
GeoPixel : Pixel Grounding Large Multimodal Model in Remote Sensing

Akashah Shabbir¹ Mohammed Zumri¹ Mohammed Bennamoun² Fahad Shahbaz Khan^{1,3} Salman Khan^{1,4}

Abstract

Recent advances in large multimodal models (LMMs) have recognized fine-grained grounding as an imperative factor of visual understanding and dialogue. However, the benefits of such representation in LMMs are limited to the natural image domain, and these models perform poorly for remote sensing (RS). The distinct overhead viewpoint, scale variation, and presence of small objects in high-resolution RS imagery present a unique challenge in region-level comprehension. Moreover, the development of the grounding conversation capability of LMMs within RS is hindered by the lack of granular, RS domain-specific grounded data. Addressing these limitations, we propose GeoPixel - the first end-to-end high-resolution RS-LMM that supports pixel-level grounding. This capability allows fine-grained visual perception by generating interleaved masks in conversation. GeoPixel supports up to 4K HD resolution in any aspect ratio, ideal for high-precision RS image analysis. To support the grounded conversation generation (GCG) in RS imagery, we curate a visually grounded dataset GeoPixelD through a semi-automated pipeline that utilizes set-of-marks prompting and spatial priors tailored for RS data to methodically control the data generation process. GeoPixel demonstrates superior performance in pixel-level comprehension, surpassing existing LMMs in both single-target and multi-target segmentation tasks. Our methodological ablation studies validate the effectiveness of each component in the overall architecture. <https://github.com/mbzuai-oryx/GeoPixel>.

¹Mohamed bin Zayed University of Artificial Intelligence

²The University of Western Australia ³Linköping University

⁴Australian National University. Correspondence to: Akashah Shabbir <akashah.shabbir@mbzuai.ac.ae>, Mohammed Zumri <mohammed.zumri@mbzuai.ac.ae>.

Can you give a thorough description of this image, including interleaved segmentation masks to highlight key objects?



Figure 1. An example of visually grounded detailed descriptions generated by the proposed GeoPixel, highlighting its ability to interpret and segment high-resolution remote sensing imagery with fine-grained precision. The model applies distinct masks to key objects (ground track field, swimming pool, soccer field) and semantic mask to smaller objects (vehicles). It effectively identifies spatial positions (center, top) and relationships (within the sports complex) while distinguishing between the global context (buildings, roads, green spaces) and localized structures.

1. Introduction

Recent large multimodal models (LMMs) (Liu et al., 2024a; Dai et al., 2023; Bai et al., 2023b; Chen et al., 2024b) have utilized the foundational capabilities of Large Language Models (LLMs) (Touvron et al., 2023; Chiang et al., 2023; Javaheripi et al.; Bai et al., 2023a) and successfully expanded their horizon to the visual modality with promising capabilities. These LMMs can not only perform visual recognition, but also excel in advanced perception and reasoning required for vision-language tasks such as visual question answers, image captioning, visual grounding, and referring expression segmentation. Grounding LMMs (Rasheed et al., 2024; Ma et al., 2025; Zhao et al., 2023) have further advanced the fine-grained context-aware interpretation of complex visual information by allowing textual outputs to be associated with object instances. Facilitated by large-scale data in the natural images domain, grounding multimodal models pre-trained on extensive datasets have shown impressive capabilities, achieving performance levels comparable to specialist models.

However, with increasing granularity of vision and language understanding, these general domain models exhibit significant limitations in adequately supporting complex earth observation tasks. The performance degradation is influ-

Table 1. Comparison of remote sensing large multimodal models (RS-LMMs), focusing on their grounding capabilities. The ‘Region Output’ column highlights the model’s ability to associate objects with specific spatial regions. Existing models primarily utilize LLMs to generate bounding box coordinates for object grounding. However, none of the current RS-LMMs possess the capability for ‘pixel grounding’, i.e., generating detailed segmentation masks, which are crucial for fine-grained spatial interpretation.

MODELS	RESOLUTION	IMAGE	REGION OUTPUT	REGION DECODER	PIXEL GROUNDING
RSGPT (HU ET AL., 2023)	224 × 224	✓	×	×	×
H2RSVLM (PANG ET AL., 2024)	336 × 336	✓	×	×	×
RS-LLAVA (BAZI ET AL., 2024)	336 × 336	✓	×	×	×
GEOCHAT (KUCKREJA ET AL., 2024)	504 × 504	✓	✓	×	×
SKYEYEPT (ZHAN ET AL., 2024)	448 × 448	✓	✓	×	×
EARTHGPT (ZHANG ET AL., 2024C)	-	✓	✓	×	×
LHRS-BOT (MUHTAR ET AL., 2024)	224 × 224	✓	✓	×	×
SKYSENSEGPT (LUO ET AL., 2024)	504 × 504	✓	✓	×	×
GEOPIXEL	DYNAMIC UPTO 4K	✓	✓	✓	✓

enced not only by the unique vantage point inherent to remote sensing (RS) images but also by large variations in the objects’ size and orientation. Moreover, in high-resolution (HR) remote sensing imagery, objects of interest may exhibit challenging-to-segment spatial footprints, such as narrow bridges that connect urban landscapes and play a critical role in city traffic planning, adding further complexity to the task.

Existing vision language models in RS (Luo et al., 2024; Zhang et al., 2024b; Kuckreja et al., 2024) use quantized coordinates in the form of bounding boxes to localize and ground objects in their response. Such a representation structure is not adequate to associate correct object semantics and also adds a computational burden to the LLM that scales with the number of distinguishable objects. Moreover, monitoring the geospatial environment and its entities demands a broader spatial perspective, now increasingly achievable through advancements in RS technologies that provide HR imagery. However, despite the availability of such rich data, current LMMs in RS struggle to fully exploit this spatial detail. These models often struggle with suboptimal resolution capabilities, hindering their ability to capture the intricate patterns present in high-resolution RS images. In addition, existing RS datasets often lack fine-grained spatial association between objects and their corresponding linguistic descriptions.

To address these issues, we present GeoPixel, a model that can generate a detailed natural language response for a high-resolution RS image with corresponding geospatial object segmentation masks. Our contributions are as follows:

- Our proposed LMM, GeoPixel, is explicitly designed for high-resolution RS image analysis with advanced multi-target pixel grounding capability. Our model adaptively divides the input images into local and global regions, enabling efficient encoding and analysis

by accommodating up to 4k resolution.

- We create GeoPixelD, a multi-modal grounded conversation generation (GCG) dataset comprising 53,816 grounded phrases linked to 600,817 object masks, specifically tailored for RS image understanding. Extensively granular annotations are created with segmentation masks through a semi-automated, scalable pipeline that integrates prior-informed visual prompting with state-of-the-art LMMs and ensures quality via rigorous verification and filtering steps.
- We introduce a comprehensive benchmark for evaluating RS LMMs in fine-grained visual understanding, assessing models’ ability to interpret complex, spatially grounded information. It comprises 5,427 manually validated referring expressions and segmentation masks, covering 61,384 annotated objects.

2. Related Work

Large Multimodal Models (LMMs): LMMs build on the success of LLMs to acquire vision capabilities. Pioneer works such as LLaVA (Liu et al., 2024b), MiniGPT-4 (Zhu et al., 2023), InstructBLIP (Dai et al., 2023) and mPLUG-Owl (Ye et al., 2023b) aligned visual features with language representations through a vision language connector, enhanced by instruction tuning to improve multimodal integration. Improving beyond image-level understanding, models such as GPT4RoI (Zhang et al., 2023), InternGPT (Liu et al., 2023b) and RegionGPT (Guo et al., 2024) introduce regional understanding by allowing inputs such as points, masks, and bounding boxes. Some models feed image coordinates directly into the language model, while others employ additional feature extraction modules to represent specific image regions’ features effectively.

Grounding LMMs: Region-level comprehension is further expanded by models such as Kosmos-2 (Peng et al., 2024), Ferret (You et al., 2023), Shikra (Chen et al., 2023), Pink (Xuan et al., 2024) and LION (Chen et al., 2024a) that allow for the precise location of objects in their outputs based on textual descriptions, a capability known as grounding. These models localize objects on a coarse scale using bounding boxes. Recent models (Lai et al., 2024; Rasheed et al., 2024; Xia et al., 2024; Ren et al., 2024; Zhang et al., 2024d; Liu et al., 2023a) focus on achieving more fine-grained visual and linguistic semantic alignment, by exploring pixel grounding. LISA (Lai et al., 2024), PixelLM (Ren et al., 2024) and GLaMM (Rasheed et al., 2024) incorporate a [SEG] token into the LLM’s vocabulary, leveraging its corresponding token embedding as a conditioning input for SAM (Kirillov et al., 2023) to enable segmentation. Additionally, GSVA (Xia et al., 2024) introduces a [REJ] token to explicitly learn to reject specified targets. Whereas

Llava-plus (Liu et al., 2023a) employs LLMs as agents to assign tasks to the segmentation expert.

Our work aligns with pixel-grounding approaches, such as those in (Lai et al., 2024; Ren et al., 2024; Rasheed et al., 2024). However, these models do not interpret the distinct top-down perspective and cannot differentiate complex spatial arrangements of remote sensing (RS) imagery. In addition, the models’ restricted input size, typically limited to dimensions such as 224×224, exacerbates this issue by constraining the field of view and spatial perception.

High-Resolution Understanding: Vision encoders, such as CLIP ViT (Radford et al., 2021), are widely utilized for various vision tasks but are typically constrained by low resolution (e.g. 224×224) restricting their applicability in HR scenarios. To address this limitation, some approaches (Dosovitskiy et al., 2021; Bai et al., 2023b; Li et al., 2023) scale positional encodings within the CLIP model through interpolation to accommodate larger input sizes, while others such as CogAgent (Hong et al., 2024) and Vary (Wei et al., 2025), employ an additional HR branch. Models such as Monkey (Li et al., 2024b), SPHINX (Lin et al., 2023), Llava-Next (Liu et al., 2024a), IXC2.5 (Zhang et al., 2024a), Textmonkey (Liu et al., 2024d) and Ureader (Ye et al., 2023a) divide the image into grids to enhance performance on HR text-centric tasks.

Remote Sensing (RS) LLMs: RSGPT (Hu et al., 2023) pioneered RS-based natural language conversation and generated detailed captions, followed by GeoChat (Kuckreja et al., 2024) which supported region-specific inputs and visual grounding through oriented bounding box coordinates. SkyEyeGPT (Zhan et al., 2024) extended its functionality to RS video captioning, while EarthGPT (Zhang et al., 2024c) and EarthDial (Soni et al., 2024) integrated multisensor RS interpretation. RS-LLaVA (Bazi et al., 2024) and H2RSVLM (Pang et al., 2024) further improved the RS data interpretation, with H2RSVLM uniquely recognizing and rejecting unanswerable questions. SkySenseGPT (Luo et al., 2024) enables image-level scene graph generation and relation reasoning, while LHRS-Bot (Muhtar et al., 2024) enhances multilevel vision-language alignment and TeoChat (Irvin et al., 2025) converses on temporal sequences. However, these models operate at low resolution and lack pixel-level understanding and grounding.

GeoGround (Zhou et al., 2024), RSUniVLM (Liu & Lian, 2024), and GeoPix (Ou et al., 2025) are works concurrent to ours and share similarities. GeoGround and RSUniVLM support pixel-level grounding by converting masks into text sequences, adding a computational burden to the LLM that scales with the number of distinguishable objects. Whereas, GeoPix incorporated both referring expressions segmentation and visual grounding/detection using box coordinates in text, which leads to task confusion.

3. Method

In the current remote sensing landscape, LMMs face significant limitations in terms of grounding and resolution capabilities (as seen in Table 1). Specifically, the outputs generated by these models lack precise spatial and semantic association with the imagery, leading to either ungrounded or only coarsely grounded text. Furthermore, most LMMs operate on relatively low-resolution data, which restricts their ability to perform fine-scale analysis essential for RS tasks such as detailed land use and transportation network extraction, infrastructure mapping, damage assessment, and environmental monitoring. To address these limitations, we present GeoPixel, a model designed to interpret high-resolution remote sensing images and generate finely detailed, pixel-grounded outputs that encompass multiple target objects.

3.1. GeoPixel Architecture Overview

GeoPixel primarily consists of 5 components (see Figure 2). (1) Adaptive Image Divider (2) Vision Encoder (3) Large Language Model (4) Grounding Vision Encoder (5) Pixel Decoder. The first three components are discussed in Section 3.2, while the latter two in Section 3.3. Jointly, these modules enable high-resolution perception, fine-grained interpretation, and grounding, as detailed below.

3.2. High Resolution Understanding

For high resolution, we adopt the dynamic image partitioning strategy of IXC-2.5 (Zhang et al., 2024a). Initially, the adaptive image divider processes the input image x_{img} , with dimensions $[h_i \times w_i]$, by up-scaling and padding it to align with the closest grid size denoted as $[g_h \times g_w]$.

$$g_h = k_1 \times \mathcal{B}, \quad g_w = k_2 \times \mathcal{B}, \quad (1)$$

$$\text{s.t., } k_1, k_2 \in \mathbb{N}, \quad k_1 \times k_2 \leq \mathcal{P}$$

where \mathcal{B} is the base resolution of the vision encoder and \mathcal{P} is the number of maximum allowable image patches. Subsequently, the image is divided into $k_1 \times k_2$ non-overlapping patches $x_{p_{i,j}}$, where $p = 0, 1, 2, \dots, (k_1 \times k_2 - 1)$, and i, j denote the row and column indices of each patch in the grid.

We employ the scaled CLIP ViT-L/14 (Zhang et al., 2024a) as our vision encoder (\mathcal{I}), with a base resolution of $\mathcal{B} = 560$, facilitating large patches for enhanced visual representation. Furthermore, a global view x_{glob} is generated by resizing x_{img} to a fixed dimension of 560×560 , aligned with the base resolution \mathcal{B} . Feature embeddings of patches $f_{p_{i,j}}$ are appended with a learnable token at the end of each row before flattening and merging (Dong et al., 2024b). Finally, global features f_{glob} and patch features f_p are concatenated (\parallel) with a special separator (s_g) inserted between them

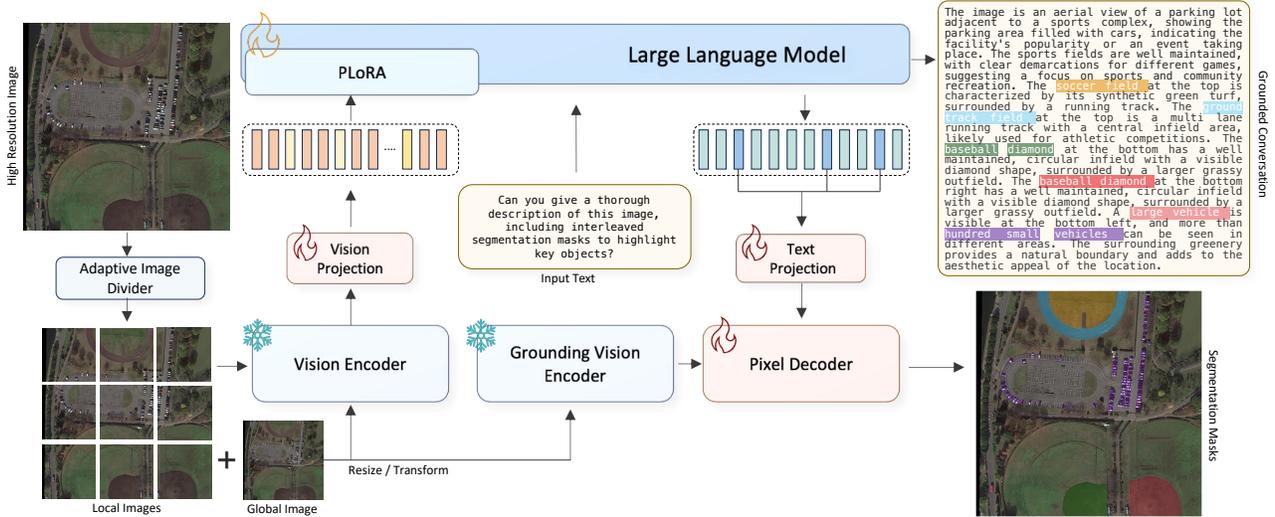


Figure 2. Overview of GeoPixel Architecture: Left: High-resolution RS images are dynamically partitioned into local patches and a resized global view, encoded by a frozen vision encoder. The encodings are projected into the language domain with separator tokens. Middle: Vision tokens, combined with text, are input into the LLM, where pLoRA is applied to vision tokens for efficient and effective multimodal alignment. Right: The corresponding embeddings for the [SEG] tokens are passed to a decoder through a text projector, along with vision embeddings from the grounding vision encoders, to generate precise segmentation masks.

(Ding et al., 2019), effectively integrating global semantics with fine-grained local details.

$$x_v = \mathcal{P}_v(f_{glob} || s_g || f_p) \quad (2)$$

$$\text{s.t. } f_{glob} = \mathcal{I}(x_{glob}), f_{p_{i,j}} = \mathcal{I}(x_{p_{i,j}})$$

We project the final unified image features onto the LLM, InternLM2 7B model (Cai et al., 2024), denoted as \mathcal{L} , through a two-layer MLP as a vision projector \mathcal{P}_v . InternLM2 is an LLM designed to process sequences of text tokens, where its input consists of discrete embeddings derived from textual data. These embeddings correspond to either natural language tokens or special placeholders inserted to represent external modalities. The placeholder $\langle \text{IMAGE} \rangle$ in the input text query x_t is a special token that represents the position of the image within the input sequence. When processing multimodal input, this placeholder is replaced with visual features x_v , extracted from the image, and projected into the same embedding space using \mathcal{P}_v .

Partial Low-Rank Adaptation (LoRA) (Dong et al., 2024a) is then applied to ensure efficient alignment of the vision tokens. Partial LoRA is a modality-specific plug-in module designed to align features from a new modality with LLM, preserving the model’s inherent capabilities while enriching it with modality-specific insights. By applying low-rank adaptations selectively to visual tokens, Partial LoRA enhances alignment efficiency while reducing the computational cost. Formally, it introduces low-rank matrices $W_A \in \mathbb{R}^{C_r \times C_{in}}$ and $W_B \in \mathbb{R}^{C_{out} \times C_r}$ within each LLM linear layer, modifying the visual token outputs x_v

without altering the language token outputs x_t , thus achieving tailored cross-modal integration.

3.3. Pixel Grounding

To establish grounding in LMM, we initialize the grounding vision encoder (\mathcal{I}_g) with a pre-trained SAM-2 (Ravi et al., 2024) encoder together with a dedicated pixel decoder module (\mathcal{D}). The SAM2 visual encoder is a Masked Autoencoder (MAE) (He et al., 2022) pre-trained Hiera (Ryali et al., 2023) image encoder having a hierarchical structure that allows the use of multiscale features during decoding. The tokenizer’s vocabulary is expanded by incorporating an additional $\langle \text{SEG} \rangle$ token, with its corresponding last-layer embedding (E) mapped to the decoder through a text projection layer \mathcal{P}_t . The text projection is a two-layer MLP that receives embeddings of dimension 4096 and transforms them into the input space of the pixel decoder, which has a dimensionality of 256.

The pixel decoder processes the image features from the frozen grounding vision encoder, along with projected LLM embeddings, to generate segmentation masks (M). The grounding vision encoder (SAM-2) is already pre-trained on large-scale datasets, making it highly effective at extracting robust, generalized image features for segmentation. Freezing the encoder ensures that these pretrained features are preserved. However, the light-weight pixel decoder and projection layer are trained to adapt pretrained vision features for segmentation tasks in GeoPixel.

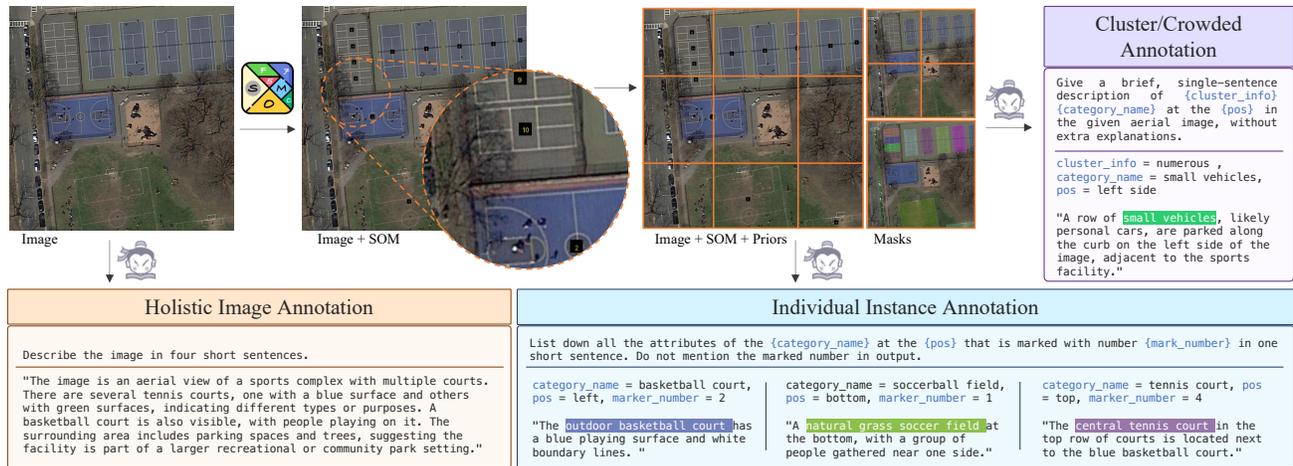


Figure 3. The GeoPixelD Annotation Pipeline provides detailed multi-tier descriptions of remote sensing imagery with object phrases aligned precisely with manually annotated masks. It begins with Holistic Image Annotation (bottom left), where an LMM generates concise scene descriptions. Individual Instance Annotation (bottom right) uses spatial({pos}) and categorical ({category_name}) priors with SOM ({mark_number}) prompting to describe key objects. Cluster Annotation (top right) organizes smaller or dense objects using refined grids for precise spatial analysis.

$$M = \mathcal{D}[\mathcal{I}_g(x_{img}), \mathcal{P}_t(E)] \quad (3)$$

Given the variable length of the input image tokens, resulting from adaptive image partitioning, the output embedding mask for <SEG> tokens is dynamically adjusted to align with these variations. This configuration ensures accurate detection of the <SEG> token and its associated embedding.

4. GeoPixelD-RS Pixel Grounding Dataset

Remote sensing imagery captures intricate semantic information and complex inter-object relationships across diverse spatial scales. To enable LMMs to acquire a detailed comprehension ability, it is essential to integrate broad contextual views with object-level distinction. Addressing the current deficit in datasets capable of facilitating a fine-grained understanding of top-down perspectives, we introduce GeoPixelD, a dataset established to provide hierarchical descriptions derived through automated multilevel image analysis. GeoPixelD structures its descriptions at three primary levels: (1) holistic scene representation, (2) individual instance observations, and (3) densely populated object groups annotations (as depicted in Figure 3).

4.1. Holistic Image Annotation

Initially, we generated descriptive captions for RS images using a robust open source model, IXC (Zhang et al., 2024a), to capture comprehensive and diverse image details. We chose the IXC model (Zhang et al., 2024a) based on a comparative study conducted with other state-of-the-art vision language mod-

els, where IXC consistently outperformed its counterparts in terms of qualitative performance. Open-ended descriptions are restricted to a limited length, integrated in prompts like "<image> Describe the image in four short sentences" (Figure 3 (bottom left)). Thus, redundancy is effectively minimized in subsequent annotations, and the model is driven to provide a holistic, context-rich depiction of each image.

4.2. Individual Instance Annotation

Next, we identify key objects for depiction and employ set-of-mark (SOM) prompting (Yang et al., 2023), which overlays distinct visual markers over specific image regions, to provide auxiliary information for visually grounded output. However, directly employing this method for aerial imagery, characterized by expansive views and diverse objects within a single frame, poses challenges such as hallucinated markers and incorrectly associated details (see Figure 7). To enhance object description accuracy in complex RS images, we spatially guide the model by introducing prior knowledge in the query in the form of category name and location, along with a marked number. Thus, effectively guiding the model to generate a precise and comprehensive description.

Specifically, we partition each image into a 3x3 grid and localize the objects by measuring their overlap with the quadrants to determine their positional reference. This quadrant-based localization, combined with categorical labels and marked numbers, serves as positional and categorical priors for LMM, enhancing focus on the target objects, a process that proves effective given the densely packed and spatially complex nature of RS imagery, where objects often vary in

scale, orientation, and proximity.

We evaluated various open-source and proprietary models for prior-informed SOM prompting in RS imagery (see Figure 8), comparing combined and individual querying approaches. ChatGPT (OpenAI, 2023) generated detailed descriptions with inferred information, while Gemini (Team et al., 2023) and InternVL (Chen et al., 2024b) exhibited repetitive output as the target objects increased. InternLM-XComposer (Zhang et al., 2024a) achieved performance comparable to ChatGPT in terms of accuracy and diversity.

4.3. Cluster/Crowd Annotation

After identifying and annotating prominent large objects, the remaining objects are grouped or identified along with their spatial properties through a structured three-stage positional analysis. First, the image is divided into a 3×3 grid, with each grid cell assigned a unique identifier. To enhance alignment with human perceptual tendencies, the central region of the grid is given a larger spatial weight. Next, a 2×2 grid is considered for the localization of more dispersed objects. Finally, a half-image grid (1×2 and 2×1) assigns broader positional information. This gridding provides a systematic framework for the localization of clusters as well as large groups of objects. LMM then describes the group attributes using quantitative and positional information.

4.4. Unifying Annotations and Language Marking

For the preprocessed iSAID (Waqas Zamir et al., 2019) train dataset (Appendix A), we derive 16,795 holistic, 36,793 instance, and 17,023 group annotations, collectively encompassing 600,817 objects. The annotations were rigorously filtered to eliminate aerial perspective inconsistencies, artifacts such as marker identifiers, fore/background references, depth cues, and inconsistent descriptors.

The key noun chunk corresponding to the object category in individual- and group-level annotations is tagged with unique identifiers ('phrase-number'), each linked to an instance or semantic mask, a process termed *text marking*. These marked annotations are combined with holistic scene representations into a single descriptive narrative. We employ a Llama-3.1-instruct 8B (Dubey et al., 2024) LLM to paraphrase concatenated annotations while preserving their semantic integrity (see Figure 9). The LLM processes the concatenated text under strict constraints to retain all marked phrases unchanged, ensuring a consistent link to associated visual masks. The outputs undergo iterative paraphrasing if any marked phrases are not preserved. By adopting this language marking strategy, the GeoPixelD dataset achieves a robust framework to generate high-quality, context-rich GCG descriptions precisely aligned with visual elements.

Following similar procedures, test set GCG descriptions

(utilizing iSAID validation set images) undergoes meticulous manual curation, requiring ~ 350 man-hours to ensure annotation completeness. We employed a robust validation protocol where each gcg description in the test set was meticulously verified by expert annotators. The process includes correcting for any omissions, inaccuracies, or partial annotations, including adjustments to object attributes that do not align with the corresponding image, thereby establishing a high-quality evaluation benchmark.

5. Experiments

Here, we explain the implementation details, present a comparative performance analysis on Remote Sensing Grounded Conversation Generation (RS-GCG) and Referring Remote Sensing Image Segmentation (RRSIS), and include an ablation study to assess the impact of key components.

5.1. Implementation Details

The model is initialized with pre-trained InternLM-XComposer-2.5 model (IXC-2.5) with 7B parameters, using LoRA for efficient fine-tuning. A fixed CLIP ViT-L vision encoder with a resolution of 560×560 is employed, along with a grounded vision encoder initialized from SAM2 weights. The trainable components include a pixel decoder (\mathcal{D}), LoRA parameters ($\alpha = 8$), a vision projector \mathcal{P}_v , and a language projector \mathcal{P}_t . The adaptive image divider uses a maximum patch number \mathcal{P} to 9 for training. In our training process, with an effective batch size of 20 over 10 epochs, the learning rate increases linearly to a maximum value of 3×10^{-4} over the initial 100 training steps, followed by a gradual cosine decay. We train GeoPixel on the GeoPixelD dataset for the GCG task on two NVIDIA A6000-48GB GPUs, which takes around 3 days.

5.2. Baselines

To rigorously evaluate the efficacy of GeoPixel, we introduce three robust baselines for comparative analysis on the GeoPixelD benchmark. First baseline LISA \dagger , an enhanced LISA (Lai et al., 2024) model, incorporates multitarget segmentation masks within its output pipeline and an updated tokenizer that includes phrase tokens ($\langle p \rangle$ and $\langle /p \rangle$) for the GCG task. The second baseline PixelLM \dagger , derived from PixelLM (Ren et al., 2024), is configured without the SAM encoder, the codebook image feature scaling factor of 2, three segmentation tokens, and the vision tower resize parameter of 448. Phrase tokens are added, and $\langle \text{SEG} \rangle$ token is replaced with multiple codebook tokens. The third baseline, GLaMM, specifically GLaMM-GCG, a model tailored for the GCG task. LISA \dagger , PixelLM \dagger and GLaMM-ft models are initialized from pretrained LISA-7B-v1, PixelLM-7B and GLaMM-GCG (7B), respectively, and additionally trained on GeoPixelD data.

Table 2. Performance Comparison on RS-GCG task. LISA[†] and PixelLM[†] denote the pretrained LISA and PixelLM models adopted for RS-GCG and finetuned on GeoPixelD training data. GLaMM represents the zero-shot performance, whereas GLaMM-FT refers to the pretrained model finetuned on GeoPixelD. GeoPixel outperforms other models across all metrics.

MODEL	CIDER	METEOR	CLAIR	UNI-TARGET			MULTI-TARGET			OVERALL		
				AP50	mIoU	RECALL	AP50	mIoU	RECALL	AP50	mIoU	RECALL
GLaMM (CVPR'24)	0.1	5.8	43.11	1.2	18.1	14.8	0.5	16.5	6.3	0.5	16.9	7.1
LISA [†] (CVPR'24)	14.6	22.3	68.96	9.5	41.7	43.1	8.3	43.1	27.5	8.5	42.7	29.0
PIXELLM [†] (CVPR'24)	18.3	22.5	73.93	13.5	41.2	44.0	10.4	42.9	28.1	10.5	42.4	29.6
GLaMM-FT (CVPR'24)	15.7	23.0	71.74	18.8	44.4	48.5	12.4	47.1	31.1	12.5	46.4	32.8
GEOPIXEL	21.6	24.0	77.50	25.5	50.8	55.6	18.0	52.9	37.0	19.0	52.3	38.8

Table 3. Performance Comparison of GeoPixel in Referring Expression Segmentation on RRSIS-D dataset. The segmentation accuracy based on referring expressions is expressed through the Precision at IoU threshold of 0.5 (P@0.5), Overall Intersection-over-Union (oIoU) and Mean Intersection-over-Union (mIoU).

METHOD	VALIDATION SET			TEST SET		
	P@0.5	oIoU	mIoU	P@0.5	oIoU	mIoU
RRN (LI ET AL., 2018)	51.09	66.53	46.06	51.07	66.43	45.64
CSMA (YE ET AL., 2019)	55.68	69.68	48.85	55.32	69.39	48.54
LSCM (HUI ET AL., 2020)	57.12	69.28	50.36	56.02	69.05	49.92
CMPC (HUANG ET AL., 2020)	57.93	70.15	50.41	55.83	69.22	49.24
BRINET (HU ET AL., 2020)	58.79	70.73	51.14	56.90	69.88	49.65
CMPC+ (LIU ET AL., 2022)	59.19	70.14	51.41	57.65	68.64	50.24
LGCE (YUAN ET AL., 2024)	68.10	76.68	60.16	67.65	76.34	59.37
LAVT (YANG ET AL., 2024)	69.54	77.59	61.46	69.52	77.19	61.04
RMSIN (LIU ET AL., 2024C)	74.66	78.27	65.10	74.26	77.79	64.20
GEOPIXEL-FT	80.00	81.77	67.99	83.33	84.90	67.30

5.3. Results

Remote Sensing Grounded Conversation Generation: A comparative analysis of the performance of various models on the RS-GCG task is evaluated across different metrics, including CIDEr, METEOR, CLAIR (Chan et al., 2023), AP50, mIoU, and recall, segmented into Uni-Target, Multi-Target, and Overall categories. CLAIR is an LLM-based similarity metric (GPT-4o in our case) that better aligns with human judgments. GeoPixel demonstrates superior performance in all metrics compared to the baselines, showing better fluency and text relevance in textual outputs while maintaining strong performance in more complex multi-target scenarios. In contrast, LISA[†] struggles with segmentation-based tasks, as evidenced by its low AP50 scores in all categories. PixelLM[†] shows a moderate improvement over LISA[†], benefiting from better image feature scaling and segmentation token adjustments. GLaMM-ft exhibits improved outcomes due to a dedicated grounding encoder and GCG pre-training, however, its performance remains inferior to that of GeoPixel (as detailed in Table 2). Qualitative results are presented in Figure 4.

Referring Remote Sensing Image Segmentation: This task focuses on segmenting specific regions in aerial

Table 4. Performance Comparison on Referring Expression Detection Task, reporting Acc@0.5 / Acc@0.7 across 3 categories: unique, non-unique and overall. GeoPixel demonstrated substantial gains over existing methods in referring expression detection/localization.

MODEL	UNIQUE		NON-UNIQUE		OVERALL	
	Acc@0.5	Acc@0.7	Acc@0.5	Acc@0.7	Acc@0.5	Acc@0.7
MINIGPT-V2	40.7	18.9	32.4	15.2	35.8	16.8
LLAVA-1.5	51.1	16.4	34.8	11.5	41.6	13.6
MINI-GEMINI	41.1	9.6	22.3	4.9	30.1	6.8
GEOCHAT	57.4	22.6	44.5	18.0	49.8	19.9
GEOPIX	57.0	22.7	44.8	18.2	49.8	20.0
GEOPIXEL-FT	70.37	41.54	65.80	40.32	67.70	40.83

imagery guided by textual descriptions. The input prompt used is: "Could you provide a segmentation mask for {referring_expression} in this image?" The model generates the response, "Sure, it is <SEG>." where the corresponding embeddings of <SEG> token is subsequently decoded to produce the segmentation mask. To address this task, we fine-tune the GeoPixel model on the RRSIS-D (Liu et al., 2024c) dataset. The resulting GeoPixel-ft model demonstrates superior performance compared to recent approaches, as shown by results on the RRSIS-D test and validation sets in Table 3. The qualitative results are provided in Figure 5.

Referring Remote Sensing Image Detection: We formulate referring expression detection (RED) as a post-processing task on GeoPixel's predictions, deriving horizontal (HBBs) and oriented bounding boxes (OBBs) from output masks, for consistent evaluation. Table 4 compares performance on VRSBench (Li et al., 2024a), a dataset based on DOTAv2 and DIOR. To prevent data leakage, GeoPixel (trained on GeoPixelD) is finetuned on VRSBench without using any data from RRSIS-D (DIOR-based dataset). Moreover, GeoPixelD and VRSBench use DOTA's training set for training and validation set for testing. We further compare GeoChat and GeoPixel using OBBs (Table 5) with both models finetuned on VRSBench, highlighting the importance of pixel alignment in LMMs. Figure 10 illustrates the qualitative results.

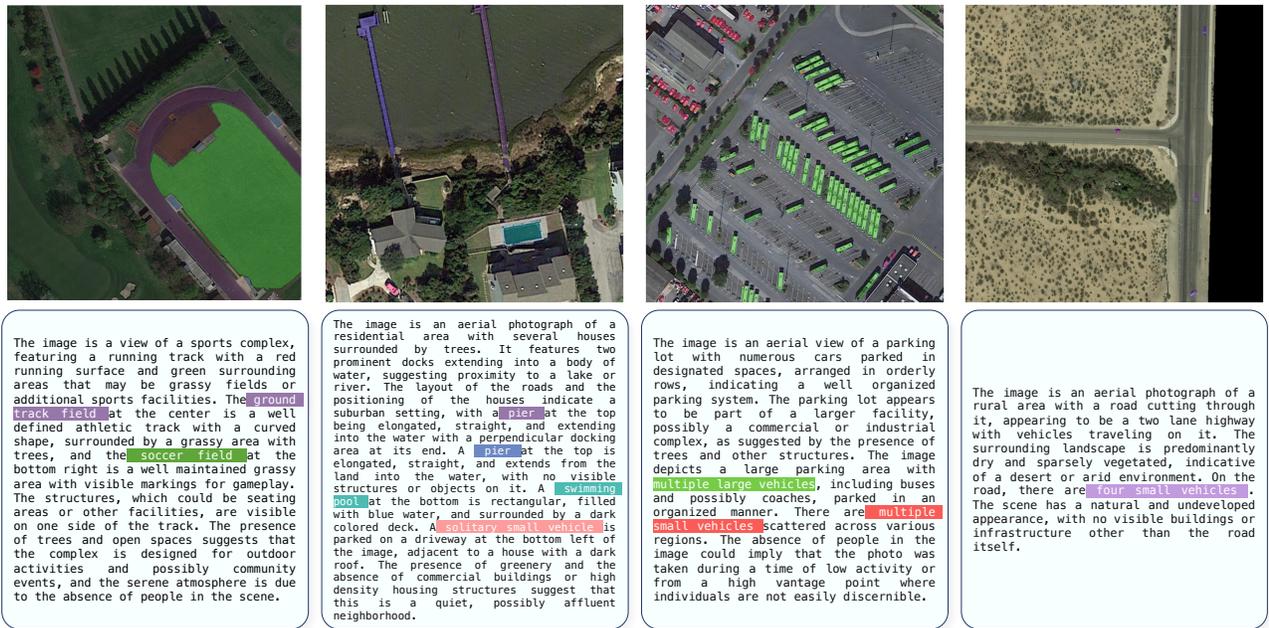


Figure 4. Qualitative results of GeoPixel on RS-GCG. Contextually rich descriptions of RS imagery with grounded object annotations. Depending on object scale and density, it employs instance masks for precise delineation of individual objects (right and middle-right images) while semantic masks capture broader categories, such as large clusters of vehicles or small objects (middle-left and left images).

Table 5. Visual grounding performance on VRSBench dataset using oriented bounding boxes for referring object localization. The observed performance improvement underscores the critical role of precise pixel-level alignment.

MODEL	UNIQUE		NON-UNIQUE		OVERALL	
	Acc@0.5	Acc@0.7	Acc@0.5	Acc@0.7	Acc@0.5	Acc@0.7
GEOCHAT	32.3	12.6	18.5	5.7	24.3	8.6
GEOPIXEL-FT	54.48	24.87	60.51	30.97	58.00	28.42

5.4. Ablation Study

Inference Resolution Effect: Increasing inference patches improves all evaluation metrics, reflecting better comprehension of visual content (Table 6). At $\mathcal{P} = 9$, CIDEr increases from 14.6 to 20.5, and METEOR improves from 23.1 to 24.3, indicating enhanced semantic understanding as image tokens scale up. The moderate mAP and mIoU gains suggest that while HR inference contributes to superior localization accuracy, competitive performance can still be maintained at lower resolutions when the model is pre-trained at HR. The superior results at ($\mathcal{P} = 9$) underscore the critical role of incorporating fine-grained spatial details during the training phase for generalized feature learning.

Annotation Complexity Effect: GeoPixel adjusts its masking based on object size and distribution (Figure 4), using instance masks for precise identification, while semantic masks for broader categories (clusters/small objects). In scenarios requiring both granularity and generalization, the model employs hybrid annotations, blending instance-level

Table 6. Effect of Inference Resolution. Reported metrics show the relationship between resolution and overall performance.

TRAINING PATCHES	INFERENCE PATCHES	CIDEr	METEOR	AP50	mIoU	RECALL
	$\mathcal{P} = 1$	14.6	23.1	12.9	47.8	32.2
	$\mathcal{P} = 4$	17.7	23.9	16.6	51.8	37.1
	$\mathcal{P} = 9$	20.5	24.3	17.6	52.1	37.4

Table 7. Effect of Annotation Complexity. Avg. Len is the average character length of captions.

DATA	OBJECTS	PHRASES	AVG. LEN	mIoU	RECALL
INSTANCES ONLY	1,740	1,740	634	58.4	48.8
SEMANTIC ONLY	21,483	698	518	44.1	37.7
MIX DATA	38,161	2,989	737	50.9	33.3

and semantic mask representations(as seen in Figure 1). Table 7 shows that this annotation complexity impacts performance, with mixed annotations yielding the lowest recall.

Remote sensing images often contain visually similar objects with subtle variations in appearance, spatial layout, and proximity, yet exhibit significant scale variations across images. This complexity challenges the model’s ability to distinguish between object presence, quantity, and the appropriate annotation type required (e.g., instance or semantic). The challenge is pronounced in the semantic-only category, where the model exhibits the lowest mIoU scores, indicating two key challenges: the model’s ability to ensure full coverage of all instances within a category and grouping



Figure 5. Qualitative results of GeoPixel’s capability in referring remote sensing expression segmentation. The figure highlights Geopixel’s ability to interpret referring expressions of varying lengths and generate precise segmentation masks, adapting to scale variations, as shown in the ground track fields. Spatial descriptors (e.g "right", "lower right"), and object characteristics (e.g "red") are interpreted with precision to achieve accurate segmentation.

Table 8. Effect of Data Complexity and Training Vision Projection (VP) Layer. T stands for Trainable and F for Frozen.

TRAINING DATA		VP	CIDER	METEOR	AP50	MIoU	RECALL
SET-1A	SET-1B						
✓		T	19.3	23.6	18.2	48.0	33.6
✓	✓	T	20.5	24.0	17.8	51.7	36.7
✓	✓	F	18.7	24.4	15.3	51.6	35.1

objects under a unified mask instead of individual instances. The low mask recall in mixed data further suggests that generalizing masking decisions in dense scenes is difficult due to scale and distributional variability of geospatial entities.

Role of Data Complexity: Table 8 compares GeoPixel’s performance across data partitions with varying masking complexity. Set-1A is less complex, with no intraclass segmentation differences. Each instance of a single class is individually masked or represented using a semantic mask. Set-1B introduces a higher complexity by assigning instance masks to larger objects while grouping smaller ones under a semantic mask within the same class. For example, larger boats may be individually described, while smaller boats in the image could be grouped under a single semantic description. This ablation evaluates the model’s ability to handle varying levels of annotation granularity, providing

insights into its ability to generalize across different scales and segmentation strategies. The results indicate that the inclusion of more complex annotation (Set-1B) enhances segmentation accuracy and descriptive detail, as the model is trained with more diverse mask configurations.

Vision Projection: Next we study the effect of training the vision projection layer by comparing the performance when the vision projection layer is fixed or trainable during the fine-tuning stage. Table 8 summarizes the results. Training the vision projection layer results in an improvement in some metrics, highlighting the role of feature alignment.

6. Conclusion

We present GeoPixel, a large multimodal model (LMM) designed specifically for the unique challenges of high-resolution remote sensing (RS) image analysis. GeoPixel introduces a robust end-to-end architecture capable of adaptive image partitioning and pixel-level grounding, enabling the precise interpretation and generation of geospatially aware descriptions in RS imagery. By addressing key limitations of current LMMs, such as low-resolution constraints and coarse object-grounding, GeoPixel provides a fine-grained visual understanding that bridges the gap between language and high-resolution RS data.

Impact Statement

GeoPixel, a large multimodal model, is designed to enhance the fine-grained spatial understanding of high-resolution remote sensing (RS) imagery. Our work contributes to the advancement of machine learning by improving pixel-level object grounding, a critical capability for geospatial applications such as urban planning, environmental monitoring, disaster response, and infrastructure assessment.

From an ethical perspective, GeoPixel operates on publicly available remote sensing datasets, ensuring transparency and reproducibility in research. However, the use of high-resolution imagery raises privacy considerations, particularly in applications that involve urban environments. Although our model does not process personally identifiable information (PII) or real-time surveillance data, we acknowledge that future extensions incorporating private datasets should adhere to strict data governance and privacy-preserving policies.

Regarding social impact, the ability of GeoPixel to generate detailed geospatially grounded descriptions can help in disaster relief efforts by improving situational awareness from satellite imagery. It can also improve climate change monitoring by facilitating precise land use classification and deforestation tracking. We emphasize that GeoPixel is developed for scientific and humanitarian purposes and encourage ethical deployment in accordance with open research principles. Overall, this work aligns with the broader goal of advancing machine learning in remote sensing, promoting open and responsible AI applications, and fostering positive societal impact while remaining mindful of potential ethical considerations.

Acknowledgement

The computations were enabled by resources provided by NAISS at Alvis, partially funded by the Swedish Research Council through grant agreement no. 2022-06725, LUMI hosted by CSC (Finland) and LUMI consortium, and by Berzelius resource provided by the Knut and Alice Wallenberg Foundation at the NSC

References

- Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., Hui, B., Ji, L., Li, M., Lin, J., Lin, R., Liu, D., Liu, G., Lu, C., Lu, K., Ma, J., Men, R., Ren, X., Ren, X., Tan, C., Tan, S., Tu, J., Wang, P., Wang, S., Wang, W., Wu, S., Xu, B., Xu, J., Yang, A., Yang, H., Yang, J., Yang, S., Yao, Y., Yu, B., Yuan, H., Yuan, Z., Zhang, J., Zhang, X., Zhang, Y., Zhang, Z., Zhou, C., Zhou, J., Zhou, X., and Zhu, T. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023a.
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., and Zhou, J. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2): 3, 2023b.
- Bazi, Y., Bashmal, L., Al Rahhal, M. M., Ricci, R., and Melgani, F. Rs-llava: A large vision-language model for joint captioning and question answering in remote sensing imagery. *Remote Sensing*, 16(9):1477, 2024.
- Cai, Z., Cao, M., Chen, H., Chen, K., Chen, K., Chen, X., Chen, X., Chen, Z., Chen, Z., Chu, P., Dong, X., Duan, H., et al. Internlm2 technical report, 2024.
- Chan, D., Petryk, S., Gonzalez, J. E., Darrell, T., and Canny, J. Clair: Evaluating image captions with large language models. *arXiv preprint arXiv:2310.12971*, 2023.
- Chen, G., Shen, L., Shao, R., Deng, X., and Nie, L. Lion: Empowering multimodal large language model with dual-level visual knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26540–26550, June 2024a.
- Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., and Zhao, R. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., Li, B., Luo, P., Lu, T., Qiao, Y., and Dai, J. 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 (CVPR)*, pp. 24185–24198, June 2024b.
- Chiang, W.-L., Li, Z., Lin, Z., Sheng, Y., Wu, Z., Zhang, H., Zheng, L., Zhuang, S., Zhuang, Y., Gonzalez, J. E., Stoica, I., and Xing, E. P. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna/>.
- Dai, W., Li, J., Li, D., Tiong, A., Zhao, J., Wang, W., Li, B., Fung, P. N., and Hoi, S. Instructblip: Towards general-purpose vision-language models with instruction tuning. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 49250–49267. Curran Associates, Inc., 2023.
- Ding, G., Khan, S., Tang, Z., Zhang, J., and Porikli, F. Towards better validity: Dispersion based clustering for unsupervised person re-identification. *arXiv preprint arXiv:1906.01308*, 2019.

- Dong, X., Zhang, P., Zang, Y., Cao, Y., Wang, B., Ouyang, L., Wei, X., Zhang, S., Duan, H., Cao, M., Zhang, W., Li, Y., Yan, H., Gao, Y., Zhang, X., Li, W., Li, J., Chen, K., He, C., Zhang, X., Qiao, Y., Lin, D., and Wang, J. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *arXiv preprint arXiv:2401.16420*, 2024a.
- Dong, X., Zhang, P., Zang, Y., Cao, Y., Wang, B., Ouyang, L., Zhang, S., Duan, H., Zhang, W., Li, Y., Yan, H., Gao, Y., Chen, Z., Zhang, X., Li, W., Li, J., Wang, W., Chen, K., He, C., Zhang, X., Dai, J., Qiao, Y., Lin, D., and Wang, J. Internlm-xcomposer2-4khd: A pioneering large vision-language model handling resolutions from 336 pixels to 4k hd. *arXiv preprint arXiv:2404.06512*, 2024b.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., and Houshy, N. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Guo, Q., De Mello, S., Yin, H., Byeon, W., Cheung, K. C., Yu, Y., Luo, P., and Liu, S. Regionpt: Towards region understanding vision language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13796–13806, June 2024.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., and Girshick, R. Masked autoencoders are scalable vision learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15979–15988, 2022. doi: 10.1109/CVPR52688.2022.01553.
- Hong, W., Wang, W., Lv, Q., Xu, J., Yu, W., Ji, J., Wang, Y., Wang, Z., Dong, Y., Ding, M., et al. Cogagent: A visual language model for gui agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14281–14290, 2024.
- Hu, Y., Yuan, J., Wen, C., Lu, X., and Li, X. Rsgpt: A remote sensing vision language model and benchmark. *arXiv preprint arXiv:2307.15266*, 2023.
- Hu, Z., Feng, G., Sun, J., Zhang, L., and Lu, H. Bi-directional relationship inferring network for referring image segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4423–4432, 2020. doi: 10.1109/CVPR42600.2020.00448.
- Huang, S., Hui, T., Liu, S., Li, G., Wei, Y., Han, J., Liu, L., and Li, B. Referring image segmentation via cross-modal progressive comprehension. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10488–10497, 2020.
- Hui, T., Liu, S., Huang, S., Li, G., Yu, S., Zhang, F., and Han, J. 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*, pp. 59–75. Springer, 2020.
- Irvin, J. A., Liu, E. R., Chen, J. C., Dormoy, I., Kim, J., Khanna, S., Zheng, Z., and Ermon, S. TeoChat: A large vision-language assistant for temporal earth observation data. In *International Conference on Learning Representations*, 2025.
- Javaheripi, M., Bubeck, S., et al. Phi-2: the surprising power of small language models (2023). URL <https://www.microsoft.com/en-us/research/blog/phi-2-the-surprising-power-of-small-language-models>.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4015–4026, 2023.
- Kuckreja, K., Danish, M. S., Naseer, M., Das, A., Khan, S., and Khan, F. S. Geochat: grounded large vision-language model for remote sensing. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 27831–27840, 2024. doi: 10.1109/CVPR52733.2024.02629.
- Lai, X., Tian, Z., Chen, Y., Li, Y., Yuan, Y., Liu, S., and Jia, J. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9579–9589, 2024.
- Li, J., Li, D., Savarese, S., and Hoi, S. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.
- Li, R., Li, K., Kuo, Y.-C., Shu, M., Qi, X., Shen, X., and Jia, J. Referring image segmentation via recurrent refinement networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5745–5753, 2018. doi: 10.1109/CVPR.2018.00602.
- Li, X., Ding, J., and Elhoseiny, M. Vrsbench: A versatile vision-language benchmark dataset for remote sensing image understanding. In Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., and Zhang, C.

- (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 3229–3242. Curran Associates, Inc., 2024a. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/05b7f821234f66b78f99e7803fffa78a-Paper-Datasets_and_Benchmarks_Track.pdf.
- Li, Z., Yang, B., Liu, Q., Ma, Z., Zhang, S., Yang, J., Sun, Y., Liu, Y., and Bai, X. Monkey: Image resolution and text label are important things for large multi-modal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26763–26773, 2024b.
- Lin, Z., Liu, C., Zhang, R., Gao, P., Qiu, L., Xiao, H., Qiu, H., Lin, C., Shao, W., Chen, K., et al. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv preprint arXiv:2311.07575*, 2023.
- Liu, H., Li, C., Li, Y., Li, B., Zhang, Y., Shen, S., and Lee, Y. J. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024a. URL <https://llava-vl.github.io/blog/2024-01-30-llava-next/>.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024b.
- Liu, S., Hui, T., Huang, S., Wei, Y., Li, B., and Li, G. Cross-modal progressive comprehension for referring segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4761–4775, 2022. doi: 10.1109/TPAMI.2021.3079993.
- Liu, S., Cheng, H., Liu, H., Zhang, H., Li, F., Ren, T., Zou, X., Yang, J., Su, H., Zhu, J., et al. Llava-plus: Learning to use tools for creating multimodal agents. *arXiv preprint arXiv:2311.05437*, 2023a.
- Liu, S., Ma, Y., Zhang, X., Wang, H., Ji, J., Sun, X., and Ji, R. Rotated multi-scale interaction network for referring remote sensing image segmentation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26648–26658, 2024c. doi: 10.1109/CVPR52733.2024.02517.
- Liu, X. and Lian, Z. Rsunivlm: A unified vision language model for remote sensing via granularity-oriented mixture of experts. *arXiv preprint arXiv:2412.05679*, 2024.
- Liu, Y., Yang, B., Liu, Q., Li, Z., Ma, Z., Zhang, S., and Bai, X. Textmonkey: An ocr-free large multi-modal model for understanding document. *arXiv preprint arXiv:2403.04473*, 2024d.
- Liu, Z., He, Y., Wang, W., Wang, W., Wang, Y., Chen, S., Zhang, Q., Lai, Z., Yang, Y., Li, Q., et al. Interngpt: Solving vision-centric tasks by interacting with chatgpt beyond language. *arXiv preprint arXiv:2305.05662*, 2023b.
- Luo, J., Pang, Z., Zhang, Y., Wang, T., Wang, L., Dang, B., Lao, J., Wang, J., Chen, J., Tan, Y., et al. Skysensegpt: A fine-grained instruction tuning dataset and model for remote sensing vision-language understanding. *arXiv preprint arXiv:2406.10100*, 2024.
- Ma, C., Jiang, Y., Wu, J., Yuan, Z., and Qi, X. Groma: Localized visual tokenization for grounding multimodal large language models. In *European Conference on Computer Vision*, pp. 417–435. Springer, 2025.
- Muhtar, D., Li, Z., Gu, F., Zhang, X., and Xiao, P. Lhrsbot: Empowering remote sensing with vgi-enhanced large multimodal language model. *arXiv preprint arXiv:2402.02544*, 2024.
- OpenAI. Chatgpt: Language model for dialogue applications. <https://openai.com/chatgpt>, 2023. Accessed: 2024-12-31.
- Ou, R., Hu, Y., Zhang, F., Chen, J., and Liu, Y. Geopix: Multi-modal large language model for pixel-level image understanding in remote sensing. *arXiv preprint arXiv:2501.06828*, 2025.
- Pang, C., Wu, J., Li, J., Liu, Y., Sun, J., Li, W., Weng, X., Wang, S., Feng, L., Xia, G.-S., et al. H2rsvlm: Towards helpful and honest remote sensing large vision language model. *arXiv preprint arXiv:2403.20213*, 2024.
- Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., Ye, Q., and Wei, F. Grounding multimodal large language models to the world. In *The Twelfth International Conference on Learning Representations*, 2024.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Rasheed, H., Maaz, M., Shaji, S., Shaker, A., Khan, S., Cholakkal, H., Anwer, R. M., Xing, E., Yang, M.-H., and Khan, F. S. Glamm: Pixel grounding large multimodal model. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13009–13018, 2024. doi: 10.1109/CVPR52733.2024.01236.
- Ravi, N., Gabeur, V., Hu, Y.-T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., Mintun, E., Pan, J., Alwala, K. V., Carion, N., Wu, C.-Y., Girshick, R., Dollár, P., and Feichtenhofer, C. Sam 2: Segment anything in images and videos. *arXiv*

- preprint *arXiv:2408.00714*, 2024. URL <https://arxiv.org/abs/2408.00714>.
- Ren, Z., Huang, Z., Wei, Y., Zhao, Y., Fu, D., Feng, J., and Jin, X. Pixellm: Pixel reasoning with large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26374–26383, June 2024.
- Ryali, C., Hu, Y.-T., Bolya, D., Wei, C., Fan, H., Huang, P.-Y., Aggarwal, V., Chowdhury, A., Poursaeed, O., Hoffman, J., Malik, J., Li, Y., and Feichtenhofer, C. Hiera: A hierarchical vision transformer without the bells-and-whistles. *ICML*, 2023.
- Soni, S., Dudhane, A., Debary, H., Fiaz, M., Munir, M. A., Danish, M. S., Fraccaro, P., Watson, C. D., Klein, L. J., Khan, F. S., et al. Earthdial: Turning multi-sensory earth observations to interactive dialogues. *arXiv preprint arXiv:2412.15190*, 2024.
- Team, G., Anil, R., Borgeaud, S., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., Millican, K., et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Waqas Zamir, S., Arora, A., Gupta, A., Khan, S., Sun, G., Shahbaz Khan, F., Zhu, F., Shao, L., Xia, G.-S., and Bai, X. isaid: A large-scale dataset for instance segmentation in aerial images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 28–37, 2019.
- Wei, H., Kong, L., Chen, J., Zhao, L., Ge, Z., Yang, J., Sun, J., Han, C., and Zhang, X. Vary: Scaling up the vision vocabulary for large vision-language model. In *European Conference on Computer Vision*, pp. 408–424. Springer, 2025.
- Xia, Z., Han, D., Han, Y., Pan, X., Song, S., and Huang, G. Gsva: Generalized segmentation via multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3858–3869, June 2024.
- Xuan, S., Guo, Q., Yang, M., and Zhang, S. Pink: Unveiling the power of referential comprehension for multi-modal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13838–13848, June 2024.
- Yang, J., Zhang, H., Li, F., Zou, X., Li, C., and Gao, J. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*, 2023.
- Yang, Z., Wang, J., Ye, X., Tang, Y., Chen, K., Zhao, H., and Torr, P. H. Language-aware vision transformer for referring segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–18, 2024. doi: 10.1109/TPAMI.2024.3468640.
- Ye, J., Hu, A., Xu, H., Ye, Q., Yan, M., Xu, G., Li, C., Tian, J., Qian, Q., Zhang, J., et al. Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model. *arXiv preprint arXiv:2310.05126*, 2023a.
- Ye, L., Rochan, M., Liu, Z., and Wang, Y. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10502–10511, 2019.
- Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023b.
- You, H., Zhang, H., Gan, Z., Du, X., Zhang, B., Wang, Z., Cao, L., Chang, S.-F., and Yang, Y. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023.
- Yuan, Z., Mou, L., Hua, Y., and Zhu, X. X. Rrsis: Referring remote sensing image segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- Zhan, Y., Xiong, Z., and Yuan, Y. Skyeyegpt: Unifying remote sensing vision-language tasks via instruction tuning with large language model. *arXiv preprint arXiv:2401.09712*, 2024.
- Zhang, P., Dong, X., Zang, Y., Cao, Y., Qian, R., Chen, L., Guo, Q., Duan, H., Wang, B., Ouyang, L., Zhang, S., Zhang, W., Li, Y., Gao, Y., Sun, P., Zhang, X., Li, W., Li, J., Wang, W., Yan, H., He, C., Zhang, X., Chen, K., Dai, J., Qiao, Y., Lin, D., and Wang, J. Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output. *arXiv preprint arXiv:2407.03320*, 2024a.
- Zhang, S., Sun, P., Chen, S., Xiao, M., Shao, W., Zhang, W., Liu, Y., Chen, K., and Luo, P. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv preprint arXiv:2307.03601*, 2023.
- Zhang, W., Cai, M., Zhang, T., Zhuang, Y., and Mao, X. Earthgpt: A universal multimodal large language model

for multisensor image comprehension in remote sensing domain. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–20, 2024b. doi: 10.1109/TGRS.2024.3409624.

Zhang, W., Cai, M., Zhang, T., Zhuang, Y., and Mao, X. Earthgpt: A universal multi-modal large language model for multi-sensor image comprehension in remote sensing domain. *IEEE Transactions on Geoscience and Remote Sensing*, 2024c.

Zhang, Y., Ma, Z., Gao, X., Shakiah, S., Gao, Q., and Chai, J. Groundhog: Grounding large language models to holistic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14227–14238, June 2024d.

Zhao, Y., Lin, Z., Zhou, D., Huang, Z., Feng, J., and Kang, B. Bubogpt: Enabling visual grounding in multi-modal llms. *arXiv preprint arXiv:2307.08581*, 2023.

Zhou, Y., Lan, M., Li, X., Ke, Y., Jiang, X., Feng, L., and Zhang, W. Geoground: A unified large vision-language model. for remote sensing visual grounding. *arXiv preprint arXiv:2411.11904*, 2024.

Zhu, D., Chen, J., Shen, X., Li, X., and Elhoseiny, M. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

A. Limitations and Challenges

While GeoPixel has demonstrated significant advances in pixel-level grounding for high-resolution RS images, several challenges remain. These challenges are particularly evident in the following failure cases (illustrated in Figure 6). The model occasionally produces erroneous masks due to ambiguities in the masking strategy, particularly in determining object presence and quantity, as well as deciding whether semantic segmentation or instance-level annotation is appropriate. An incorrect decision in this regard can result in repetitive descriptions of visually similar objects, leading to inconsistencies in the generated output. Furthermore, such errors may manifest as fragmented or overlapping masks, introducing confusion in object delineation and undermining the overall segmentation quality. Moreover, the model often confuses instance masks within the same spatial location, particularly in densely populated or crowded images.

Future work may focus on addressing these challenges by incorporating more robust masking strategies and dynamic resolution adjustment techniques to improve segmentation accuracy in complex scenes. Additionally, extending GeoPixel’s capabilities to integrate multimodal data, such as Synthetic Aperture Radar (SAR) or infrared imagery, could significantly enhance its ability to analyze diverse remote sensing datasets. GeoPixel is a significant step forward in leveraging the potential of LMMs for remote sensing, opening new avenues for research in this critical domain.

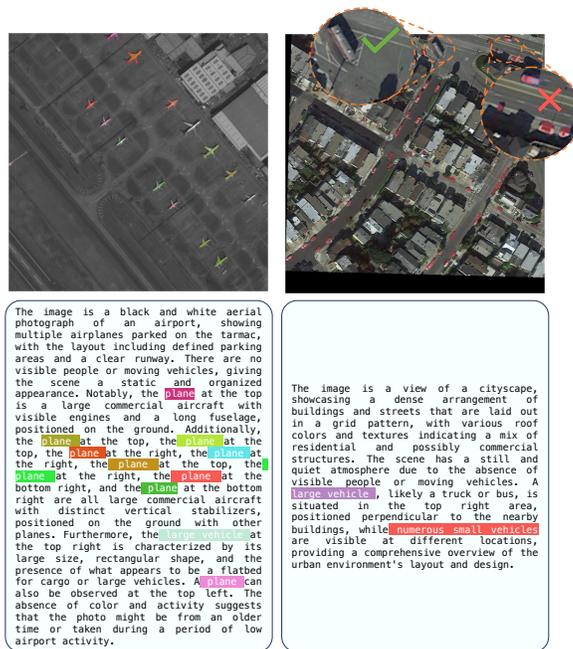


Figure 6. Failure case due to incorrect mask association (left) and wrong instance segmentation in the same spatial region (right).

B. GeoPixelD dataset

Preprocessing and Marking: We utilize the instance-level annotated dataset, iSAID (Waqas Zamir et al., 2019), to generate grounded conversations through our annotation pipelines. The images undergo a preprocessing step in which they are cropped into 800 x 800 pixel patches. Objects for instance annotations are selected based on an area threshold to ensure their reasonable size, therefore preventing the marker from obscuring a significant portion of the object and maintaining its distinguishability. A 14 x 14 pixels fixed size marker is used, regardless of the actual dimensions of the object. However, the marker’s placement is determined based on the segmentation mask’s area and shape. For large objects, the marker is positioned at the center of the mask if the calculated center falls within the mask boundaries; otherwise, it is adjusted to the nearest point on the object’s border. For small objects, the center of the bounding box is aligned with a point on the polygon mask boundary, which typically results in an average marker overlap of 50% with the object.

In addition, multiple marking techniques were explored, including bounding boxes, masks, contours, and numerical markers, to determine their impact on accuracy and object fidelity. Our findings reveal that bounding boxes and contours add superfluous visual information that can obscure the fine details of the object. In contrast, simple numerical markers placed

directly on the object effectively signal its presence without compromising visual clarity or introducing noise, thereby preserving the integrity of object details for enhanced model performance.

Set-of-Mark Prompting without Prior

Describe the object/objects of given image marked by number {markers_str} as if analyzing a map. The description should be enumerated in the following format: 1. <description> 2. <description> _

1. The roundabout at the center of the intersection is marked with a red star-shaped sculpture, which serves as a focal point for the traffic flow.
2. To the north of the roundabout, there is a multi-story building with a flat roof, which appears to be an institutional or commercial structure.
3. To the east of the roundabout, there is another multi-story building with a similar architectural style to the one in the north.
4. To the south of the roundabout, there is a smaller building with a flat roof, which could be an office or a service building.
5. To the west of the roundabout, there is a parking lot with several cars, indicating the presence of a commercial or public facility nearby.
6. There is a grassy area to the southeast of the roundabout, which provides a green space amidst the urban environment.



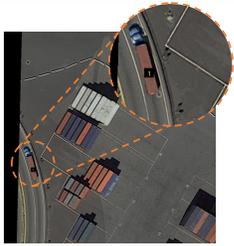
Set-of-Mark Prompting with Prior

List down all the attributes of the {category_name} at the {pos} that is marked with number {mark_number} in one short sentence. Do not mention the marked number in output.

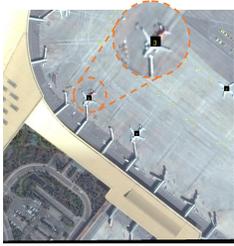
1. The roundabout at the center is marked by a distinctive red and orange star pattern.
2. The vehicle at the top right is a dark-colored sedan.
3. The vehicle at the top is a small, dark-colored car.
4. The vehicle at the top right is a dark-colored sedan.
5. The vehicle at the top right is a small, dark-colored car parked perpendicular to the curb, adjacent to building number 4.
6. The vehicle at the top is a blue car parked in a designated parking spot.

Figure 7. Comparative effectiveness of SOM prompting methods, highlighting the critical role of priors. Without priors, SOM relies solely on the VLM to detect and describe marked objects independently, resulting in inaccurate descriptions and hallucinated markers in complex remote sensing scenes. In contrast, SOM with priors utilizes explicit marker positions ({pos}) and predefined object categories ({category_name}) as priors, providing structured prompts that reduce ambiguity and guide the VLM to produce precise and reliable descriptions. Incorrect parts are noted in red whereas correct parts are noted in green.

List down all the attributes of the vehicle at the left that is marked with number 1 in one short sentence. Do not mention the marked number in the output.



- ✳️ The vehicle on the left is blue, appears to be a compact car, and is traveling on a curved road adjacent to a parking lot or storage area with shipping containers.
- ✦ The vehicle is a white truck driving on a curved road.
- 🚚 The vehicle at the left is a semi-truck with a blue cab and an attached orange container.
- 🚚 The vehicle is a blue truck.



- ✳️ 1.The plane in the center is parked at a gate, has a white fuselage with wings extended, appears to be a commercial airliner, and casts a shadow on the tarmac due to sunlight from the top left.
- 2.The plane on the right is parked at a gate, has a white fuselage with extended wings, appears to be a commercial airliner, and casts a shadow to the lower right side due to the sunlight from the top left.
- 3.The plane in the center is parked at a gate, has a white fuselage with extended wings, appears to be a commercial airliner, and casts a shadow toward the lower left side due to sunlight coming from the upper left.

- ✦ 1.The plane at the center is white and parked on a tarmac.
- 2.The plane at the right is white and parked on a tarmac.
- 3.The plane at the center is white and parked on a tarmac.

- ✳️ 1.The plane at the center is white with a red stripe and is parked at a gate.
- 2.The plane on the right is white with a blue stripe and is parked at a gate.
- 3.The plane at the center is white with a red stripe and is parked at a gate.

- 🚗 1.The plane at the center is a large commercial aircraft with a predominantly white body and blue engines.
- 2.The plane at the right is a large commercial aircraft with a predominantly white body and blue engines.
- 3.The plane at the center is a large commercial aircraft with two engines positioned on the tarmac facing a terminal gate.

Figure 8. Comparison of open-source and proprietary models for prior-informed set of marks (SOM) prompting for RS imagery. Incorrect parts are noted in red whereas correct parts are noted in green.

Prompt Template: Paraphrase the following description of an image viewed from a satellite in a single paragraph while ensuring that all the words enclosed in single quotes are included and preserved exactly as they are. Remove any mention of background, foreground, or camera position. Image Description: {Holistic Image Annotation}{0:last} {Individual Instance Annotation} {Cluster/Crowded Annotation} {Holistic Image Annotation}{last}

Prompt: Paraphrase the following description of an image viewed from a satellite in a single paragraph while ensuring that all the words enclosed in single quotes are included and preserved exactly as they are. Remove any mention of background, foreground, or camera position. Image Description: The image is an aerial photograph of a residential area with docks and boats. The water appears calm, and the docks are equipped with vehicles, suggesting recreational or transportation use. The surrounding land features greenery and private property, indicating a peaceful, possibly affluent neighborhood. The **object-1** at the center has a single boat docked to it, with a structure that includes a roof and appears to be made of wood. The **object-2** at the right is equipped with a boat and has a structure extending over the water, connected to a docking area on land. The **object-3** on the left is elongated with a small structure at its end, extending into the water and connected to the shore by a wooden plank walkway. The **object-4** is rectangular, has a clear blue surface, and is surrounded by a deck area with lounge chairs. The **object-5** at the bottom left is rectangular, filled with blue water, and surrounded by a deck area with lounge chairs. The **object-6** at the right is white, appears to be a motor boat, and is docked at a wooden pier. The **object-7** at the right is white with a blue stripe, appears to be a motor boat, and is docked at a wooden pier. The **object-8** at the center is a small motor boat with a single outboard motor, docked perpendicular to the pier. The **object-9** at the right is white with a single cabin and is docked at wooden piers. There is no visible activity, giving the scene a tranquil atmosphere.



GCG Description: The aerial photograph of a residential area features docks and boats, with calm water suggesting recreational or transportation use. The land surrounding the docks has greenery and private property, indicating a peaceful, possibly affluent neighborhood. The **object-1** at the center has a single boat docked to it, with a structure including a roof made of wood. Nearby, the **object-2** at the right is equipped with a boat and has a structure extending over the water, connected to a docking area on land. The **object-3** on the left is elongated with a small structure at its end, extending into the water and connected to the shore by a wooden plank walkway. Two swimming pools, **object-4** and **object-5**, are rectangular and surrounded by deck areas with lounge chairs. Various boats are visible at right, including a white motor **object-6**, **object-7**, a white motor boat, and **object-8** a white boat with a single cabin, all docked at wooden piers. A small motor **object-9** with a single outboard motor is also docked, perpendicular to the pier. The scene is tranquil, with no visible activity.

Prompt: Paraphrase the following description of an image viewed from a satellite in a single paragraph while ensuring that all the words enclosed in single quotes are included and preserved exactly as they are. Remove any mention of background, foreground, or camera position. Image Description: The image is an aerial view of a tennis court complex with multiple courts. The courts are arranged in a grid pattern, each with its own net and markings. The surrounding area includes greenery, indicating the presence of trees or landscaping. An aerial view of **object-1** with blue playing surfaces, marked for play, surrounded by greenery and adjacent to a parking area. There are no people visible on the courts, suggesting that play may not be in session at the time the photo was taken.



GCG Description: The image is a view of **object-1** with blue playing surfaces, marked for play, surrounded by greenery and adjacent to a parking area. It features multiple tennis courts arranged in a grid pattern, each with its own net and markings, situated amidst greenery, indicating the presence of trees or landscaping.

Figure 9. Unifying Annotations through LLM Paraphrasing and Text Marking to track associated masks. Objects are indexed numerically (e.g., "object-N"), and holistic (blue), individual (teal), and cluster (green) annotations are concatenated into a single image description. Paraphrasing instructions with combined description produce a concise, consistent GCG description that eliminates redundancy while preserving object-mask associations, even with reordering.



Figure 10. Qualitative results of GeoPixel's capability in Referring Expression Detection on VRSBench. Segmentation masks generated by GeoPixel are used to derive horizontal bounding boxes (shown in blue) and oriented bounding boxes (shown in red).