

# RemoteReasoner: Towards Unifying Geospatial Reasoning Workflow

Liang Yao<sup>1\*</sup>, Fan Liu<sup>1†</sup>, Hongbo Lu<sup>2,3</sup>, Chuanyi Zhang<sup>1</sup>,  
Rui Min<sup>1</sup>, Shengxiang Xu<sup>4</sup>, Shimin Di<sup>4</sup>, Pai Peng<sup>3†</sup>

<sup>1</sup>Hohai University

<sup>2</sup>Shanghai Jiao Tong University

<sup>3</sup>COWARobot Co. Ltd.

<sup>4</sup>Southeast University

fanliu@hhu.edu.cn, penmpai\_sh@163.com

## Abstract

Remote sensing imagery presents vast, inherently unstructured spatial data, necessitating sophisticated reasoning to interpret complex user intents and contextual relationships beyond simple recognition tasks. In this paper, we aim to construct an Earth observation workflow to handle complex queries by reasoning about spatial context and user intent. As a reasoning workflow, it should autonomously explore and construct its own inference paths, rather than being confined to predefined ground-truth sequences. Ideally, its architecture ought to be unified yet generalized, possessing capabilities to perform diverse reasoning tasks through one model without requiring additional fine-tuning. Existing remote sensing approaches rely on supervised fine-tuning paradigms and task-specific heads, limiting both autonomous reasoning and unified generalization. To this end, we propose RemoteReasoner, a unified workflow for geospatial reasoning. The design of RemoteReasoner integrates a multi-modal large language model (MLLM) for interpreting user instructions and localizing targets, together with task transformation strategies that enable multi-granularity tasks, including object-, region-, and pixel-level. In contrast to existing methods, our framework is trained with reinforcement learning (RL) to endow the MLLM sufficient reasoning autonomy. At the inference stage, our transformation strategies enable diverse task output formats without requiring task-specific decoders or further fine-tuning. Experiments demonstrated that RemoteReasoner achieves state-of-the-art performance across multi-granularity reasoning tasks. Furthermore, it retains the MLLM’s inherent generalization capability, demonstrating robust performance on unseen tasks and categories.

**Code** — <https://github.com/1e12Leon/RemoteReasoner>

## 1 Introduction

Remote sensing analysis has evolved beyond simple object detection (Liu et al. 2025a; Li et al. 2025a) or classification (Adegun, Viriri, and Tapamo 2023), increasingly requiring stronger capabilities to support complex real-world

\*This work was done during a research internship at COWARobot. Liang Yao and Fan Liu contribute equally.

†Corresponding Author

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

## Existing Method

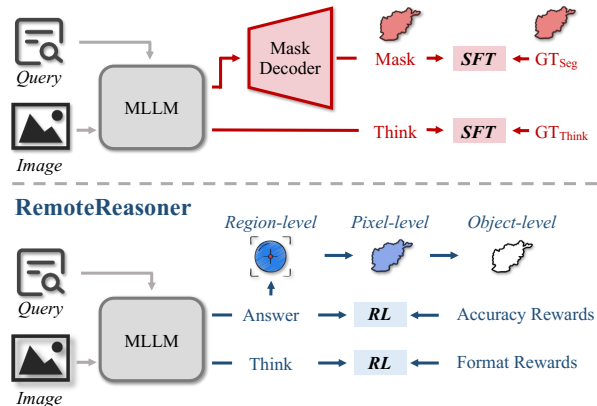


Figure 1: Comparison between Existing Frameworks and RemoteReasoner. Existing remote sensing reasoning approach (e.g., SegEarth-R1) requires SFT with annotated reasoning processes, and is limited to single-task outputs and task decoder. In contrast, our framework supports unsupervised reasoning and multi-granularity tasks.

decision-making. Although existing approaches (Yao et al. 2025b; Liu et al. 2024a) are effective for multiple tasks, they struggle to interpret nuanced user intent and infer implicit spatial relationships for intricate queries. Reasoning (Jaech et al. 2024; Xu et al. 2025) bridges this gap by enabling systems to dynamically decipher the underlying goals of ambiguous or high-level instructions (e.g., “Find areas at risk of flooding near critical infrastructure after heavy rain”). Such a geospatial workflow would conveniently accommodate varied geoscience applications (Zhao et al. 2024).

Numerous efforts (Lai et al. 2024; Ren et al. 2024; Chen et al. 2024b) have explored reasoning tasks in the natural image domain. For instance, LISA (Lai et al. 2024) utilizes multi-modal large language models (MLLMs) (Carolan, Fennelly, and Smeaton 2024) for reasoning-driven segmentation tasks. However, these approaches rely on supervised fine-tuning (SFT) paradigms to train the models. Such frameworks constrain model autonomy, ultimately compromising generalization capabilities. This situation can be

manifested in substantial performance degradation on out-of-distribution samples and catastrophic forgetting of general capabilities (Liu et al. 2025b). Similarly, the remote sensing community has seen pioneering works (Li et al. 2025b; Zhang et al. 2023; Wang et al. 2024). SegEarth-R1 (Li et al. 2025b) introduces a geospatial pixel reasoning task that integrates a hierarchical MLLM and a mask decoder to generate both reasoning processes and segmentation outputs. However, it remains constrained by SFT paradigms and exclusively supports segmentation tasks.

These limitations motivate us to design a reasoning workflow that autonomously interprets user intent through self-directed thinking. It should also efficiently generate geospatial interpretations in diverse output formats, which are tailored for various downstream intelligence analysis. Firstly, the requirement for self-thinking should be fulfilled through an open-ended learning paradigm. This aligns naturally with the emerging reinforcement learning (RL) frameworks (Guo et al. 2025; Shen et al. 2025). Rather than SFT approaches that demand strictly annotated reasoning processes, RL explores the model’s self-evolution potential to cultivate reasoning capabilities. Secondly, to flexibly support multi-granularity tasks, we should avoid redundant task decoders and training through decision-level task transformation.

Followed by these insights, we propose **RemoteReasoner**, a reasoning workflow that integrates a robust MLLM and a flexible inference strategy. As shown in Fig. 1, building upon the *pixel-level* reasoning task proposed in SegEarth-R1 (Li et al. 2025b), RemoteReasoner further implements two novel tasks through autonomous reasoning: 1) *region-level*: geospatial region reasoning, and 2) *object-level*: geospatial contour reasoning. These three distinct tasks can be achieved through a single MLLM inference followed by a unified decision-level task transformation pipeline, significantly improving output efficiency for multi-granularity reasoning tasks.

During training, we utilize Group Relative Policy Optimization (GRPO) (Shao et al. 2024) to fine-tune the reasoning model. Benefiting from our sophisticated reward function that balances output formats, localization accuracy, and quantity, RemoteReasoner demonstrates strong capabilities in autonomous reasoning and precise *region-level* localization. This training approach maintains the inherent generalization capability of MLLM, thereby endowing it with the ability to recognize out-of-distribution categories. During inference, we design a task transformation pipeline inspired by RemoteSAM (Yao et al. 2025b), which enables the model to perform *object-* and *pixel-level* reasoning tasks, eliminating redundant computations across granularities.

Extensive experimental results demonstrate that RemoteReasoner achieves state-of-the-art performance on multi-granularity reasoning tasks, evidencing its enhanced autonomous reasoning capabilities and superior generalization. For instance, RemoteReasoner achieves 3.67% accuracy improvement over SegEarth-R1 in pixel-level geospatial reasoning tasks. On unseen tasks (*e.g.*, Referring Segmentation), our framework significantly outperforms competing models (*e.g.*, PixelLM, LISA) in performance.

The contributions are summarized as follows:

- We propose RemoteReasoner, a unified yet generalized reasoning workflow for Earth observation, which reveals that integrating pure RL and task transformation strategies could handle multi-granularity reasoning tasks.
- We introduce two novel geospatial reasoning tasks, region reasoning and contour reasoning. We also provide the corresponding precisely annotated datasets to facilitate the advancement of remote sensing research.
- Holistic evaluations demonstrate that RemoteReasoner achieves superior generalization capabilities and efficiency in handling multi-granularity reasoning tasks through a single forward pass.

## 2 Related Work

### 2.1 Reasoning in Large Language Models

Recent efforts (Jaech et al. 2024; Guo et al. 2025; Ren et al. 2024; Lai et al. 2024; Li et al. 2025b) have shown that reasoning is central to aligning large language models (LLMs) with complex instructions. Models like OpenAI-o1 (Jaech et al. 2024) and DeepSeek-R1 (Guo et al. 2025) incorporate chain-of-thought (CoT) (Wei et al. 2022) reasoning into their generation process, enabling flexible and interpretable multi-step inference. Both adopt reinforcement learning (RL) (Kaelbling, Littman, and Moore 1996; Li 2017; Arulkumaran et al. 2017) to induce generalizable reasoning behaviors—OpenAI-o1 through RLHF and CoT monitoring, and DeepSeek-R1 via large-scale RL from scratch followed by multi-stage pretraining. These approaches exhibit strong reasoning capabilities but remain limited to natural language and vision tasks.

In remote sensing, models such as PixelLM (Ren et al. 2024), LISA (Lai et al. 2024), and SegEarth-R1 (Li et al. 2025b) extend reasoning to spatial domains. While effective for pixel- or region-level inference, they often support narrow task types and fixed granularity. Notably, SegEarth-R1 relies on supervised fine-tuning, making it less robust to distribution shifts and less adaptable in open-ended settings. These limitations underscore the need for a more versatile geospatial framework.

### 2.2 Remote Sensing Multi-modal Models

Recent remote sensing multimodal models (Liu et al. 2024a; Zhang et al. 2024b; Hu et al. 2025; Ma et al. 2025; Li et al. 2024; ?) demonstrate strong capabilities in handling diverse tasks across spatial granularities. Falcon (Yao et al. 2025a) supports 14 vision-language tasks via multi-resolution encoding; GeoChat (Kuckreja et al. 2024) and LHRS-Bot (Muhtar et al. 2024) incorporate spatial prompts and geographic priors for improved region-level grounding; RemoteSAM (Yao et al. 2025b) adopts a unified encoder for dense pixel-level segmentation; SkyEyeGPT (Zhan, Xiong, and Yuan 2024) reformulates task input/output as natural language for broad task generalization. While effective in geospatial understanding, these models lack robust reasoning capabilities, limiting performance under ambiguous or compositional prompts. While existing models demonstrate strong perception capabilities, they lack the reasoning ability to interpret complex prompts or adapt across granularities.

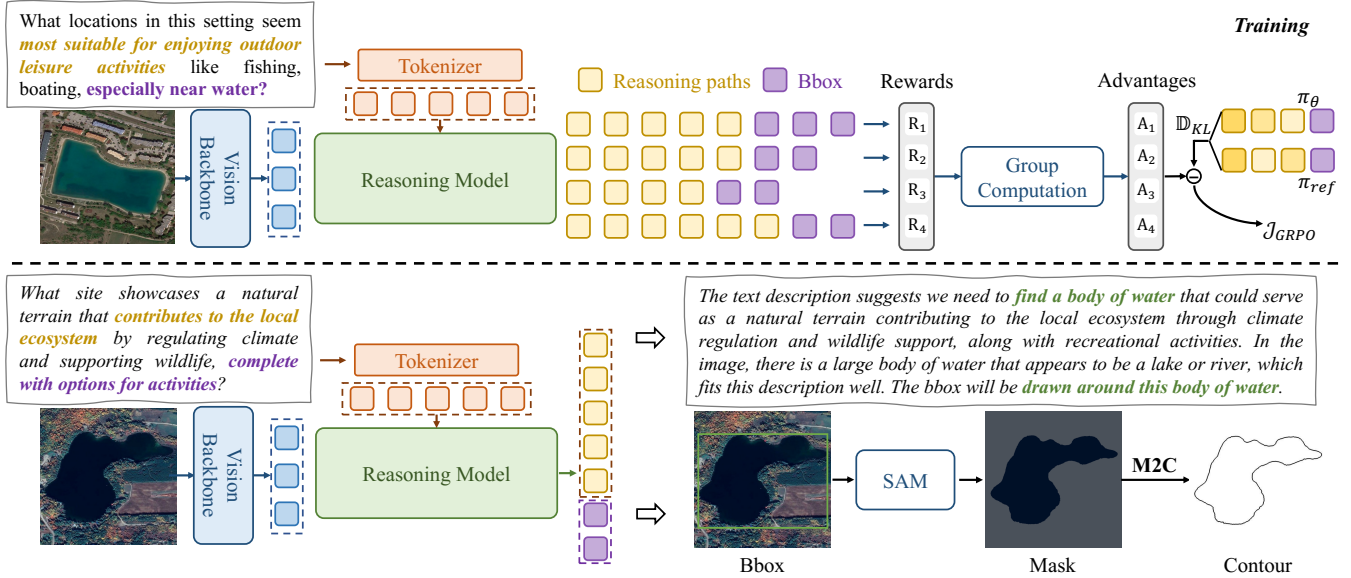


Figure 2: Overview of our RemoteReasoner. We utilize GRPO (Shao et al. 2024) to explore the model’s self-thinking capability. Then we design an inference workflow to perform multi-granularity reasoning tasks.

### 3 RemoteReasoner

In this section, we introduce our method of building a workflow for multi-granularity geospatial reasoning tasks. The overall framework is represented in Fig. 2.

#### 3.1 Data

The EarthReason (Li et al. 2025b) dataset is proposed alongside SegEarth-R1, containing 6 questions per sample with an average length of 20.86 words. It comprises 2,371 training, 1,135 validation, and 1,928 testing samples. However, this dataset is limited to pixel tasks. To address this, we generate detection and contour annotations through Mask2Bbox (M2B) (Liu et al. 2024a) and Mask2Contour (M2C) (Miao et al. 2025) strategies, respectively. Additionally, we design prompt templates combined with original questions and generated bounding boxes to form training samples. The segmentation and contour annotations are exclusively reserved for downstream testing evaluation.

#### 3.2 Reasoning Model

We utilize Qwen2.5-VL-7B-instruct (Bai et al. 2025) as our reasoning model  $\mathcal{F}_\theta$ . Despite Qwen2.5-VL’s superior performance in general-domain applications, direct adaptation to remote sensing imagery introduces domain knowledge discrepancy. Moreover, its reasoning capabilities require significant enhancement to address complex geospatial tasks. Therefore, we require fine-tuning it through reinforcement learning on remote sensing data.

#### 3.3 Training

We utilize Group Relative Policy Optimization (GRPO) (Shao et al. 2024) to train RemoteReasoner’s reasoning model. Unlike traditional reinforcement learning

methods like PPO (Schulman et al. 2017), GRPO directly compares groups of candidate outputs without an additional critic model. Given a query  $Q$ , GRPO samples  $N$  candidate outputs  $o_1, o_2, \dots, o_N$  from the policy  $\pi_\theta$ . Then it computes each output  $o_i$  through a Reward function  $R(Q, o_i)$ . To determine the relative quality of the outputs, GRPO further normalizes the obtained rewards and subsequently derives the advantage:

$$A_i = \frac{r_i - \text{mean}\{r_1, r_2, \dots, r_N\}}{\text{std}\{r_1, r_2, \dots, r_N\}}, \quad (1)$$

where  $A_i$  denotes the advantage of  $o_i$  relative to other sampled outputs. GRPO aims to make the model generate outputs that obtain higher advantages, then utilize them to optimize the policy  $\pi_\theta$ :

$$\begin{aligned} \mathcal{J}_{GRPO}(\theta) = & \mathbb{E}\{\{o_i\}_{i=1}^N \sim \pi_{\theta_{old}}(q)\} \\ & \frac{1}{N} \sum_{i=1}^N \{ \min[c_1 \cdot A_i, c_2 \cdot A_i] \\ & - \beta \mathbb{D}_{KL}[\pi_\theta || \pi_{ref}] \}, \end{aligned} \quad (2)$$

where,  $c_1$  and  $c_2$  can be formulated as follows:

$$\begin{aligned} c_1 = & \frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}, \\ c_2 = & \text{clip}\left(\frac{\pi_\theta(o_i|q)}{\pi_{\theta_{old}}(o_i|q)}, 1 + \epsilon, 1 - \epsilon\right). \end{aligned} \quad (3)$$

To adopt GRPO for our reasoning model, we design a composite reward function to optimize localization precision, object count accuracy, and output format. The total reward  $R$  combines 3 critical components:

- **Accuracy reward** measures bounding box alignment using Intersection-over-Union (IoU) with optimal Hungarian matching. For predicted boxes  $B = \{b_k\}_{k=1}^K$  and ground truth  $G = \{g_j\}_{j=1}^J$ :

$$R_{\text{IoU}} = \frac{1}{|G|} \sum_{j=1}^{|G|} \max_k \left[ \frac{|b_k \cap g_j|}{|b_k \cup g_j|} \right]. \quad (4)$$

- **Count reward** penalizes deviations in object quantity prediction using an exponential decay function:

$$R_{\text{count}} = \begin{cases} 1 & (J = 0) \wedge (K = 0) \\ 0 & (J = 0) \wedge (K > 0), \\ \exp\left(-2 \cdot \frac{|K-J|}{J}\right) & J > 0 \end{cases}, \quad (5)$$

where  $J = |G|$  and  $K = |B|$  denote ground-truth and predicted counts respectively. The exponential term  $(-2 \cdot)$  creates a smooth penalty gradient that becomes stricter as relative count error increases.

- **Format reward**  $R_{\text{format}}$  checks whether RemoteReasoner’s outputs follow the specified format, returning 1 or 0 based on compliance. It is the JSON-style with reasoning paths in the  $\langle \text{think} \rangle \dots \langle / \text{think} \rangle$  tag and bboxes in the  $\langle \text{answer} \rangle [[x_1, y_1, x_2, y_2] \dots] \langle / \text{answer} \rangle$  tag.

Finally, the overall reward function is:

$$R = R_{\text{IoU}} + R_{\text{count}} + R_{\text{format}}. \quad (6)$$

### 3.4 Inference Workflow

After training, the reasoning model itself generates bounding boxes (bbox) indicating target regions. To further enhance the workflow and enable both pixel-level and object-level reasoning tasks, we need to leverage established tools. Specifically, we first input the bbox and its centroid coordinates to the SAM2 (Ravi et al. 2024) to generate pixel-level masks. Subsequently, morphological operations are applied to analyze the mask images, extracting the outermost polygonal vertices to obtain target contours. The mathematical formulation of this process is presented as follows:

Given a remote sensing image  $I \in \mathbb{R}^{H \times W \times 3}$  and a textual query  $Q$ , we define a multi-stage workflow for multi-task reasoning.

- **Geospatial Region Reasoning:** A vision-language model  $\mathcal{F}_\theta$  processes the image-query pair  $(I, Q)$  to jointly produce:

$$(T, b) = \mathcal{F}_\theta(I, Q), \quad (7)$$

where  $T$  denotes the textual reasoning trace, and  $b = [x_{\min}, y_{\min}, x_{\max}, y_{\max}]$  is the normalized bounding box of the target region.

- **Geospatial Pixel Reasoning:** The bounding box  $b$  and its center points  $v$  serve as prompts for the Segment Anything Model (SAM)  $\mathcal{S}$ . SAM generates a binary mask  $M \in \{0, 1\}^{H \times W}$  isolating the target:

$$M = \mathcal{S}(I; b, v), \quad (8)$$

Method	<i>cIoU</i>		<i>gIoU</i>	
	Val	Test	Val	Test
<i>SFT Methods</i>				
LISA	57.39	59.10	61.04	60.88
PixelLM	57.79	59.22	57.94	60.01
PSALM	62.03	64.61	66.61	68.30
SegEarth-R1	64.13	68.25	68.60	70.75
<i>RL Method</i>				
<b>RemoteReasoner</b>	<b>67.80</b>	<b>69.13</b>	<b>69.02</b>	<b>70.96</b>

Table 1: Geospatial Pixel Reasoning Results on EarthReason. It is worth noting that RemoteReasoner was not trained directly utilizing its segmentation labels.

Method	<i>Acc@0.5</i>		<i>gIoU</i>	
	Val	Test	Val	Test
DeepSeek-VL2-tiny	12.08	12.67	17.51	18.62
GeoChat	10.10	8.89	12.57	11.44
Qwen2.5-VL-7B	41.21	45.82	38.77	41.80
<b>RemoteReasoner</b>	<b>66.51</b>	<b>68.11</b>	<b>67.04</b>	<b>69.29</b>

Table 2: Geospatial Region Reasoning Results.

where  $M_{(i,j)} = 1$  indicates pixel  $(i, j)$  belongs to the target.

- **Geospatial Contour Reasoning:** The target contour  $C$  is derived  $M$  via morphological boundary detection:

$$C = \Gamma(M) = \{(x_k, y_k)\}_{k=1}^K, \quad (9)$$

where  $\Gamma(\cdot)$  extracts the polygonal vertices  $\{(x_k, y_k)\}$  of the mask’s external boundary.

## 4 Experiments

### 4.1 Dataset and Metrics

**Dataset.** We evaluate on the validation and test sets of EarthReason (Li et al. 2025b). Pixel-level tasks directly utilize the original annotations, while region-level and object-level tasks are transformed using M2B (Liu et al. 2024a) and M2C (Miao et al. 2025) annotation formats based on the original labels, respectively. Additionally, we employ the RSVG (Zhan, Xiong, and Yuan 2023) test set for Visual Grounding evaluation and the RRSISD (Liu et al. 2024b) test set for Referring Expression Segmentation tasks, with their training sets explicitly excluded from model training. For unseen category evaluation, we manually selected 100 images covering 10 distinct unseen categories from the RemoteSAM-270K (Yao et al. 2025b) dataset.

**Metrics.** Following previous works (Yao et al. 2025b; Li et al. 2025b), we adopt *gIoU*, *cIoU*, and *Acc@0.5* as region- and pixel-level evaluation metrics. For contour reasoning, we followed (Jiao et al. 2025) to employ F1 score with 1-pixel and 3-pixel dilated kernel ( $F1@1$  &  $F1@3$ ) as evaluation metrics, and incorporate two geometric error measures: Average Symmetric Distance (*ASD*) and Hausdorff Distance (*HD*).

Method	$F1@1\uparrow$		$F1@3\uparrow$		$Recall@1\uparrow$		$Recall@3\uparrow$		$ASD\downarrow$		$HD\downarrow$	
	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test	Val	Test
EPOC	0.10	0.10	0.19	0.20	0.14	0.15	0.28	0.29	60.07	64.38	243.98	255.81
<b>RemoteReasoner</b>	<b>0.46</b>	<b>0.44</b>	<b>0.55</b>	<b>0.58</b>	<b>0.40</b>	<b>0.47</b>	<b>0.53</b>	<b>0.64</b>	<b>41.50</b>	<b>45.20</b>	<b>160.56</b>	<b>165.11</b>

Table 3: Geospatial Contour Reasoning Results on EarthReason.

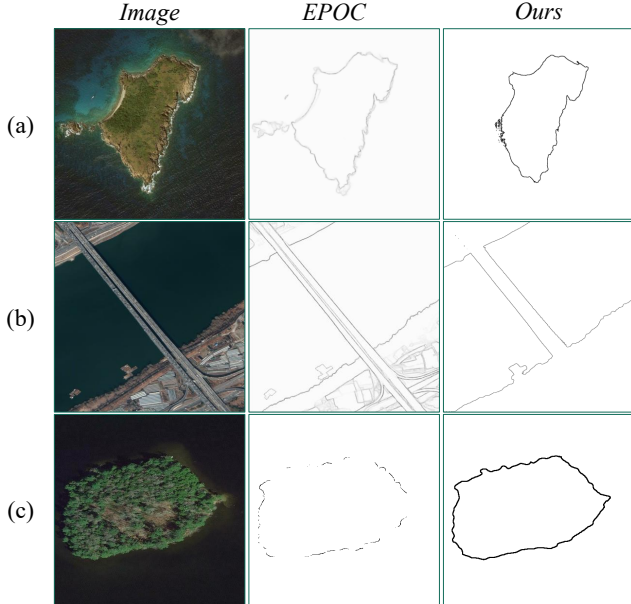


Figure 3: Geospatial Contour Reasoning Results. We select EPOC (Chen et al. 2024a) for comparison.

## 4.2 Implementations Details

We employ GRPO (Shao et al. 2024) to fine-tune the Qwen2.5-VL-7B-Instruct (Bai et al. 2025) model using LoRA (Hu et al. 2022) (rank=8, alpha=16, targeting all linear layers). Training is conducted for 24 epochs on  $8 \times$  NVIDIA L20 GPUs with a global batch size of 512 (achieved via per-device batch size 8 and gradient accumulation steps 8), using the AdamW (Loshchilov and Hutter 2017) optimizer with a learning rate of  $1e-6$ , bf16 precision, and DeepSpeed (Rasley et al. 2020) Zero3 for memory efficiency.

## 4.3 Main Results

**Geospatial Pixel Reasoning** We evaluate RemoteReasoner’s performance on pixel-level reasoning tasks using the EarthReason dataset. We compare against state-of-the-art general-domain methods (LISA (Lai et al. 2024), PixelLM (Ren et al. 2024), PSALM (Zhang et al. 2024a)) and the remote sensing-domain SegEarth-R1 (Li et al. 2025b), with results summarized in Tab. 1. Experimental results demonstrate that despite not directly utilizing pixel-level labels during training, our approach remarkably surpasses SFT-based methods that leverage both pixel annotations and reasoning trajectories. For instance, RemoteReasoner achieves 10.41% higher  $cIoU$  than LISA on the validation

set. Compared to the best-performing SegEarth-R1, we observe a 3.67% improvement. The primary reason lies in the superior robustness of reinforcement learning (RL) frameworks compared to supervised fine-tuning (SFT) paradigms. Even without direct exposure to pixel-level annotations, our model effectively utilizes bounding boxes to prompt the Segment Anything Model (SAM), achieving precise segmentation results through this indirect yet efficient interaction mechanism.

**Geospatial Region Reasoning** This section validates RemoteReasoner’s performance on region-level reasoning tasks, with results presented in Tab. 2. Compared to the remote sensing multi-modal foundation model GeoChat (Kuckreja et al. 2024), RemoteReasoner demonstrates at least a 54.47% improvement in  $Acc@0.5$ . This significant gain stems from GeoChat’s weak reasoning capabilities, which fail to interpret ambiguous queries. Similarly, RemoteReasoner achieves over 22.29% improvement against general-purpose MLLMs, *e.g.*, Qwen-2.5VL-7B (Bai et al. 2025) and DeepSeek-VL2-tiny (3B) (Guo et al. 2025). Although these MLLMs possess basic reasoning abilities, their insufficient domain-specific knowledge in remote sensing leads to suboptimal performance on geospatial region-level reasoning tasks.

**Geospatial Contour Reasoning** This section validates the contour extraction performance of RemoteReasoner. We select EPOC (Chen et al. 2024a) as a comparative baseline. Since EPOC lacks reasoning capabilities, we crop the bounding boxes (bbox) generated by RemoteReasoner and input them into EPOC to obtain contours, as shown in Fig. 3. As EPOC is class-agnostic, it extracts contours for all objects within the region (*e.g.*, Fig. 3 (a) and (b)). In contrast, RemoteReasoner’s contours are derived from segmentation masks, resulting in tighter alignment with target boundaries. Furthermore, under low-resolution conditions (Fig. 3 (c)), EPOC produces blurry or disconnected contours, whereas our method maintains sharp and closed boundaries.

Quantitative results (as shown in Tab. 3) demonstrate that RemoteReasoner surpasses EPOC by at least 0.34 in both F1 scores, with significantly lower distance errors compared to EPOC. This performance gap can be attributed to EPOC’s tendency to extract extraneous contours and its failure to handle indistinct target boundaries effectively.

## 4.4 Further Analysis

After reinforcement learning training, can the reasoning model maintain the original capabilities of the MLLM? To explore this, we evaluate RemoteReasoner’s extrinsic capabilities, including its ability to handle unseen task types and

Task	Method	Acc@0.5	gIoU
VG	<i>Trained</i>		
	MGVLF	76.78	68.04
	Falcon	56.87	-
	SkyEyeGPT	70.50	-
	LHRS-Bot	73.45	-
	<i>Zero-shot</i>		
	<b>RemoteReasoner</b>	<b>50.76</b>	<b>49.85</b>
Task	Method	cIoU	gIoU
RES	<i>Trained</i>		
	GeoGround	-	60.50
	SegEarth-R1	67.56	66.40
	<i>Zero-shot</i>		
	LISA	-	26.78
	PixelLM	-	31.65
	<b>RemoteReasoner</b>	<b>54.29</b>	<b>50.97</b>

Table 4: Performance on unseen tasks. VG and RES are Visual Grounding and Referring Expression Segmentation.

Category	Car	Ship	Field	Apron	Grass
<i>gIoU</i>	21.56	50.03	36.72	54.09	38.27
Category	Trucks	Hangar	Parking	Roof	Warehouses
<i>gIoU</i>	25.58	58.34	30.52	52.85	41.00

Table 5: Out-of-distribution category recognition performance. We manually select 10 novel classes that are absent in EarthReason.

recognize out-of-distribution categories.

**Can RemoteReasoner Do Other Tasks?** We evaluate RemoteReasoner on Visual Grounding and Referring Expression Segmentation tasks, with results summarized in Tab. 4. Comparative experiments are conducted against state-of-the-art LLM-based approaches, including general-domain methods (MiniGPT-V2 (Zhu et al. 2023), LISA (Lai et al. 2024), PixelLM (Ren et al. 2024)) and remote sensing-specific frameworks (Falcon (Yao et al. 2025a), LHRS-Bot (Muhtar et al. 2024), GeoGround (Zhou et al. 2024)). Notably, all remote sensing methods were trained on task-specific data, while RemoteReasoner’s zero-shot performance achieves 80% of their supervised results. Under identical zero-shot settings, our framework achieves gains of 4.12% on Visual Grounding and 19.32% on Referring Expression Segmentation tasks.

Furthermore, as the workflow’s core component is a multi-modal large language model, we verify preservation of its inherent capabilities through multi-turn dialogue evaluation (Fig. 4). The model first performs image captioning, then accurately answers content-based questions about the image. Additional results are provided in the supplementary material. These results demonstrate our framework’s strong

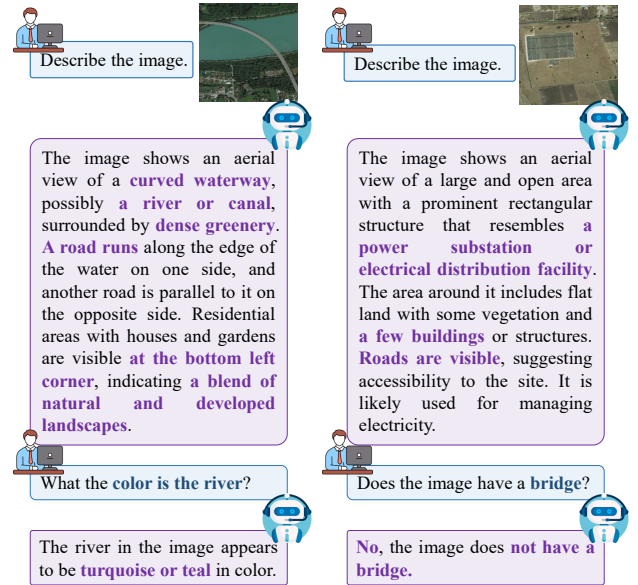


Figure 4: Quantitative results of other tasks (Image Captioning & VQA).

Output Format	Training Method	cIoU	gIoU
Text2Mask	SFT	22.01	22.97
	RL	18.74	19.21
Bbox2Mask	SFT	60.25	62.19
	RL	<b>69.13</b>	<b>70.96</b>

Table 6: Ablation on different pixel-level output formats. Text2Mask denotes LLM-generated textual polygon.

generalization capabilities with RL.

**Can RemoteReasoner Recognize Other Category?** Given that the EarthReason (Li et al. 2025b) dataset contains only 28 categories, which is significantly insufficient to cover common remote sensing targets. Therefore, we further investigate RemoteReasoner’s recognition capability for unseen categories. We manually select 100 images spanning 10 novel categories from RemoteSAM-270K (Yao et al. 2025b) that are absent in EarthReason for evaluation. As represented in Tab. 5, RemoteReasoner demonstrates robust recognition of out-of-distribution categories with consistent accuracy. It demonstrates that our RL strategy avoids overfitting to category distributions in the training data. However, the model still exhibits insufficient capabilities in detecting small objects (e.g., cars), indicating a direction for future research.

**Qualitative Analysis** The qualitative results of the three supported tasks in our workflow are presented in Fig. 5. In the first sample row, the user seeks a venue for group sports, while a baseball field nearly blends into the background in the image. RemoteReasoner correctly reasons and locates the target, simultaneously yielding its segmentation mask and contour. The second row shows a relatively blurred image of a medical helipad. Based on its function of providing

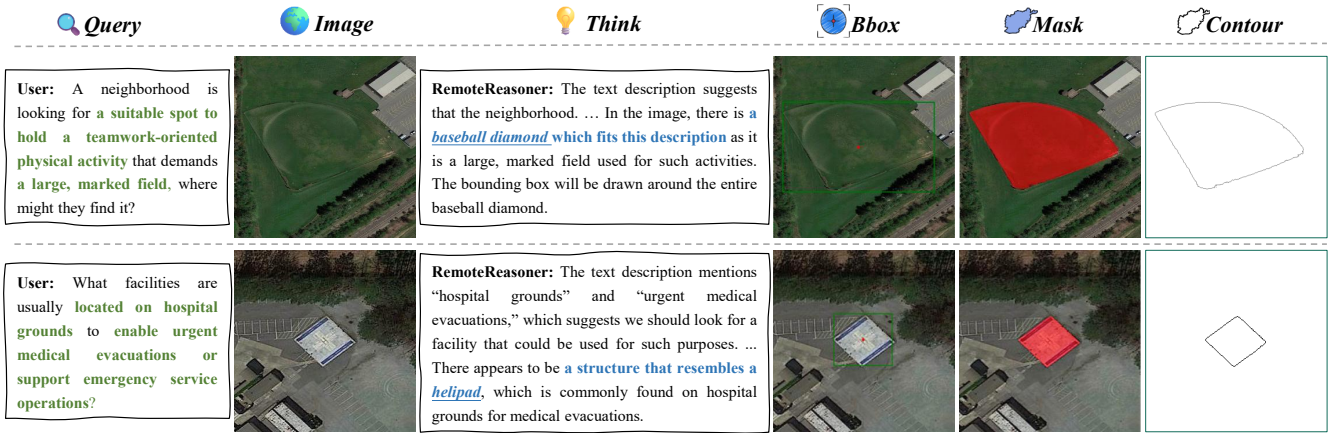


Figure 5: Qualitative Results. Given an image and its corresponding implicit query, RemoteReasoner autonomously identifies the user-intended target category through reasoning and accurately executes visual-centric tasks across three granularities.

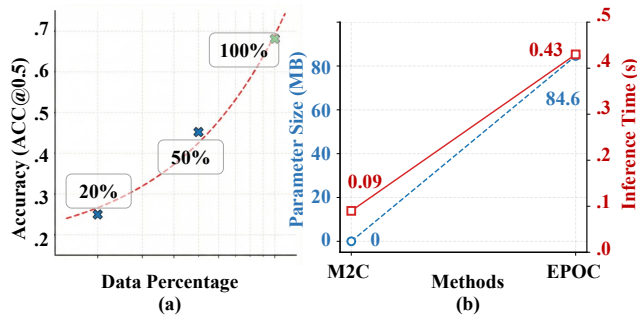


Figure 6: (a) Data scaling experiment. We observed that the model performance increases as the data scales exponentially. (b) Inference cost experiment. Mask2Contour method achieved significantly faster inference speed.

emergency rescue services, RemoteReasoner accurately locates and segments the target area along with its contour. These visualizations indicate that our method can adapt to challenging reasoning tasks.

#### 4.5 Ablation Studies

**Different Pixel-level Output Formats** Recently, RemoteSAM (Yao et al. 2025b) validates that segmentation, as a fundamental visual output unit, enables bottom-up transformation to other vision tasks. Therefore, if LLMs could directly generate masks in textual representations, the architectural design holds promise for greater unification. To explore this hypothesis, we conducted experiments using such formulations, as summarized in Tab. 6. However, textual outputs prove ill-suited for this predictive task, with both SFT and RL strategies exhibiting significant performance degradation compared to visual-based approaches. In contrast, directly outputting bounding boxes and then generating masks through SAM significantly outperforms textual mask representations (22.01 vs. 60.25). When incorporating RL strategies, the accuracy further improves from 60.25 to

69.13. Based on these findings, we adopt the Mask2Contour approach in our workflow design.

**Data Scaling Experiment** To validate data impact on model performance, we evaluate it by training with 20%, 50%, and 100% subsets of the full training data. As illustrated in Fig. 6 (a), the results demonstrate that model performance improves with increasing data scale. The full dataset (100%) achieves the best performance, indicating the current data scale contains no redundancy. These findings suggest that geospatial reasoning tasks retain untapped potential that can be unlocked through larger data scales.

**Inference Cost Ablation** We further evaluate inference speed across different contour extraction strategies, as shown in Fig. 6 (b). The Mask2Contour approach directly applies morphological operations without introducing additional parameters, achieving significantly faster processing speed (0.09s per image). In contrast, EPOC, as a deep learning model with 84.6MB parameters, demonstrates inferior inference speed (0.43s per image). This comparison highlights the efficiency advantages of our proposed workflow.

## 5 Conclusion

We present RemoteReasoner, a novel workflow advancing autonomous geospatial reasoning. Overcoming limitations of SFT-based approaches, it leverages RL strategy to cultivate robust reasoning capabilities while preserving MLLM generalization. Its core innovation is a unified inference pipeline enabling efficient multi-granularity task transformation (pixel-, region-, and object-level) from a single MLLM output, eliminating redundant per-task decoders. RemoteReasoner achieves SOTA results, including  $> 15\%$  and  $30\%$  accuracy gains in region- and object-level reasoning, and demonstrates superior generalization across diverse, complex queries. It provides a flexible foundation for downstream geospatial intelligence requiring nuanced intent interpretation and multi-format outputs.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62372155 and Grant 62302149, in part by Basic Research Program of Jiangsu under Grant BK20250188, in part by the Research Funds of Jiangsu Hydraulic Research Institute under Grant 2025z065, and in part by Changzhou Science and Technology Bureau Project No. 20231313, and in part by Research and Application of the Next-Generation General-Purpose Intelligent Robot Brain (Robo-GPT).

## References

- Adegun, A. A.; Viriri, S.; and Tapamo, J.-R. 2023. Review of deep learning methods for remote sensing satellite images classification: experimental survey and comparative analysis. *Journal of Big Data*, 10(1): 93.
- Arulkumaran, K.; Deisenroth, M. P.; Brundage, M.; and Bharath, A. A. 2017. Deep reinforcement learning: A brief survey. *IEEE SPM*, 34(6): 26–38.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-v1 technical report. *arXiv preprint arXiv:2502.13923*.
- Carolan, K.; Fennelly, L.; and Smeaton, A. F. 2024. A review of multi-modal large language and vision models. *arXiv preprint arXiv:2404.01322*.
- Chen, D.; Cahyawijaya, S.; Liu, J.; Wang, B.; and Fung, P. 2024a. Subobject-level image tokenization. *arXiv preprint arXiv:2402.14327*.
- Chen, Y.-C.; Li, W.-H.; Sun, C.; Wang, Y.-C. F.; and Chen, C.-S. 2024b. Sam4mllm: Enhance multi-modal large language model for referring expression segmentation. In *ECCV*, 323–340. Springer.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, Y.; Yuan, J.; Wen, C.; Lu, X.; Liu, Y.; and Li, X. 2025. Rsgpt: A remote sensing vision language model and benchmark. *ISPRS Journal of Photogrammetry and Remote Sensing*, 224: 272–286.
- Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A.; et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Jiao, W.; Cheng, H.; Vosselman, G.; and Persello, C. 2025. RolPoly: Vectorized building outline extraction using vertex and logit embeddings. *ISPRS Journal of Photogrammetry and Remote Sensing*, 224: 317–328.
- Kaelbling, L. P.; Littman, M. L.; and Moore, A. W. 1996. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4: 237–285.
- 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 *CVPR*, 27831–27840.
- Lai, X.; Tian, Z.; Chen, Y.; Li, Y.; Yuan, Y.; Liu, S.; and Jia, J. 2024. Lisa: Reasoning segmentation via large language model. In *CVPR*, 9579–9589.
- Li, K.; Dong, F.; Wang, D.; Li, S.; Wang, Q.; Gao, X.; and Chua, T.-S. 2024. Show me what and where has changed? question answering and grounding for remote sensing change detection. *arXiv preprint arXiv:2410.23828*.
- Li, K.; Wang, D.; Wang, T.; Dong, F.; Zhang, Y.; Zhang, L.; Wang, X.; Li, S.; and Wang, Q. 2025a. RSVG-ZeroOV: Exploring a Training-Free Framework for Zero-Shot Open-Vocabulary Visual Grounding in Remote Sensing Images. *arXiv preprint arXiv:2509.18711*.
- 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*.
- Li, Y. 2017. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*.
- 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 TGRS*, 62: 1–16.
- Liu, F.; Yao, L.; Zhang, C.; Wu, T.; Zhang, X.; Jiang, X.; and Zhou, J. 2025a. Boost UAV-based Object Detection via Scale-Invariant Feature Disentanglement and Adversarial Learning. *IEEE TGRS*.
- 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 *CVPR*, 26658–26668.
- Liu, Y.; Peng, B.; Zhong, Z.; Yue, Z.; Lu, F.; Yu, B.; and Jia, J. 2025b. Seg-zero: Reasoning-chain guided segmentation via cognitive reinforcement. *arXiv preprint arXiv:2503.06520*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ma, X.; Zhang, X.; Pun, M.-O.; and Huang, B. 2025. A Unified Framework with Multimodal Fine-tuning for Remote Sensing Semantic Segmentation. *IEEE TGRS*.
- Miao, S.; Chen, D.; Liu, F.; Zhang, C.; Gu, Y.; Guo, S.; and Zhou, J. 2025. Prompting directsam for semantic contour extraction in remote sensing images. In *ICASSP 2025*, 1–5. IEEE.
- Muhtar, D.; Li, Z.; Gu, F.; Zhang, X.; and Xiao, P. 2024. Lhrs-bot: Empowering remote sensing with vgi-enhanced large multimodal language model. In *ECCV*, 440–457. Springer.
- Rasley, J.; Rajbhandari, S.; Ruwase, O.; and He, Y. 2020. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 3505–3506.

- 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*.
- Ren, Z.; Huang, Z.; Wei, Y.; Zhao, Y.; Fu, D.; Feng, J.; and Jin, X. 2024. Pixellm: Pixel reasoning with large multimodal model. In *CVPR*, 26374–26383.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shen, H.; Liu, P.; Li, J.; Fang, C.; Ma, Y.; Liao, J.; Shen, Q.; Zhang, Z.; Zhao, K.; Zhang, Q.; et al. 2025. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Wang, J.; Zheng, Z.; Chen, Z.; Ma, A.; and Zhong, Y. 2024. Earthvqa: Towards queryable earth via relational reasoning-based remote sensing visual question answering. In *AAAI*, volume 38, 5481–5489.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *NeurIPS*, 35: 24824–24837.
- Xu, F.; Hao, Q.; Zong, Z.; Wang, J.; Zhang, Y.; Wang, J.; Lan, X.; Gong, J.; Ouyang, T.; Meng, F.; et al. 2025. Towards large reasoning models: A survey of reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*.
- Yao, K.; Xu, N.; Yang, R.; Xu, Y.; Gao, Z.; Kitrungrotsakul, T.; Ren, Y.; Zhang, P.; Wang, J.; Wei, N.; et al. 2025a. Falcon: A Remote Sensing Vision-Language Foundation Model. *arXiv preprint arXiv:2503.11070*.
- Yao, L.; Liu, F.; Chen, D.; Zhang, C.; Wang, Y.; Chen, Z.; Xu, W.; Di, S.; and Zheng, Y. 2025b. RemoteSAM: Towards Segment Anything for Earth Observation. *arXiv preprint arXiv:2505.18022*.
- Zhan, Y.; Xiong, Z.; and Yuan, Y. 2023. Rsvg: Exploring data and models for visual grounding on remote sensing data. *IEEE TGRS*, 61: 1–13.
- Zhan, Y.; Xiong, Z.; and Yuan, Y. 2024. SkyEyeGPT: Unifying Remote Sensing Vision-Language Tasks via Instruction Tuning with Large Language Model. *arXiv:2401.09712*.
- Zhang, Z.; Jiao, L.; Li, L.; Liu, X.; Chen, P.; Liu, F.; Li, Y.; and Guo, Z. 2023. A spatial hierarchical reasoning network for remote sensing visual question answering. *IEEE TGRS*, 61: 1–15.
- Zhang, Z.; Ma, Y.; Zhang, E.; and Bai, X. 2024a. Psalm: Pixelwise segmentation with large multi-modal model. In *ECCV*, 74–91. Springer.
- Zhang, Z.; Zhao, T.; Guo, Y.; and Yin, J. 2024b. Rs5m and georsclip: A large scale vision-language dataset and a large vision-language model for remote sensing. *IEEE TGRS*.
- Zhao, T.; Wang, S.; Ouyang, C.; Chen, M.; Liu, C.; Zhang, J.; Yu, L.; Wang, F.; Xie, Y.; Li, J.; et al. 2024. Artificial intelligence for geoscience: Progress, challenges, and perspectives. *The Innovation*, 5(5).
- Zhou, Y.; Lan, M.; Li, X.; Ke, Y.; Jiang, X.; Feng, L.; and Zhang, W. 2024. Geoground: A unified large vision-language model for remote sensing visual grounding. *arXiv preprint arXiv:2411.11904*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.