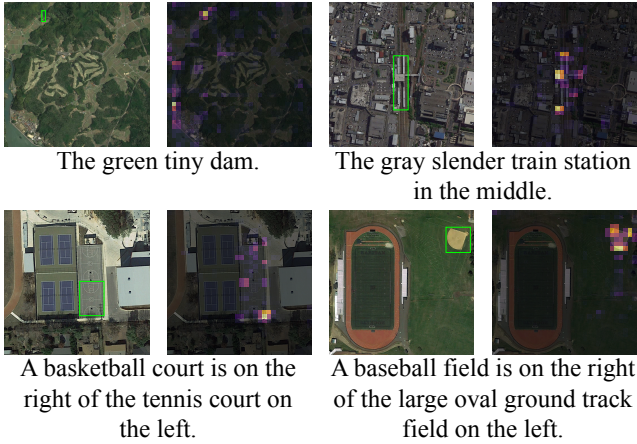


Figure 4: Visualization of some cross-attention maps from VLM on RRSIS-D test set.

Focus Stage. To mitigate the first *issue (i)*, referring to the aforementioned **Guidelines 2&3**, we propose an attention interaction module that integrates cross-attention with structural priors derived from self-attention. Considering the diverse scales of objects in remote sensing images, we first extract multi-scale self-attention maps from the U-Net in DM, and fuse them to form a unified structural prior $\mathcal{A}_S \in \mathbb{R}^{H \times W \times H \times W}$, which is denoted as:

$$\mathcal{A}_S = 1/L \sum_{l \in L} (\mathcal{A}_S^l), \quad (2)$$

where \mathcal{A}_S^l denotes the self-attention map from layer l . Next, we calculate the correlation score² between \mathcal{A}_C and \mathcal{A}_S through cosine similarity:

$$\mathcal{A}_{CS} = \frac{\mathcal{A}_C \cdot \mathcal{A}_S}{\|\mathcal{A}_C\|_2 \|\mathcal{A}_S\|_2}, \quad (3)$$

where $\|\cdot\|_2$ denotes the L_2 norm. This interaction introduces object-centric structural guidance into the original cross-attention map, allowing the model to shift focus from incomplete boundaries to semantically coherent regions.

Evolve Stage. Although this interaction strategy remarkably fills the gaps within the object area, it also exacerbates

²In fact, \mathcal{A}_C does not match the spatial resolution of \mathcal{A}_S , hence we interpolate \mathcal{A}_C to the same size as \mathcal{A}_S before interaction.

the impact of *issue (ii)*. To address this, we introduce an attention evolution module based on recursive region expansion to suppress scattered activation signals that are outside the referred object area. Specifically, we first select the top- K pixels from the interaction map \mathcal{A}_C with the highest response values:

$$\mathcal{S} = \text{TopK}(\mathcal{A}_C, K), \quad (4)$$

where $\mathcal{S} = \{(i_k, j_k)\}_{k=1}^K$ denotes the set of starting seed positions. For each seed (i_k, j_k) , we define a local neighborhood kernel $\mathcal{N}(i, j)$ of size 3×3 , and iteratively expand the region by selecting the neighbor with the highest response:

$$(i^*, j^*) = \arg \max_{(u, v) \in \mathcal{N}(i, j)} \mathcal{A}_{CS}[u, v]. \quad (5)$$

We then perform a depth-first search (DFS) starting from each seed to recursively grow the region. A pixel (u, v) is included in the evolved attention map \mathcal{A}_E if it satisfies:

$$\mathcal{A}_{CS}[u, v] \geq \tau, \ \&\& \ (u, v) \in \text{DFS}(\mathcal{S}), \quad (6)$$

where τ is a response threshold that controls the expansion boundary, and the process continues until no further qualified neighbors can be added. In addition, we conduct an ablation study to analyze the impact of the ordering between the **Focus** and **Evolve** stages on performance (see Tab. 7 for details). Finally, we perform binarization on the evolved attention map \mathcal{A}_E to obtain the segmentation mask $\mathbf{M} \in \{0, 1\}^{H \times W}$, which is defined as:

$$\mathbf{M}(i, j) = \begin{cases} 1, & \text{if } \mathcal{A}_E(i, j) > \alpha \\ 0, & \text{otherwise} \end{cases}, \quad (7)$$

where $\alpha \in [0, 1]$ is a predefined binarization threshold.

Refine Stage (Optional). Following prior work (Liu et al. 2024c; Dong et al. 2024), we additionally introduce an optional refinement stage to further improve the segmentation mask \mathbf{M} . We evaluate the effectiveness of several refinement methods, including DenseCRF and various prompt types for SAM. Our results show that box prompts yield the best performance, whereas point prompts are less effective.

Experiments

Experimental Settings

Datasets. We evaluate our method on two benchmarks tailored for RSVG: RRSIS-D (Liu et al. 2024c) and RISBench

Method		RSREC					RSRES				
		Pr@0.3	Pr@0.5	Pr@0.7	mIoU	oIoU	Pr@0.3	Pr@0.5	Pr@0.7	mIoU	oIoU
W-S	TRIS (ICCV'23)	14.62	3.79	0.37	13.20	13.14	15.02	4.91	1.26	13.11	15.40
	SAG (ICCV'23)	9.39	2.10	0.29	9.22	9.21	11.63	4.02	0.60	11.10	11.13
	QueryMatch (MM'24)	22.04	16.22	12.10	17.21	15.26	20.97	15.54	10.62	15.73	10.73
Z-S	VLM (Baseline)	43.35	28.66	14.27	30.90	29.53	-	-	-	-	-
	DiffSegmenter (TIP'25)	8.22	1.69	0.20	8.70	8.63	2.18	0.23	0.00	4.47	4.35
	DiffSegmenter (w/ VLM)	39.80	24.29	13.01	28.47	26.92	24.99	9.48	1.61	19.03	15.19
	DiffPNG (ECCV'24)	18.93	8.45	3.30	14.84	12.54	17.70	8.27	3.02	14.86	14.71
	DiffPNG (w/ VLM)	37.92	23.18	11.35	26.93	25.65	19.25	6.78	0.95	15.49	10.55
	OV-VG (arXiv'24)	27.78	19.16	12.96	22.07	18.34	-	-	-	-	-
	GroundVLP (AAAI'24)	21.66	16.83	12.27	17.14	17.51	-	-	-	-	-
RSVG-ZeroOV (Ours)	45.71	30.15	16.74	32.92	32.94	30.97	12.84	2.36	21.85	18.85	
w/ Refine	VLM (Baseline)	43.60	30.35	16.76	31.93	28.38	38.18	24.68	11.00	25.72	20.46
	DiffSegmenter (TIP'25)	8.16	1.55	0.17	8.60	8.58	3.02	0.32	0.03	4.60	4.70
	DiffSegmenter (w/ VLM)	38.41	25.11	13.96	28.50	24.57	32.92	19.42	9.14	23.73	18.00
	DiffPNG (ECCV'24)	19.59	10.31	5.03	15.71	13.11	19.65	11.32	6.23	16.22	13.77
	DiffPNG (w/ VLM)	32.43	21.29	13.04	24.89	20.79	27.03	17.64	11.00	20.99	13.53
	OV-VG (arXiv'24)	27.23	16.20	13.27	21.62	17.18	20.68	15.51	9.05	16.17	9.71
	GroundVLP (AAAI'24)	21.26	16.20	11.32	16.51	16.88	18.07	13.33	7.24	13.14	11.23
RSVG-ZeroOV (Ours)	45.70	31.39	17.63	34.49	31.28	40.01	27.39	13.38	28.35	22.83	

Table 4: Comparisons with SOTA methods on **RRSIS-D** test set. ‘W-S’ and ‘Z-S’ denote the weakly-supervised and zero-shot settings, respectively. ‘w/ VLM’ indicates that the cross-attention from the VLM is used to replace the original cross-attention. ‘w/ Refine’ indicates that post-processing operations are applied to further refine the segmentation results.

Method	RSREC		RSRES	
	Pr@0.5	mIoU	Pr@0.5	mIoU
GeoChat (CVPR'24) + SAM	27.61	32.53	18.85	24.92
Qwen2.5-VL (arXiv'25) + SAM	30.35	31.93	24.68	20.46
LISA (CVPR'24)	25.80	27.78	24.51	26.78
NExT-Chat (ICML'24)	-	-	26.37	24.98
RSVG-ZeroOV (Ours)	31.39	34.49	27.39	28.35

Table 5: Comparisons with SOTA generic pixel-level VLMs and a remote sensing VLM on **RRSIS-D** test set.

(Dong et al. 2024). Each sample comprises an RGB image, a referring expression, and a segmentation mask. Both datasets are extended from the RSREC dataset, inheriting its bounding box annotations, which allows for comprehensive evaluation of both REREC and RSRES tasks.

Implementation Details. Our RSVG-ZeroOV offers a zero-shot solution, requiring only an inference process, without the need for any training images or annotations. We employ the pre-trained Qwen2.5-VL (Bai et al. 2025) and Stable Diffusion V1.4 (Rombach et al. 2022) as our VLM and DM, respectively. For Stable Diffusion, we use a guidance scale of 7.5 with 1,000 total diffusion steps, and perform DDIM-based sampling with 20 steps. The DDIM noise schedule follows a scaled linear progression from 0.00085 to 0.012. Following prior works (Li et al. 2024b; Liu et al. 2024c), we report mIoU, oIoU and Precision@X with $X \in \{0.3, 0.5, 0.7\}$. We set $K=7$ for seed selection, $\tau=0.3$ for response thresholding, and $\alpha=0.4$ for binarization. All experiments are conducted on a single RTX4090 GPU.

Main Results

Comparison Results on RRSIS-D. As shown in Tab. 4, RSVG-ZeroOV (w/ Refine) achieves SOTA performance under the zero-shot setting, surpassing all weakly-supervised and zero-shot methods. Compared to the best weakly-supervised method QueryMatch, our approach achieves substantial gains in mIoU **+17.28%** on RSREC and **+12.62%** on RSRES, highlighting the strong grounding capability of our zero-shot framework. Compared with other zero-shot approaches, *i.e.*, DiffSegmenter (w/ VLM) and DiffPNG (w/ VLM), our method also achieves significant improvements in mIoU by **4.45%**-**5.99%** on RSREC, and **2.82%**-**6.36%** on RSRES, respectively. This highlights the effectiveness of our framework and its related components beyond VLM alone. Furthermore, we compare our method against several SOTA generic VLMs and a remote sensing VLM, as shown in Tab. 5. RSVG-ZeroOV achieves the best performance across all metrics on both RSREC and RSRES, demonstrating superior grounding ability even compared to models tailored for pixel-level reasoning or remote sensing tasks.

Comparison Results on RISBench. We further evaluate RSVG-ZeroOV on the more challenging RISBench, which features longer and more semantically complex referring expressions compared to RRSIS-D. As shown in Tab. 6, RSVG-ZeroOV (w/ Refine) demonstrates exceptional performance, establishing new SOTA benchmarks for Pr@0.5 and mIoU at **38.90%/38.87%** and **31.03%/31.84%** on both tasks, respectively. Notably, the proposed method outperforms the two zero-shot methods, OV-VG and GroundVLP by **+18.8%/18.13%** and **+17.15%/16.27%** in oIoU, respectively, indicating its superior ability to capture object structures and preserve overall shape integrity.

Method		RSREC					RSRES				
		Pr@0.3	Pr@0.5	Pr@0.7	mIoU	oIoU	Pr@0.3	Pr@0.5	Pr@0.7	mIoU	oIoU
W-S	TRIS (ICCV'23)	19.27	8.89	2.86	14.63	15.74	14.23	4.40	0.92	11.46	16.06
	SAG (ICCV'23)	5.83	1.78	0.33	7.12	7.10	8.79	2.55	0.35	9.31	9.36
	QueryMatch (MM'24)	31.79	27.74	23.32	26.72	15.86	31.06	26.68	20.27	24.59	10.77
Z-S	VLM (Baseline)	48.62	35.08	20.26	34.60	30.15	-	-	-	-	-
	DiffSegmenter (TIP'25)	6.03	1.87	0.40	7.14	7.00	2.44	0.46	0.06	4.00	3.74
	DiffSegmenter (w/ VLM)	27.17	15.93	7.98	21.61	19.03	28.65	11.04	2.50	21.37	21.87
	DiffPNG (ECCV'24)	12.14	6.03	2.72	11.10	9.01	13.44	5.79	1.93	11.85	11.79
	DiffPNG (w/ VLM)	34.60	21.54	11.62	26.18	23.58	23.50	8.73	1.76	18.49	14.45
	OV-VG (arXiv'24)	26.83	22.16	18.31	22.65	15.64	-	-	-	-	-
	GroundVLP (AAAI'24)	23.13	19.90	17.15	19.76	16.46	-	-	-	-	-
	RSVG-ZeroOV (Ours)	50.17	35.08	19.03	36.20	34.37	32.71	14.74	3.73	22.69	21.75
Z-S w/ Refine	VLM (Baseline)	49.07	38.09	26.57	37.74	30.68	43.03	30.38	18.04	30.45	21.30
	DiffSegmenter (TIP'25)	6.05	1.99	0.45	7.11	6.96	2.65	0.73	0.17	3.88	3.79
	DiffSegmenter (w/ VLM)	26.74	16.74	9.47	22.01	17.91	21.39	12.02	6.40	17.65	12.87
	DiffPNG (ECCV'24)	9.34	3.76	1.38	8.43	8.66	3.99	1.73	0.76	4.53	5.81
	DiffPNG (w/ VLM)	22.07	12.19	6.00	17.16	17.40	14.74	5.65	1.60	12.19	8.44
	OV-VG (arXiv'24)	26.90	22.40	18.50	22.85	15.50	23.67	17.75	12.25	18.16	9.20
	GroundVLP (AAAI'24)	23.13	19.91	16.37	19.19	16.17	20.78	15.82	10.95	15.58	10.08
	RSVG-ZeroOV (Ours)	50.77	38.90	24.93	38.87	34.30	44.30	31.03	18.61	31.84	26.35

Table 6: Comparisons with SOTA methods on **RISBench** test set. ‘W-S’ and ‘Z-S’ denote the weakly-supervised and zero-shot settings, respectively. ‘w/ VLM’ indicates that the cross-attention from the VLM is used to replace the original cross-attention. ‘w/ Refine’ indicates that post-processing operations are applied to further refine the segmentation results.

Method	RSREC		RSRES	
	Pr@0.5	mIoU	Pr@0.5	mIoU
RSVG-ZeroOV (Ours)	30.15	32.92	12.84	21.85
w/o VLM	16.22	18.82	11.43	15.81
w/o DM	21.49	26.26	1.18	6.15
w/o Evolve	22.63	26.65	10.26	20.56
O-F-E Stage	30.15	32.92	12.84	21.85
O-E-F Stage	27.34	29.51	7.18	15.89

Table 7: Effect of components in RSVG-ZeroOV.

Ablation Study

In this section, we report ablation study results on the RRSIS-D test set to investigate its effectiveness.

Effect of RSVG-ZeroOV’s components. As shown in Tab. 7, removing VLM (replaced by a cross-attention map in DM) leads to significant drops of **14.10%** and **6.04%** in mIoU on both tasks. Excluding DM results in failure on RSRES, as the model loses structural priors necessary for segment object regions. Exp.(w/o Evolve) highlight the benefit of attention evolve module, increasing the Pr@0.5 and mIoU of **7.52%** and **6.27%** on RSREC. The improved performance of O-F-E over O-E-F indicates that embedding the DM’s self-attention into the VLM’s cross-attention better preserves structural and shape information of objects, which facilitates more accurate region aggregation via DFS and suppresses noisy activations.

Effect of self-attention resolutions. As shown in Tab. 8, we investigate the impact of varying resolutions of self-attention

Resolution	RSREC		RSRES	
	Pr@0.5	mIoU	Pr@0.5	mIoU
16	26.16	28.39	12.01	19.98
32	29.84	31.97	12.53	21.11
64	27.70	30.76	10.55	20.13
16, 32	28.23	29.86	12.70	20.41
32, 64	30.15	32.92	12.84	21.85
16, 32, 64	27.86	30.51	11.98	20.36

Table 8: Effect of varying self-attention map resolutions.

maps. We find that using multiple resolutions [32, 64] leads to the best performance, suggesting that multi-scale self-attention maps not only preserve high-resolution details but also effectively capture contextual semantics.

Conclusion

In this paper, we introduce RSVG-ZeroOV, a zero-shot framework that integrates a VLM and a DM via attention maps to tackle the open-vocabulary remote sensing visual grounding. Without additional training, we design a novel overview-focus-evolve strategy to progressively enhance attention toward referred objects. Extensive experiments on existing benchmarks demonstrate that RSVG-ZeroOV achieves SOTA performance, outperforming existing weakly-supervised and zero-shot methods. To the best of our knowledge, RSVG-ZeroOV is the first zero-shot framework for remote sensing visual grounding. We hope this work can potentially inspire researchers to address open-vocabulary remote sensing tasks from a fresh perspective.

Acknowledgments

This work was supported in part by the National Science and Technology Major Project under Grant 2022ZD0117103, in part by the National Natural Science Foundation of China under Grants 62577041, in part by the Outstanding Youth Science Foundation of Shaanxi Province under Grant 2025JC-JCQN-083, in part by the Key Research and Development Program of Shaanxi Province under Grants 2025CY-YBXM-047 and 2024GX-YBXM-140.

References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.
- Chen, K.; Liu, C.; Chen, H.; Zhang, H.; Li, W.; Zou, Z.; and Shi, Z. 2024a. RSPrompter: Learning to prompt for remote sensing instance segmentation based on visual foundation model. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–17.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 24185–24198.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3): 6.
- Ding, Y.; Xu, H.; Wang, D.; Li, K.; and Tian, Y. 2024. Visual Selection and Multistage Reasoning for RSVG. *IEEE Geoscience and Remote Sensing Letters*, 21: 1–5.
- Dong, Z.; Sun, Y.; Liu, T.; Zuo, W.; and Gu, Y. 2024. Cross-modal bidirectional interaction model for referring remote sensing image segmentation. *arXiv preprint arXiv:2410.08613*.
- Huang, S.; He, S.; and Wen, B. 2025. ZoRI: Towards discriminative zero-shot remote sensing instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 3724–3732.
- 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.
- Kuckreja, K.; Danish, M. S.; Naseer, M.; Das, A.; Khan, S.; and Khan, F. S. 2024. Geochat: Grounded large vision-language model for remote sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27831–27840.
- Lan, M.; Rong, F.; Jiao, H.; Gao, Z.; and Zhang, L. 2024. Language query-based transformer with multiscale cross-modal alignment for visual grounding on remote sensing images. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–13.
- Li, K.; Dong, F.; Wang, D.; Li, S.; Wang, Q.; Gao, X.; and Chua, T.-S. 2024a. Show me what and where has changed? question answering and grounding for remote sensing change detection. *arXiv preprint arXiv:2410.23828*.
- Li, K.; Liu, R.; Cao, X.; Bai, X.; Zhou, F.; Meng, D.; and Wang, Z. 2025a. Segearth-ov: Towards training-free open-vocabulary segmentation for remote sensing images. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 10545–10556.
- Li, K.; Wang, D.; Xu, H.; Zhong, H.; and Wang, C. 2024b. Language-Guided Progressive Attention for Visual Grounding in Remote Sensing Images. *IEEE Transactions on Geoscience and Remote Sensing*.
- Li, K.; Xin, Z.; Pang, L.; Pang, C.; Deng, Y.; Yao, J.; Xia, G.; Meng, D.; Wang, Z.; and Cao, X. 2025b. Segearth-r1: Geospatial pixel reasoning via large language model. *arXiv preprint arXiv:2504.09644*.
- Liu, F.; Chen, D.; Guan, Z.; Zhou, X.; Zhu, J.; Ye, Q.; Fu, L.; and Zhou, J. 2024a. Remoteclip: A vision language foundation model for remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–16.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 26296–26306.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, S.; Ma, Y.; Zhang, X.; Wang, H.; Ji, J.; Sun, X.; and Ji, R. 2024c. 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.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024d. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, 38–55. Springer.
- Pan, J.; Liu, Y.; Fu, Y.; Ma, M.; Li, J.; Paudel, D. P.; Van Gool, L.; and Huang, X. 2025. Locate anything on earth: Advancing open-vocabulary object detection for remote sensing community. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 6281–6289.
- Pnvr, K.; Singh, B.; Ghosh, P.; Siddiquie, B.; and Jacobs, D. 2023. Ld-znet: A latent diffusion approach for text-based image segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4157–4168.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

Shan, Z.; Liu, Y.; Zhou, L.; Yan, C.; Wang, H.; and Xie, X. 2025. Ros-sam: High-quality interactive segmentation for remote sensing moving object. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 3625–3635.

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, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; et al. 2024a. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.

Wang, W.; Lv, Q.; Yu, W.; Hong, W.; Qi, J.; Wang, Y.; Ji, J.; Yang, Z.; Zhao, L.; XiXuan, S.; et al. 2024b. Cogvlm: Visual expert for pretrained language models. *Advances in Neural Information Processing Systems*, 37: 121475–121499.

Xu, J.; Liu, S.; Vahdat, A.; Byeon, W.; Wang, X.; and De Mello, S. 2023. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2955–2966.

Yao, L.; Liu, F.; Chen, D.; Zhang, C.; Wang, Y.; Chen, Z.; Xu, W.; Di, S.; and Zheng, Y. 2025. RemoteSAM: Towards Segment Anything for Earth Observation. In *Proceedings of the 33th ACM International Conference on Multimedia*.

Ye, C.; Zhuge, Y.; and Zhang, P. 2025. Towards open-vocabulary remote sensing image semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 9436–9444.

Yuan, Z.; Mou, L.; Hua, Y.; and Zhu, X. X. 2023. Rrsis: Referring remote sensing image segmentation. *arXiv preprint arXiv:2306.08625*.

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.

Zheng, Y.; Wu, W.; Li, Q.; Wang, X.; Zhou, X.; Ren, A.; Shen, J.; Zhao, L.; Li, G.; and Yang, X. 2025. InstructSAM: A Training-Free Framework for Instruction-Oriented Remote Sensing Object Recognition. *arXiv preprint arXiv:2505.15818*.