

# UniGeoSeg: Towards Unified Open-World Segmentation for Geospatial Scenes

## Supplementary Material

### 1. Overview

This supplementary material provides further details for the proposed dataset GeoSeg-1M and benchmark GeoSeg-Bench, as well as the proposed model UniGeoSeg. These details were omitted from the main body of the paper due to space constraints.

The supplementary material is organized as follows:

- Section 2. Details of the GeoSeg-1M Dataset Creation.
- Section 3. Statistics and Additional Samples of GeoSeg-1M.
- Section 4. More Details about GeoSeg-Bench.
- Section 5. More Details about Experiments.
- Section 6. Additional Examples of Model Predictions.
- Section 7. Evaluating UniGeoSeg with Alternative Language Model.
- Section 8 Datasheets for GeoSeg-1M and GeoSeg-Bench.

### 2. Details of the GeoSeg-1M Dataset Creation

#### 2.1. Data Sources of GeoSeg-1M

To construct GeoSeg-1M, we aggregate a wide range of publicly available remote-sensing datasets that provide pixel-level or region-level annotations across diverse spatial resolutions, geographic regions, and semantic categories. These datasets collectively supply the raw images and masks from which our unified multimodal corpus is built. Below, we summarize the data sources used in GeoSeg-1M and provide brief descriptions of each dataset.

**Chesapeake Land Cover** [20] provides a high-resolution (1 m) land-cover classification for the Chesapeake Bay watershed, covering regions in Maryland, Virginia, West Virginia, Pennsylvania, Delaware, and Washington D.C. The dataset includes raster layers derived from NAIP (Red, Green, Blue, NIR) imagery as well as Landsat-8 bands, with land-cover labels for classes such as water, tree canopy / forest, low vegetation, barren land, other impervious surfaces, roads, and no-data. The data was produced in collaboration between the Chesapeake Conservancy and USGS, and reflects coverage for multiple epochs (2013/2014, 2017/2018, and 2021/2022) with detailed land-use and land-cover change products.

**DeepGlobe** [7] dataset originates from the DeepGlobe 2018 Satellite Image Understanding Challenge and includes a land-cover segmentation track. For the segmentation / land-cover component, it provides 803 RGB satellite images of size  $2448 \times 2448$  pixels with a ground sampling distance of approximately 0.5 m. The annotation consists of seven land-cover classes, including urban, agriculture,

rangeland, forest, water, barren, and unknown.

**FLAIR** [8] is a country-scale land-cover semantic segmentation dataset for France. It is built from multi-source optical imagery with 20 cm ground sampling distance, and contains more than 20 billion labeled pixels across more than 817 km<sup>2</sup> of aerial acquisitions. In addition to spatial data, FLAIR integrates temporal and spectral information from satellite time series, providing fine-grained annotations for land-use monitoring and segmentation research.

**GID-15** [30] is a large-scale semantic segmentation dataset based on Gaofen-2 satellite images. It consists of 150 GF-2 images with pixel-level annotations covering 15 land-cover classes. The dataset provides two types of ground-truth format: .png (grayscale labels) and .tiff (RGB palette), to facilitate different usage scenarios.

**Globe230K** [22] is a globally distributed land-cover dataset consisting of approximately 232,819 image tiles of size 512×512 pixels. The tiles have a 1-m ground sampling distance and are annotated with a set of land-use/land-cover categories covering diverse geographic regions.

**LoveDA** [28] is a semantic segmentation dataset constructed from 0.3 m spatial-resolution images collected via the Google Earth platform. It covers both urban and rural regions, providing scenes of varied spatial layouts. The dataset includes six foreground classes, and offers pixel-level annotations for all categories.

**MiniFrance** [4] is a very-high-resolution aerial semantic segmentation dataset released for the IEEE Data Fusion Contest 2022 (DFC2022). It contains approximately 2,000 VHR images annotated with 12 land-use/land-cover classes, based on the Urban Atlas project. The training split of MiniFrance includes both labeled and unlabeled images to support semi-supervised learning.

**Potsdam** [1] is a benchmark aerial segmentation dataset consisting of 38 orthorectified tiles, each of size 6000×6000 pixels, acquired with a ground sampling distance of 5 cm. The images include four spectral bands (R, G, B, NIR), and the corresponding ground-truth labels cover six semantic categories: impervious surfaces, buildings, low vegetation, trees, cars, and background/clutter.

**Vaihingen** [1] is an aerial semantic segmentation dataset composed of 33 image patches, captured at a ground sampling distance of 9 cm. The dataset provides both true orthophotos and corresponding digital surface models, and its annotated classes include six categories: impervious surfaces, buildings, low vegetation, trees, cars, and clutter/background.

**Five-Billion-Pixels** [25] is a land-cover classification dataset derived from Gaofen-2 satellite imagery with a spa-

tial resolution of 4 m. It contains more than five billion labeled pixels distributed across 150 large-area scenes and is annotated with 24 land-cover classes.

**FAIR1M** [24] is a large-scale fine-grained object detection dataset in high-resolution remote sensing imagery. It contains over 15,000 images with a spatial resolution between 0.3 m and 0.8 m, collected from Gaofen-2 satellites and Google Earth. The dataset comprises more than 1 million object instances, annotated with rotated bounding boxes across 5 main categories (airplanes, ships, vehicles, courts, roads) and 37 sub-categories.

**DIOR** [10] is a large-scale object detection dataset comprising 23,463 remote-sensing images annotated with 192,472 object instances from 20 categories. The images exhibit variations in viewpoints, scenes, and spatial resolutions, and each object is annotated using axis-aligned bounding boxes.

**DOTA** [29] is an aerial image dataset for object detection containing 2,806 images with sizes ranging from 800×800 to 4000×4000 pixels. It includes 188,282 annotated object instances from 15 categories, where each object is labeled with a bounding box allowing arbitrary orientations.

For FAIR1M, DIOR, and DOTA, which provide region-level annotations in the form of bounding boxes rather than pixel-level masks, we convert the box annotations into segmentation masks using the procedure of SAMRS[27]. This conversion allows all region-level annotations from FAIR1M, DIOR, and DOTA to be integrated into our unified pixel-based segmentation framework.

**RRSIS-D** [14] consists of 17,402 image–text–mask triplets intended for the referring segmentation task in remote sensing. The images in RRSIS-D have a size of 800×800 pixels, covering a variety of ground sampling distances (from 0.5 m to 0.3 m), and the dataset includes 20 object categories described in natural-language expressions.

**EarthReason** [11] is a benchmark for a novel remote-sensing task called geospatial pixel reasoning. The dataset contains 5,434 manually annotated image masks and over 30,000 implicit question-answer pairs, where each question describes a target region in an image in a reasoning style rather than direct naming. The dataset supports implicit querying, requiring models to infer masks from contextual and spatial clues rather than explicit instructions.

**RemoteSAM-270K** [31] is a large-scale generalized referring expression segmentation dataset for remote sensing, consisting of 270K image-text-mask triplets. It covers 297 object categories and 16 fine-grained attribute types with an average of 3.17 attributes per sample, supporting multi-target, no-target, and single-target expressions.

## 2.2. Data Pre-processing of GeoSeg-1M

The raw data used to construct GeoSeg-1M originate from multiple remote sensing benchmarks with substantial varia-

tions in spatial resolution, sensor characteristics, annotation formats, and category taxonomies. These datasets typically provide large aerial images paired with semantic segmentation masks encoded in heterogeneous ways (*e.g.*, indexed masks, palette masks, or RGB-encoded labels). Such inconsistencies make unified processing and subsequent multimodal instruction generation difficult. To ensure compatibility across sources and create a standardized corpus suitable for multimodal segmentation, we perform a unified preprocessing pipeline on all images and masks.

**Patch Extraction.** We adopt a sliding-window strategy to extract fixed-size patches from both images and masks. Each image is divided into patches of  $512 \times 512$  pixels with a stride of 256. The corresponding masks are cropped using identical window parameters to maintain spatial alignment.

**Mask Standardization.** All annotations are converted into single-channel categorical masks with integer IDs. Regardless of the original encoding format, every pixel is mapped to a unified category index. This enables uniform downstream parsing while preserving the semantics of each dataset. We then consolidate all category label definitions into a unified taxonomy. Through synonym merging, we produce a standardized set of 117 semantic categories. This harmonization ensures that patches originating from different datasets share a consistent semantic space and can be used jointly for training a single segmentation model.

**Instance Candidate Extraction.** To support object-centric instruction-driven tasks such as referring segmentation, we further decompose the semantic masks into instance-level candidates. Since most remote sensing datasets provide class-wise semantic masks rather than instance annotations, we apply connected-component analysis (using 8-connectivity) to each semantic category map. Each connected region is treated as a potential instance for subsequent evaluation and instruction generation. Only minimal sanity checks—such as removing extremely tiny isolated regions or invalid labels—are applied at this stage; the full quality assessment is deferred to the mask filtering pipeline described in the next subsection.

This preprocessing stage yields a unified collection of standardized image patches, normalized masks, harmonized categories, and instance-level region proposals, which together form the basis for the subsequent high-quality mask filtering and instruction construction procedures.

## 2.3. Mask Filtering of GeoSeg-1M

After obtaining instance-level region proposals from connected-component decomposition, we perform a two-stage mask filtering pipeline to ensure that only high-quality, semantically meaningful regions are preserved for instruction generation. This process is designed to remove noisy, ambiguous, or low-value regions that commonly arise from fragmented masks and inconsistent labeling in

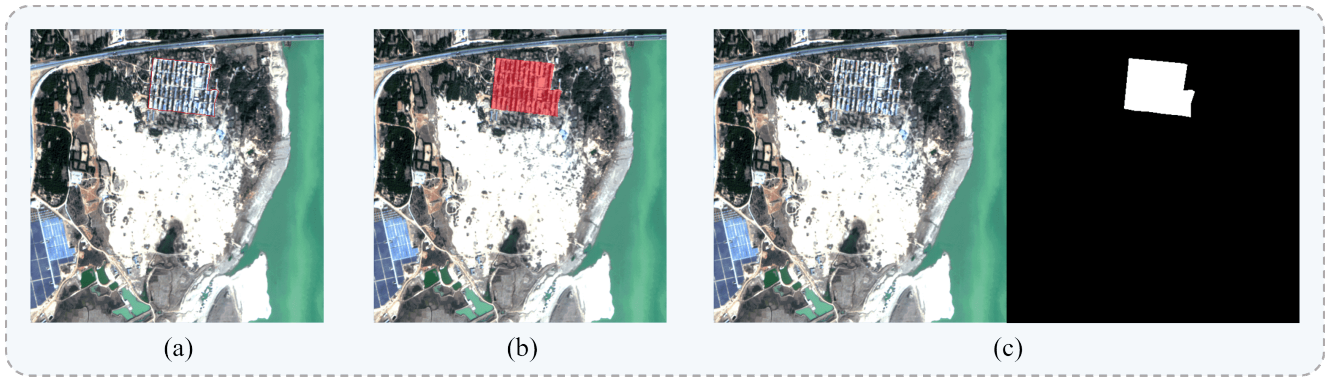


Figure 1. The three mask marking strategies we tried in model-based mask filtering. (a) Boundary-only highlight. (b) Semi-transparent filled-mask overlay. (c) The original image and the binary mask

**Prompt:** You are a remote sensing and vision-language expert. You are given an image where a candidate region is outlined with a thin contour. Assess whether this region is suitable to be used as an independent segmentation target. Evaluate the region according to the following integrated criteria:

- Semantic Soundness:  
The region should form a complete, meaningful object or land-cover unit, rather than a fragmented, noisy, or semantically incoherent area.
- Visual Distinctiveness:  
The region must be visually recognizable and separable from its surroundings based on shape, texture, or structural cues.  
Its approximate boundaries should appear interpretable and not excessively irregular.
- Unambiguous Localization  
The region should be uniquely identifiable within the image and not easily confusable with nearby similar patterns or other instances of the same category.

**\*\*Decision Rule:**  
If all criteria are fully satisfied → output “generate”.  
If any criterion fails → output “pass”.

**\*\*Output Format:**  
A single word: generate or pass.

Figure 2. The prompt of InternVL3 for mask filtering. The images marked with ✓ in the top-right corner are retained, as their mask regions exhibit clear semantic structure, while the images marked with × in the bottom-right corner are filtered out due to ambiguous semantics.

remote sensing datasets.

**Rule-Based Filtering.** In the first stage, we apply rule-based filtering guided by geometric and contextual constraints. Each instance mask must occupy between 0.5% and 70% of the patch area; extremely small regions provide insufficient visual cues while excessively large ones often cover entire landcover zones that are unsuitable for instruction-driven segmentation tasks. To avoid ambiguity in referring and reasoning tasks, instances from the same category are restricted as follows:

(i) Each region must not belong to a category with more

than six instances in the same patch, preventing confusion when many small, similar objects cluster together.

(ii) For each candidate region, no other instance of the same category may lie within a 15-pixel radius, ensuring that fragmented or closely adjacent subregions are not incorrectly treated as distinct objects.

(iii) To maintain dataset diversity and prevent over-representation of dominant classes, at most two instances per category in each patch are randomly sampled to enter the next evaluation stage. These constraints effectively suppress low-quality regions while preserving category variety

**Prompt:** You are a remote sensing and vision-language expert. You are tasked with generating a reasoning-based segmentation question and its corresponding justification for a remote sensing image.

The image contains one target region that is highlighted in red and belongs to the category `<class>`. It also includes other objects from different categories [`<class list>`].

Your goal is to produce a question that implicitly requires inferring the target's function, use, or intrinsic attributes (e.g., material, structure, status). The question should be non-obvious and reasoning-driven.

**\*\*Requirements:**

- Question

- Must include the word "segment" or the phrase "please segment".

- Should be based on one or more inferred aspects of the target's attribute(s), function(s), or role(s).

- Can be about typical or intended use (e.g., emergency access or play sport), or intrinsic characteristics (e.g., vegetation status, surface material).

- Do not mention visual markers like "red" or "highlighted".

- Do not reference or compare to the other categories.

- Be about 30 words and phrased naturally.

- Reason

- Max 30 words.

- Must state the region belongs to the `<class>` category.

- Justify how this region fulfills the functional or intrinsic criteria in the question.

- Do not mention visual markers like "red" or "highlighted".

**\*\*Output Format:**

Question: `<your inference-based segmentation prompt>`

Reason: `<your concise justification mentioning <class>>`

**Answer:**

Question: Please segment the area typically designed for high-density living spaces with potential access to local amenities and public transportation.

Reason: The region belongs to the `<urban_residential>` category, which supports dense living and access to essential services and transport.



Figure 3. The prompt of GPT-4o for attribute reasoning instruction generation.

across the dataset.

**Model-Based Filtering.** In the second stage, we perform a model-based assessment using InternVL3 to further evaluate whether each region can serve as a meaningful standalone segmentation target. Rather than feeding raw binary masks—which completely fill the region and obscure internal appearance—we generate an outline-based visualization that draws a thin contour around the candidate region while preserving all underlying image content. This representation keeps both the object and its surroundings fully visible, providing richer cues for quality assessment.

Before settling on this approach, we tried three mask marking strategies: (i) boundary-only highlighting, (ii) semi-transparent filled-mask overlays, and (iii) supplying the original image and the mask as separate inputs, as shown in Fig. 1. Human inspection and initial model responses indicated that boundary-only visualization offers the most reliable evaluations. Filled overlays tend to partially hide texture and boundary details, while separate-mask inputs reduce contextual awareness. In contrast, the outline-based representation maintains precise localization cues without compromising scene context. Using this visualization, InternVL3 is prompted to judge semantic soundness, visual

distinctiveness, and unambiguous localization of the region, based on the instruction shown in Fig. 2. Only regions validated by the model proceed to instruction construction.

Together, this two-stage filtering pipeline substantially improves annotation quality by removing ambiguous, noisy, or low-information regions and ensuring that the retained masks correspond to visually clear, contextually meaningful, and well-defined objects suitable for instruction-driven segmentation tasks.

## 2.4. Reasoning Instruction Generation

To construct high-quality reasoning instructions, we design a multi-stage pipeline that selects category-diverse images, differentiates attribute reasoning from contextual reasoning cases, and employs both GPT-4o and open-source VLMs for generation and cross-evaluation, respectively.

**Semantic Diversity Filtering.** We first ensure that the images entering this stage possess sufficiently diverse semantic content. For each dataset, we compute the empirical distribution of object categories based on the corresponding masks and rank images by the number of distinct categories they contain. Only images within the high-diversity tier of the distribution of each dataset are preserved for fur-

**Prompt:** You are a remote sensing and vision-language expert. In the remote sensing image, one object is highlighted in red and others in blue. Please identify 2 to 3 concise cues that clearly distinguish the red object from all blue ones.

Focus on spatial and contextual features, such as:

- **Proximity to landmarks** (e.g., near (or far away from) a canal or a fenced compound)
- **Neighborhood contrast** (e.g., located in a sparse area while others cluster tightly)
- **Embedded environment** (e.g., inside a residential block or surrounded by vegetation)
- **Relative spatial position** (e.g., at the southern edge of the cluster or between the highway and river)

Avoid relying solely on size, shape, or color unless they explain spatial or contextual differences. Composite relations like “between A and B” are allowed if they help differentiate the red object.

**Output format:**

Provide a numbered list of 2 to 3 short and clear phrases or clauses describing distinguishing features. Do not include explanations.

- 1.
- 2.
3. (optional)

**Answer:**

1. Located at the southern edge of the urban cluster.
2. Adjacent to a roundabout.
3. Surrounded by a more residential environment.




Figure 4. The prompt of GPT-4o for the first step of context reasoning instruction generation.

**Prompt:** You are a remote sensing and vision-language expert. You are given a remote sensing image and 2–3 short phrases that describe the function, use, or attributes of a target region in a remote sensing image. These clues should be interpreted together and grounded in the image context. Your task is to act as an expert of remote sensing, refer to the image and clues, and comprehensively generate a reasoning-based segmentation task.

The task is suggested to consist of two aspects. One aspect is spatial cues, inspired by the input phrases. The other aspect is attribute cues, generated based on image information and the category of target.

The target region is red-highlight in the image and it belongs to the category <class>, but its category name must not appear in the output. Instead, describe it with a reasoning-based expression that reflects its inferred role or use. For example:

“the place where large aircraft are stationed and serviced” (instead of “airport”),

“an area where goods are loaded and unloaded for transportation” (instead of “harbor”).

The expression must remain semantically meaningful, specific, and reasoning-driven.

Some input phrases may contain visual hints such as “red area” or “highlighted region.” You must ignore and exclude any such references in your final output.

**Requirements:**

--The question must include the word “segment” or the phrase “please segment”.

--Do not include any reference to visual cues such as “red”, “blue”, or “highlighted”.

--Do not mention or imply differences between the target and other objects or classes.

--The purpose of the task is to synthesize contextual clues from the image and implicitly refer to the target area.

--Use natural language and ensure the question is at least 25 words.

**Input phrases:**

<phrases>

**Output Format:**


Question: <a reasoning-based segmentation prompt that refers to the <class> region without naming it or comparing it to others>

Reason: <a concise and formal justification grounded in the region’s spatial configuration and functional cues>

**Answer:**

Question: Please segment the area located at the southern edge of the urban cluster, adjacent to a roundabout, where materials are processed or manufactured within a predominantly residential setting.

Reason: The region is industrial land situated at the southern boundary of the settlement and directly connected to a roundabout, and its layout indicates localized processing or production activity within the surrounding urban environment.



1. Located at the southern edge of the urban cluster.
2. Adjacent to a roundabout.
3. Surrounded by a more residential environment.

Figure 5. The prompt of GPT-4o for the second step of context reasoning instruction generation.

ther processing. Although exact thresholds differ due to dataset-specific statistics, they generally correspond to selecting roughly the top decile.

**Reasoning Type Assignment.** For every selected image,

we determine the reasoning type by examining how many regions belong to same semantic category. Images containing a single instance of the target category are treated as attribute reasoning cases, whereas images containing two

**Prompt:** You are a remote sensing and vision-language expert. Given a remote sensing image and a corresponding segmentation question + justification, your task is to evaluate the overall quality of the image-text pair from six distinct dimensions.

The image contains:

- One red-highlighted object (the intended segmentation target)
- One or more blue-highlighted distractors (belonging to the same category as the red object)
- All highlighted objects (red and blue) belong to the <class> category.

The red-highlighted region was automatically selected and may not always be meaningful or clean.

Your goal is to assess how well the question and justification identify and support this target through reasoning based on the image's visual content.

Please provide a numeric score (integer only) and a concise comment for **each** of the following six dimensions:

- Task Complexity & Implicitness (Score: 1–10)\*\***
  - Does the question involve multi-step reasoning?
  - Does it avoid directly revealing the red-highlighted object's location?
  - Is it indirect, scenario-driven, and non-trivial?
- Image-Text Consistency & Reasoning Support (Score: 1–10)\*\***
  - Can each condition in the question be clearly identified or inferred from the image?
  - Is there sufficient visual support for the reasoning process?
- Uniqueness of Target (Score: 1–10)\*\***
  - When combining all conditions, does the reasoning uniquely lead to the red-highlighted region? - Are blue-highlighted objects reasonably excluded?
- Mask Quality & Semantic Significance (Score: 1–10)\*\***
  - Is the red-highlighted region clearly bounded, visually meaningful, and semantically distinct (e.g., building, road, land parcel)? - Does it have clear real-world interpretation?
- Question Language Clarity & Naturalness (Score: 1–5)\*\***
  - Is the question grammatically correct, clear, and at least 30 words long? - Does it use natural language and contain the verb "segment" or "please segment"?
- Reason Quality & Conciseness (Score: 1–5)\*\***
  - Is the reason concise (≤30 words), accurate, and explicitly states that the target belongs to <class>? - Does it clearly explain why the target satisfies the question's conditions?

### Output Format:

Return the result as a JSON object with the following fields:

```
```json
{
  "task_complexity": {"score": 0, "comment": "<comment>"},
  "image_text_consistency": {"score": 0, "comment": "<comment>"},
  "target_uniqueness": {"score": 0, "comment": "<comment>"},
  "mask_quality": {"score": 0, "comment": "<comment>"},
  "question_clarity": {"score": 0, "comment": "<comment>"},
  "reason_quality": {"score": 0, "comment": "<comment>"}
}
```




Figure 6. The prompt of InternVL3 and QwenVL2 for evaluating the quality of reasoning image–mask–instruction triplets.

or three instances are designated as contextual reasoning cases. This separation enables us to distinguish samples that require reasoning about intrinsic properties from those involving relational or discriminative understanding.

**Attribute Reasoning Prompt.** For attribute reasoning samples, GPT-4o receives two inputs—the image with the target region outlined using a thin contour and the list of all categories present in the image—and is instructed to produce a question centered on attributes, functional roles, or other intrinsic characteristics. The corresponding prompt template is shown in Fig. 3.

**Context Reasoning Prompt.** Context reasoning samples involve a two-stage generation process. All same-category regions are highlighted using distinct colors to clearly separate them. GPT-4o first produces two to three explicit visual cues that distinguish the highlighted regions in terms of spatial arrangement, geometry, or appearance (prompt in Fig. 4). These cues, together with the original image, are then used in a second GPT-4o step, where the model generates a reasoning question and a concise explanation grounded in contextual or relational understanding (prompt in Fig. 5).

**Quality Filtering.** All resulting triplets of image–mask–instruction are further validated through cross-model evaluation using QwenVL2-72B and InternVL3-

78B. Both models score each sample along multiple dimensions, including task complexity and implicitness, image–text consistency and reasoning soundness, uniqueness of the target, mask quality and semantic significance, clarity and naturalness of the question, and conciseness of the explanation. The scoring prompt is provided in Fig. 6. Each dimension has a task-specific threshold, and only samples that exceed all thresholds across both evaluators are retained. After this filtering, the final reasoning subset contains 104,989 samples, including 63,567 attribute reasoning samples and 41,422 context reasoning samples.

## 2.5. Referring Instruction Generation

To construct the referring segmentation subset, we employ a unified prompting template that guides GPT-4o to produce concise and unambiguous referential expressions for each candidate region. As shown in Fig. 7, the prompt provides the image patch, the semantic class of the target region, and a strict output format beginning with “Question:”. The only emphasized requirement in the template is that the expression must clearly and uniquely identify the region based on interpretable cues.

A notable effect of this design is that the generated instructions naturally rely on contextual and spatial information—such as relative position, surrounding structures, or

**Prompt:** You are an AI assistant and a remote sensing expert. I am currently building a dataset for referential expression segmentation in remote sensing imagery. Your task is to generate a clear and specific referential expression (i.e., a question or instruction) that precisely directs attention to a specific region within the image. The region is highlighted with a red outline in the image and is known to correspond to the class <class>

For example:

"Segment the forest in the upper-left corner of the image."

"Please segment the red-colored house at the bottom."

Guidelines:

The red outline is for dataset construction only. Do not mention, imply, or rely on it in the question.

The expression should be unambiguous and precise. It may reference attributes such as location, color, shape, or size to help clearly identify the region. Avoid vague references.

Only output the referential expression in the following format:

Question: [Your referential expression here]

**Answer:**

Question: Segment the large passenger ship docked along the waterfront near the industrial area



Figure 7. The prompt of GPT-4o for referring instruction generation.

**Prompt:** You are a remote-sensing expert. Given an image and a referring-segmentation instruction, evaluate how well the instruction uniquely and accurately identifies the red-outlined target object.

Provide an integer score and a brief comment for each of the following three dimensions:

1. Grounding Accuracy (1–10)

- Are the spatial/attribute cues correct, image-grounded, and unambiguous?

2. Target Uniqueness (1–10)

- Do the cues reliably single out the target rather than the distractors?

3. Language Quality (1–5)

- Is the instruction clear, natural, and does it contain “segment” or “please segment”?

Output JSON only:

```
{
  "grounding_accuracy": {"score": 0, "comment": "<comment>"},
  "target_uniqueness": {"score": 0, "comment": "<comment>"},
  "language_quality": {"score": 0, "comment": "<comment>"}
}
```



Figure 8. The prompt of InternVL3 and QwenVL2 for evaluating the quality of referring image–mask–instruction triplets.

functional cues—rather than superficial attributes like color or size. This leads to more meaningful and semantically grounded referring expressions, consistent with real-world remote sensing interpretation.

Following instruction generation, all samples undergo cross-model validation using InternVL3-78B and QwenVL2-72B. Both models independently evaluate each instruction–image pair using the scoring prompt shown in Fig. 8, assessing clarity, grounding, consistency, and uniqueness of the reference. Only samples that satisfy the required thresholds across both evaluators are retained. After this stage, the final referring subset contains 336,311 high-quality region–instruction pairs.

## 2.6. Interactive Instruction Generation

For interactive segmentation, we generate point- and box-based instructions from mask annotations. For each mask

region, we sample  $k \in \{1, 2, 3\}$  points and normalize their coordinates to the  $[0, 1]$  range along both image axes. The number of points  $k$  is determined adaptively according to the mask size: regions with fewer than 200 pixels are assigned a single point, whereas larger regions follow a distribution in which single-point prompts account for 60% of cases, and two- or three-point prompts each account for 20%. Denoting the normalized coordinates of the sampled points as  $(x_i, y_i)$ ,  $i = 1, \dots, k$ , the points are inserted into a fixed template: “Please segment the region/target corresponding to the points  $\{(x_1, y_1), \dots, (x_k, y_k)\}$ .”

For each mask, we compute the tight bounding box defined by the top-left corner  $(x_{\min}, y_{\min})$  and bottom-right corner  $(x_{\max}, y_{\max})$ , and normalize all coordinates to  $[0, 1]$ . These coordinates are then inserted into the template: “Please segment the region/target corresponding to the box  $x_0, y_0=[x_{\min}, y_{\min}]$ ,  $x_1, y_1=[x_{\max}, y_{\max}]$ .”

Table 1. Statistics of GeoSeg-1M.

Subset	Samples (generated)	Samples (total)	Avg Text Length	Categories	Average Mask Size	Other Statistics
Referring	336,311	560,179	12.05 words	100	18.4 K	-
Interactive	480,949	480,949	9.80 words	106	15.5 K	Point:242,488 Box: 238,461
Reasoning	104,989	119,214	23.93 words	108	14.8 K	Attribute: 63,567 Contextual: 41,422
Overall	922,249	1,148,504	12.18 words	117	17.0 K	Resolution: 0.05–153 m Sources <sup>1</sup> : C, DG, FB, FL, GI, GL, L M, P, Va, FA, DO, DI, EarthReason [11] RemoteSAM [31], RRSIS-D [14]

1 C:Chesapeake [20], DG:DeepGlobe [7], FB:Five-Billion-Pixel [25], FL:FLAIR [8], GI:GID-15 [30], GL:Globe230K [22], L:LoveDA [28], M:MiniFrance [4], P:Potsdam [1], Va:Vaihingen [1], FA:FAIR1M [24], DO:DOTA [29], DI:DIOR [10].

This procedure produces 480,949 interactive samples, providing consistent and structured supervision for both point- and box-driven segmentation tasks.

### 3. Statistics and Additional Samples of GeoSeg-1M

Tab. 1 presents detailed statistics of the GeoSeg-1M dataset. The dataset is divided into three subsets corresponding to the main segmentation tasks: referring, interactive, and reasoning. For each subset, we report the number of samples generated in our pipeline, as well as the total number of samples including integrated external datasets. The average text length indicates the typical instruction length for each task, while the number of categories reflects the semantic diversity covered. We also provide the average mask size per subset and task-specific statistics, such as point/box distribution for interactive samples and attribute/context counts for reasoning samples. Overall, GeoSeg-1M contains more than 1.14 million image–mask–instruction triplets spanning 117 categories, aggregated from multiple remote sensing sources. The semantic categories are listed in Tab. 7, and their overall quantity distribution is shown in Fig. 9. The distributions of categories for the reasoning, referring, and interactive tasks are also presented separately.

To complement the main paper, Fig. 11 presents additional GeoSeg-1M samples that highlight the visual richness and instruction diversity of the dataset.

### 4. More details about GeoSeg-Bench

GeoSeg-Bench comprises 6,892 samples organized into three subsets corresponding to reasoning, interactive, and referring segmentation tasks. The reasoning subset contains 1,711 samples, including 1,150 attribute reasoning instances and 561 contextual reasoning instances. The interactive subset includes 2,870 samples, with 1,435 samples using point-based prompts and 1,435 samples using box-based prompts. The referring subset consists of 2,311 samples. The benchmark covers 102 semantic categories. The

full list of categories and the corresponding quantity distributions are provided in Tab. 8. These statistics illustrate the overall diversity and balance of GeoSeg-Bench, supporting comprehensive evaluation across reasoning, referring, and interactive segmentation tasks.

Tab. 2 reports instruction diversity on GeoSeg-1M and GeoSeg-Bench, excluding interactive template instructions. Referring instructions contain common words but maintain clear syntactic diversity, while reasoning instructions exhibit substantially higher linguistic diversity. The intra-class n-gram and syntactic diversity (Tab. 2, Row 3) indicates no obvious shortcut reasoning driven by superficial language patterns. In addition, Fig. 10 shows a failure case in hard negative reasoning scene. Our model focuses on the *northern boundary* and *living zone* while ignoring the crucial cues *close to an open green space* and *near a road intersection*, which leads to a fail case.

Table 2. N-gram and Syntactic Diversity Statistics (%).


	2-gram		3-gram		Syntactic (POS)	
	Top50	Top100	Top50	Top100	Top50	Top100
GeoSeg-1M Referring	53.89	64.56	37.89	47.85	36.83	43.23
GeoSeg-1M Reasoning	24.39	30.35	13.36	17.70	1.00	1.60
GeoSeg-1M Class Rural-Residential (398 samples)	34.43	44.73	20.40	27.70	14.07	26.63
GeoSeg-Bench Referring	56.46	68.57	41.25	52.39	24.59	34.45
GeoSeg-Bench Reasoning	26.87	34.01	14.95	19.69	3.48	6.28

To further illustrate the composition and annotation style of GeoSeg-Bench, we present additional representative samples in Fig. 12.


## 5. More Details about Experiments

### 5.1. Implementation Details for UniGeoSeg

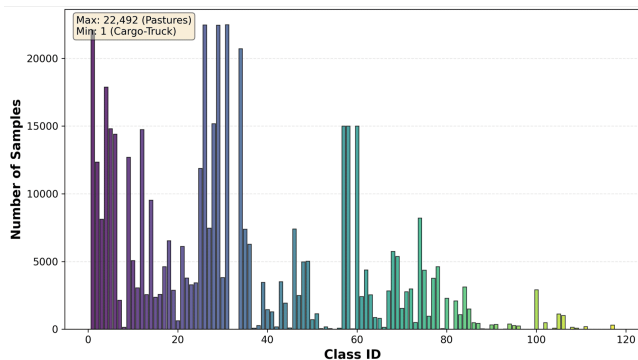
**The subset with DIOR-derived data excluded.** To better examine the generalization behavior of UniGeoSeg under unfamiliar data sources, we additionally construct a subset of GeoSeg-1M by removing all samples originating from the DIOR [10] family of datasets. The resulting sub-




(a) Class distribution of all data in GeoSeg-1M



(b) Class distribution of reasoning segmentation task in GeoSeg-1M



(c) Class distribution of referring segmentation task in GeoSeg-1M



(d) Class distribution of interactive segmentation task in GeoSeg-1M

Figure 9. Class distribution in GeoSeg-1M. (a) Overall class distribution across the entire dataset. (b–d) Class distributions for each specific task: (b) Reasoning segmentation, (c) Referring segmentation, and (d) Interactive segmentation.




Figure 10. Example of a hard fail case.

set contains 970,238 samples, comprising 109,274 reasoning, 431,204 referring, and 429,760 interactive instances. Since the downstream benchmarks SIOR [27] and DIOR-RSVG [32] are themselves derived from DIOR, excluding DIOR-based training data mitigates this source-specific overlap and provides a more independent setting for evaluating the model’s cross-task generalization.

**The subset used in ablations.** To conduct the ablation studies of task-adaptive text enhancement (TATE) and latent knowledge memory (LKM) under a manageable computational budget, we further construct a compact subset of 141,132 samples within the DIOR-excluded por-

tion of GeoSeg-1M. This subset comprises 42,699 reasoning, 55,481 referring, and 42,952 interactive instances randomly sampled from the remaining data. We then evaluate the resulting models on the reasoning and referring tasks of GeoSeg-Bench, as well as on zero-shot interactive segmentation on SIOR, enabling a focused examination of how each module affects both the accuracy and the generalization capability of UniGeoSeg.

**Training Configuration.** All experiments are conducted with bfloat16 precision on eight NVIDIA A800 GPUs. Input images are resized to  $512 \times 512$ . We use the AdamW optimizer with an initial learning rate of  $1 \times 10^{-4}$ , a cosine decay schedule, and a warmup ratio of 0.03. Weight decay is set to 0.0. The progressive task scheduling (PTS) strategy is applied: the sampling weight of interactive segmentation gradually decreases to 0.7, with the remaining weight assigned to reasoning samples. Batch size is set to 16, and each experiment is trained for 3 epochs. Data loading uses 4 workers per GPU, and lazy preprocessing is enabled to accelerate training. Additional hyper-parameters and implementation details are summarized in Tab. 3.

Table 3. Supplementary training hyper-parameters and setup.

Parameter	Value	Parameter	Value
Precision	bfloat16	GPUs	8 × NVIDIA A800
Image Size	512 × 512	Batch Size	16
Training Epochs	3	Optimizer	AdamW
Initial LR	$1 \times 10^{-4}$	LR Schedule	Cosine decay
Warmup Ratio	0.03	Weight Decay	0.0
Gradient Checkpointing	Enabled	Data Loader Workers	4
PTS Sampling	Decays to 0.7	ZeRO Stage	ZeRO 2

## 5.2. Comparison with Other Architectures

We compare UniGeoSeg with several baseline models in terms of their visual encoder, language backbone, and the use of dual visual encoders. Notably, some baselines incorporate Segment Anything Model (SAM), employing a dual-encoder setup: one encoder provides image features to the segmentation decoder, while a second encoder aligned with the LLM backbone supplies features for language-guided understanding. UniGeoSeg, in contrast, does not use SAM and relies on a single visual encoder. Table 4 summarizes these architectural differences across models.

Building on this comparison, UniGeoSeg employs a streamlined and efficient architecture. It relies on a single Swin-B[15] visual encoder, reducing structural redundancy and computational overhead while maintaining strong visual feature representation. The language backbone of UniGeoSeg is relatively small, which further contributes to efficiency. Despite not using SAM encoder, UniGeoSeg achieves core SAM-like functionality by integrating interactive segmentation instructions and embedding visual prompts directly into the textual input space during training. Consequently, UniGeoSeg provides a compact and lightweight architecture that captures both language-guided and pixel-level cues, demonstrating advantages over both dual-encoder models and other single-encoder baselines.

## 5.3. Fine-tuning Protocol

We summarize fine-tuning protocols on GeoSeg-1M for all baselines, as shown in Tab. 5.

## 6. Additional Examples of Model Predictions

This section presents additional qualitative results to further illustrate the behavior of the model in various scenarios, as shown in Fig. 13.

## 7. Evaluating UniGeoSeg with Alternative Language Model

In this section, we examine the flexibility of UniGeoSeg by replacing its original Phi-1.5 [12] language backbone with DeepSeek-7B-Chat [2]. The model is trained on

GeoSeg-1M under the same settings as the main configuration and evaluated on the EarthReason [11] and RRSIS-D [14] benchmarks. While switching the language backbone naturally introduces some variation in overall performance, the DeepSeek-based variant still maintains competitive segmentation quality.

To further validate the general applicability of our proposed modules TATE, LKM, and PTS, we perform additional ablation studies using an alternative backbone. Following identical training and evaluation protocols, we remove each module in turn and report the results in Tab. 6. The consistent performance gains observed for UniGeoSeg across different language model architectures indicate that these modules are largely model-agnostic and provide robust improvements independent of the underlying LLM.

## 8. Datasheets

### 8.1. Motivation

1. “*For what purpose was the dataset created?*”

**A:** The GeoSeg-1M and GeoSeg-Bench datasets were created to address the lack of large-scale, high-quality multimodal segmentation data in remote sensing, a bottleneck that limits the development of unified vision–language segmentation models. GeoSeg-1M provides a million-level collection of image–mask–instruction triplets covering referring, reasoning, and interactive segmentation, enabling instruction-driven multimodal understanding at scale. GeoSeg-Bench complements it with a carefully designed evaluation suite for assessing fine-grained reasoning ability, context-aware grounding, and interactive segmentation performance. Together, these two datasets aim to support the training and rigorous benchmarking of general-purpose remote sensing segmentation models capable of robust instruction following and comprehensive scene understanding.

2. “*Who created the dataset (e.g., which team, research group) and on behalf of which entity?*”

**A:** The dataset was created by the following authors:

- Anonymous authors

3. “*Who funded the creation of the dataset?*”

Table 4. Comparison of UniGeoSeg with baseline model architectures. The “Encoder Configuration” column denotes the visual encoder setup: Dual-encoder indicates a SAM-based encoder for the segmentation decoder plus a second LLM-aligned encoder, while Single-encoder indicates only one visual encoder is used.

Method	Vision Encoder	LLM Type	Encoder Configuration
LISA[9]	CLIP-L[18]	Vicuna-7B[6]	Dual-encoder
PixelLM[19]	CLIP-L	Vicuna-7B	Dual-encoder
PSALM[33]	Swin-B[15]	Phi-1.5(1.3B)[12]	Single-encoder
Geopixel[21]	CLIP-L	InternLM2-7B[3]	Dual-encoder
Geopix[16]	CLIP-L	Llama-2-7B[26]	Single-encoder
Earthmind[23]	InternViT[5]	InternVL2-4B[5]	Dual-encoder
LISAT[17]	RemoteCLIP[13]	Vicuna-7B	Dual-encoder
Segearth-R1[11]	Swin-B	Phi-1.5(1.3B)	Single-encoder
Ours (UniGeoSeg)	Swin-B	Phi-1.5(1.3B)	Single-encoder

Table 5. Fine-tuning protocol for all baselines on GeoSeg-1M.

Method	Epoch	Resolution	Encoder	Optimizer	Frozen	Finetune	LR	Batch
PSALM[33]	3	1024×1024	Single	AdamW	vision encoder	LLM, mask decoder	1e-4	8
GeoPixel[21]	3	560×560	Dual	AdamW	vision encoder, SAM encoder, LLM	vision/text projector mask decoder	3e-4	8
Earthmind[23]	3	512×512	Dual	Adam	vision encoder, SAM encoder	LLM-lora, projector, mask decoder	4e-5	16
LISAT[17]	3	512×512	Dual	AdamW	vision encoder	LLM-lora, projector, mask decoder	3e-4	16
SegEarth-R1[11]	3	1024×1024	Single	AdamW	vision encoder	LLM, mask decoder	1e-4	8
UniGeoSeg	3	512×512	Single	AdamW	vision encoder	LLM, mask decoder	1e-4	8

**A:** The dataset creation was funded by the affiliations of the authors involved in this work.

## 8.2. Composition

Most of the questions in this section are intended to provide dataset consumers with the information they need to make informed decisions about using the dataset for their chosen tasks. Some of the questions are designed to elicit information about compliance with the EU’s General Data Protection Regulation (GDPR) or comparable regulations in other jurisdictions. Questions that apply only to datasets that relate to people are grouped together at the end of the section. We recommend taking a broad interpretation of whether a dataset relates to people. For example, any dataset containing text that was written by people relates to people.

1. “What do the instances that comprise our datasets represent (e.g., documents, photos, people, countries)?”

**A:** The dataset primarily consists of remote sensing images captured by satellites and drones, along with their corresponding textual annotations. All datasets utilized in GeoSeg-1M and GeoSeg-Bench are publicly accessible and nonprofit.

2. “How many instances are there in total (of each type, if appropriate)?”

**A:** GeoSeg-1M includes 1,148,504 image-mask-instruction triplets. Details could be found in the main text. GeoSeg-Bench consists of 6,892 samples, including 2,870 interactive, 2,311 referring, and 1,711

reasoning segmentation samples.

3. “Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?”

**A:** The images in GeoSeg-1M and GeoSeg-Bench are sourced from existing detection and segmentation datasets. Except for the samples from RemoteSAM [31], RRSIS-D [14], and EarthReason [11] in GeoSeg-1M, all textual annotations were independently created by us.

4. “Is there a label or target associated with each instance?”

**A:** Yes, for these images, we have provided image-mask-instruction triplets instances.

5. “Is any information missing from individual instances?”

**A:** No, each individual instance is complete.

6. “Are relationships between individual instances made explicit (e.g., users’ movie ratings, social network links)?”

**A:** Yes, the relationship between individual instances is explicit.

7. “Are there recommended data splits (e.g., training, development/validation, testing)?”

**A:** The GeoSeg-1M is designed to train the RS MLLMs for instruction-driven segmentation, and the GeoSeg-Bench is designed to evaluation.

8. “Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)?”

Table 6. Results on EarthReason and RRSIS-D of UniGeoSeg-DeepSeek7B.

Method	EarthReason (Val)		EarthReason (Test)		RRSIS-D	
	cIoU	gIoU	cIoU	gIoU	cIoU	gIoU
UniGeoSeg-DeepSeek7B (w/o TATE, LKM, and PTS)	67.37	65.03	65.92	64.23	71.52	60.02
UniGeoSeg-DeepSeek7B	71.41 (+4.04)	68.88 (+3.85)	71.93 (+6.01)	69.17 (+4.94)	75.12 (+3.60)	65.75 (+5.73)

**A:** GeoSeg-1M and GeoSeg-Bench are self-contained and will be open-sourced on platforms like Hugging Face for easy use.

9. “Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor–patient confidentiality, data that includes the content of individuals’ non-public communications)?”

**A:** No, all data are clearly licensed.

10. “Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety?”

**A:** No, GeoSeg-1M and GeoSeg-Bench do not contain any data with negative information.

### 8.3. Collection Process

In addition to the goals outlined in the previous section, the questions in this section are designed to elicit information that may help researchers and practitioners create alternative datasets with similar characteristics. Again, questions that apply only to datasets that relate to people are grouped together at the end of the section.

1. “How was the data associated with each instance acquired?”

**A:** The images in GeoSeg-1M and GeoSeg-Bench are sourced from existing detection and segmentation datasets. We enrich these with annotations. Details are shown in the Section 3 in main text.

2. “What mechanisms or procedures were used to collect the data (e.g., hardware apparatuses or sensors, manual human curation, software programs, software APIs)?”

**A:** GeoSeg-1M is constructed by integrating multiple publicly available remote sensing datasets collected from airborne and satellite imaging platforms with diverse spatial resolutions. All images and pixel-level masks originate from their respective sources without additional sensor deployment. Data consolidation, preprocessing, and annotation standardization were performed through automated software pipelines, including unified format conversion, image tiling, connected-region decomposition, and large-scale mask filtering. Instruction annotations were generated using GPT-4o through API-based prompting, followed by cross-

evaluation with open-source vision-language models. No manual labeling was introduced beyond quality verification of a small subset for prompt and pipeline validation. GeoSeg-Bench was further curated by two domain experts in remote sensing, who independently reviewed and cross-validated candidate samples to ensure high-quality and unambiguous ground truth.

3. “If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?”

**A:** Please refer to the details listed in the main text Section 3.

### 8.4. Preprocessing, Cleaning, and Labeling

The questions in this section are intended to provide dataset consumers with the information they need to determine whether the “raw” data has been processed in ways that are compatible with their chosen tasks. For example, text that has been converted into a “bag-of-words” is not suitable for tasks involving word order.

1. “Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?”

**A:** Yes. Extensive preprocessing and cleaning were applied to unify heterogeneous remote sensing datasets. All images were tiled into  $512 \times 512$  patches with a stride of 256, and masks were converted into a unified binary format. Each mask was decomposed into connected regions to isolate individual objects. A two-stage filtering pipeline was used to remove low-quality or noisy regions, combining rule-based screening with model-based quality assessment. Categories across datasets were merged into a harmonized taxonomy. Instruction annotations were automatically generated and subsequently cross-evaluated by large vision–language models. No manual relabeling was added beyond limited spot-checking for validation.

2. “Was the ‘raw’ data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)?”

**A:** Yes, raw data is accessible.

3. “Is the software that was used to preprocess/clean/label

*the data available?”*

**A:** Yes, the necessary software used to preprocess and clean the data is publicly available.

## 8.5. Uses

The questions in this section are intended to encourage dataset creators to reflect on tasks for which the dataset should and should not be used. By explicitly highlighting these tasks, dataset creators can help dataset consumers make informed decisions, thereby avoiding potential risks or harms.

1. *“Has the dataset been used for any tasks already?”*

**A:** No.

2. *“Is there a repository that links to any or all papers or systems that use the dataset?”*

**A:** Yes, we will provide such links in the GitHub and the Huggingface repository.

3. *“What (other) tasks could the dataset be used for?”*

**A:** GeoSeg-1M provides extensive annotations for instruction-driven segmentation tasks. It could be used to train the MLLMs. GeoSeg-Bench provides high-quality samples in interactive, referring, and reasoning segmentation tasks. It could be used to evaluate the MLLMs.

4. *“Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?”*

**A:** No.

5. *“Are there tasks for which the dataset should not be used?”*

**A:** N/A.

## 8.6. Distribution

Dataset creators should provide answers to these questions prior to distributing the dataset either internally within the entity on behalf of which the dataset was created or externally to third parties.

1. *“Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?”*

**A:** The datasets will be made publicly accessible to the research community.

2. *“How will the dataset be distributed (e.g., tarball on website, API, GitHub)?”*

**A:** We will provide GeoSeg-1M and GeoSeg-Bench in the GitHub and the Huggingface repository.

3. *“When will the dataset be distributed?”*

**A:** We will create a repository to release the data once the paper is officially published.

4. *“Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?”*

**A:** Yes, the dataset will be released under the Creative

Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

5. *“Have any third parties imposed IP-based or other restrictions on the data associated with the instances?”*

**A:** No.

6. *“Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?”*

**A:** No.

## 8.7. Maintenance

As with the questions in the previous section, dataset creators should provide answers to these questions prior to distributing the dataset. The questions in this section are intended to encourage dataset creators to plan for dataset maintenance and communicate this plan to dataset consumers.

1. *“Who will be supporting/hosting/maintaining the dataset?”*

**A:** The authors of this work serve to support, host, and maintain the datasets.

2. *“How can the owner/curator/manager of the dataset be contacted (e.g., email address)?”*

**A:** They can be contacted via the email addresses listed on the paper or webpage.

3. *“Is there an erratum?”*

**A:** There is no explicit erratum; updates and known errors will be specified in future versions.

4. *“Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?”*

**A:** Future updates (if any) will be posted on the dataset website.

5. *“Will older versions of the dataset continue to be supported/hosted/maintained?”*

**A:**

Yes. This initial release will be updated in the future, with older versions replaced as new updates are posted.

6. *“If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?”*

**A:** Yes, we will provide detailed instructions for future extensions.

Table 7. The 117 semantic categories included in GeoSeg-1M.

---

1. Urban_Land	2. Agriculture_Land	3. Rangeland
4. Forest_Land	5. Water	6. Barren_Land
7. Low Vegetation / Field	8. Impervious (Other)	9. Road
10. Industrial_Land	11. Urban_Residential	12. Rural_Residential
13. Paddy_Field	14. Irrigated_Land	15. Dry_Cropland
16. Garden_Plot	17. Arbor_Woodland	18. Shrub_Land
19. Natural_Grassland	20. Artificial_Grassland	21. River
22. Lake	23. Pond	24. Background
25. Building	26. Industrial/Commercial/Public/ Military/Private/Transport Units	27. Mine/Dump/Construction Sites
28. Artificial Non-Agricultural Vegetated Areas	29. Arable Land (Annual Crops)	30. Permanent Crops
31. Pastures	32. Complex/Mixed Cultivation Patterns	33. Orchards at Urban Fringe
34. Herbaceous Vegetation Associations	35. Open Spaces w/ Little Vegetation	36. Wetlands
37. Clouds/Shadows	38. Park	39. Snow
40. Fish Pond	41. Stadium	42. Square
43. Overpass	44. Railway Station	45. Airport
46. Pervious Surface	47. Coniferous	48. Deciduous
49. Brushwood	50. Vineyard	51. Plowed Land
52. Swimming Pool	53. Clear Cut	54. Mixed
55. Ligneous	56. Greenhouse	57. Cropland
58. Grass	59. Tundra	60. Impervious
61. Low Vegetation	62. Tree	63. Car
64. Large-Vehicle	65. Swimming-Pool	66. Helicopter
67. Bridge	68. Plane	69. Ship
70. Soccer-Ball-Field	71. Basketball-Court	72. Ground-Track-Field
73. Small-Vehicle	74. Baseball-Diamond	75. Tennis-Court
76. Roundabout	77. Storage-Tank	78. Harbor
79. Container-Crane	80. Airport	81. Helipad
82. Chimney	83. Expressway-Service-Area	84. Expressway-Toll-Station
85. Dam	86. GolfField	87. Windmill
88. A220	89. A321	90. A330
91. A350	92. ARJ21	93. Boeing737
94. Boeing747	95. Boeing777	96. Boeing787
97. Bus	98. C919	99. Cargo-Truck
100. Dry-Cargo-Ship	101. Dump-Truck	102. Engineering-Ship
103. Excavator	104. Fishing-Boat	105. Intersection
106. Liquid-Cargo-Ship	107. Motorboat	108. Other-Airplane
109. Other-Ship	110. Other-Vehicle	111. Passenger-Ship
112. Tractor	113. Trailer	114. Truck-Tractor
115. Tugboat	116. Van	117. Warship

---

Table 8. Category Statistics of GeoSeg-Bench

ID	Name	Count (%)	ID	Name	Count (%)	ID	Name	Count (%)
1	Urban_Land	37 (0.53%)	2	Agriculture_Land	207 (2.97%)	3	Rangeland	24 (0.34%)
4	Forest_Land	97 (1.39%)	5	Water	260 (3.74%)	6	Barren_Land	226 (3.25%)
7	Low Vegetation / Field	51 (0.73%)	8	Impervious (Other)	2 (0.03%)	9	Road	478 (6.87%)
10	Industrial_Land	329 (4.73%)	11	Urban_Residential	196 (2.82%)	12	Rural_Residential	378 (5.43%)
13	Paddy_Field	118 (1.70%)	14	Irrigated_Land	398 (5.72%)	15	Dry_Cropland	77 (1.11%)
16	Garden_Plot	101 (1.45%)	17	Arbor_Woodland	176 (2.53%)	18	Shrub_Land	69 (0.99%)
19	Natural_Grassland	92 (1.32%)	20	Artificial_Grassland	82 (1.18%)	21	River	240 (3.45%)
22	Lake	85 (1.22%)	23	Pond	149 (2.14%)	24	Background	11 (0.16%)
25	Building	154 (2.21%)	26	Industrial, Commercial, Public, Military, Private And Transport Units	51 (0.73%)	27	Mine, Dump And Construction Sites	17 (0.24%)
28	Artificial Non-Agricultural Vegetated Areas	30 (0.43%)	29	Arable Land (Annual Crops)	29 (0.42%)	30	Permanent Crops	18 (0.26%)
31	Pastures	50 (0.72%)	34	Herbaceous Vegetation Associations	31 (0.45%)	35	Open Spaces With Little Or No Vegetation	24 (0.34%)
36	Wetlands	31 (0.45%)	38	Park	10 (0.14%)	39	Snow	15 (0.22%)
40	Fish Pond	13 (0.19%)	41	Stadium	34 (0.49%)	42	Square	13 (0.19%)
43	Overpass	76 (1.09%)	44	Railway Station	27 (0.39%)	45	Airport	3 (0.04%)
46	Pervious Surface	29 (0.42%)	47	Coniferous	10 (0.14%)	48	Deciduous	29 (0.42%)
49	Brushwood	19 (0.27%)	50	Vineyard	20 (0.29%)	51	Plowed Land	19 (0.27%)
52	Swimming Pool	37 (0.53%)	53	Clear Cut	3 (0.04%)	56	Greenhouse	8 (0.11%)
57	Cropland	92 (1.32%)	58	Grass	67 (0.96%)	60	Impervious	59 (0.85%)
61	Low Vegetation	15 (0.22%)	62	Tree	50 (0.72%)	63	Car	62 (0.89%)
64	Large-Vehicle	134 (1.93%)	65	Swimming-Pool	37 (0.53%)	66	Helicopter	5 (0.07%)
67	Bridge	67 (0.96%)	68	Plane	304 (4.37%)	69	Ship	151 (2.17%)
70	Soccer-Ball-Field	65 (0.93%)	71	Basketball-Court	72 (1.03%)	72	Ground-Track-Field	47 (0.68%)
73	Small-Vehicle	65 (0.93%)	74	Baseball-Diamond	85 (1.22%)	75	Tennis-Court	245 (3.52%)
76	Roundabout	53 (0.76%)	77	Storage-Tank	80 (1.15%)	78	Harbor	199 (2.86%)
79	Container-Crane	6 (0.09%)	80	Airport	45 (0.65%)	82	Chimney	20 (0.29%)
83	Expressway-Service-Area	8 (0.11%)	84	Expressway-Toll-Station	55 (0.79%)	85	Dam	19 (0.27%)
86	GolfField	10 (0.14%)	87	Windmill	3 (0.04%)	90	A330	4 (0.06%)
94	Boeing747	3 (0.04%)	95	Boeing777	1 (0.01%)	96	Boeing787	1 (0.01%)
97	Bus	3 (0.04%)	99	Cargo-Truck	15 (0.22%)	100	Dry-Cargo-Ship	32 (0.46%)
101	Dump-Truck	14 (0.20%)	102	Engineering-Ship	4 (0.06%)	103	Excavator	4 (0.06%)
104	Fishing-Boat	2 (0.03%)	105	Intersection	22 (0.32%)	106	Liquid-Cargo-Ship	13 (0.19%)
107	Motorboat	7 (0.10%)	108	Other-Airplane	1 (0.01%)	110	Other-Vehicle	1 (0.01%)
111	Passenger-Ship	10 (0.14%)	113	Trailer	3 (0.04%)	114	Truck-Tractor	3 (0.04%)
115	Tugboat	1 (0.01%)	116	Van	5 (0.07%)	117	Warship	6 (0.09%)

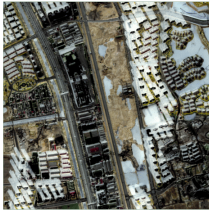


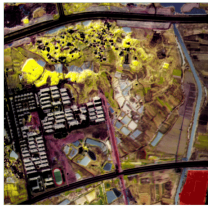
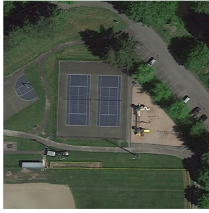





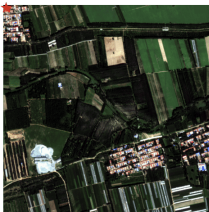
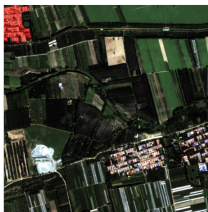
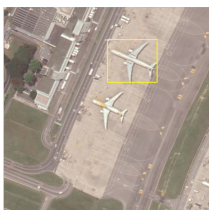
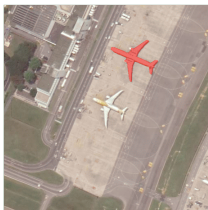
		Image	Ground Truth
<b>Attribute Reasoning:</b>	<p><b>Question:</b> Cactivities or public gatherings based on its spatial layout and structural attributes.</p> <p><b>Answer:</b> The region belongs to the &lt;Square&gt; category, typically used for public events or recreation due to its open and organized structure.</p>		
<b>Attribute Reasoning:</b>	<p><b>Question:</b> Please segment the area designed for manufacturing processes or large-scale production activities, considering its infrastructure suitable for industrial operations.</p> <p><b>Answer:</b> The region belongs to the &lt;industrial_land&gt; category, indicating its infrastructure supports manufacturing and production functions.</p>		
<b>Context Reasoning:</b>	<p><b>Question:</b> Please segment the area designed for the practice and play of a sport involving rackets, located near the area designated for basketball activities and adjacent to the open green space.</p> <p><b>Answer:</b> The region belongs to the &lt;tennis_court&gt; category, typically identified by its marked rectangular layout and its placement adjacent to basketball facilities and surrounding open green areas.</p>		
<b>Referring:</b>	<p><b>Question:</b> Segment the leafless tree located near the sidewalk adjacent to the building.</p>		
<b>Referring:</b>	<p><b>Question:</b> Segment the barren land area located centrally in the image, characterized by its irregular shape and lack of vegetation.</p>		
<b>Interactive (Point):</b>	<p><b>Question:</b> Please segment the region/target corresponding to the points [0.01,0.035].</p>		
<b>Interactive (Box):</b>	<p><b>Question:</b> Please segment the region/target corresponding to the box <math>x_0, y_0=[0.514, 0.168], x_1, y_1=[0.738, 0.367]</math>.</p>		

Figure 11. Additional samples of GeoSeg-1M.





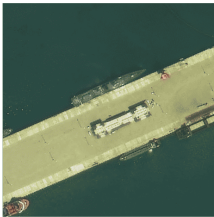


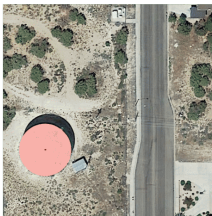
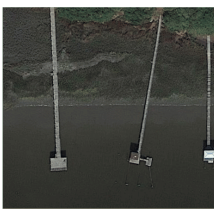
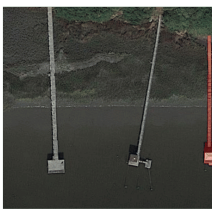



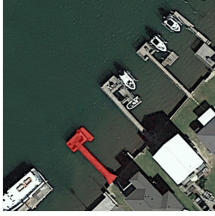
		Image	Ground Truth
<b>Attribute Reasoning:</b>	<p><b>Question:</b> Please segment the area primarily designed for facilitating vehicular and pedestrian movement across another pathway without interruption.</p> <p><b>Answer:</b> The region belongs to the &lt;Overpass&gt; category, serving the function of allowing traffic to flow over another route efficiently.</p>		
<b>Attribute Reasoning:</b>	<p><b>Question:</b> Please segment the area situated north of the railway line, adjacent to the large building complex, within the densely constructed zone, where liquid resources are likely stored or managed.</p> <p><b>Answer:</b> The region belongs to the &lt;water&gt; category, as it is located north of the railway line, sits beside the large building complex, and appears as a distinct liquid-holding area within the densely built environment.</p>		
<b>Context Reasoning:</b>	<p><b>Question:</b> Please segment the target object that is designed specifically for transporting bulk liquid materials across long distances, ensuring safe containment and specialized handling during transit.</p> <p><b>Answer:</b> The region belongs to the &lt;Liquid-Cargo-Ship&gt; category, designed for transporting liquid cargo efficiently and safely.</p>		
<b>Referring:</b>	<p><b>Question:</b> Segment the large circular white storage tank located near the road on the left side of the image.</p>		
<b>Referring:</b>	<p><b>Question:</b> Segment the rectangular structure with a white roof located on the right side of the image.</p>		
<b>Interactive (Point):</b>	<p><b>Question:</b> Please segment the region/target corresponding to the points [0.85,0.068].</p>		
<b>Interactive (Box):</b>	<p><b>Question:</b> Please segment the region/target corresponding to the box <math>x_0, y_0=[0.295, 0.607]</math>, <math>x_1, y_1=[0.529, 0.873]</math>.</p>		

Figure 12. Additional samples of GeoSeg-Bench.








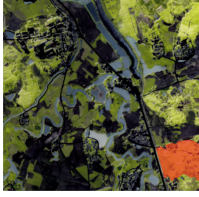
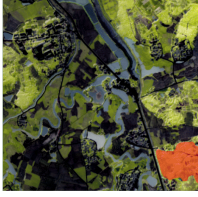












		Image	Ground Truth	Prediction of UniGeoSeg
<b>Attribute Reasoning:</b>	<b>Question:</b> Please segment the area designed for water retention and flood control in the landscape, considering its structural integrity and essential role in regional water management.			
<b>Attribute Reasoning:</b>	<b>Question:</b> Please segment the region designed for facilitating transportation and connectivity among various areas, essential for vehicular movement and access.			
<b>Context Reasoning:</b>	<b>Question:</b> Please segment the area located at the southern edge of the cluster where trees are prominent, situated between a river and open fields, and surrounded by lush vegetation.			
<b>Referring:</b>	<b>Question:</b> Segment the winding road that runs horizontally across the middle of the image, passing through the rural area.			
<b>Referring:</b>	<b>Question:</b> Segment the roundabout located near the intersection of the main road and the smaller road on the right side of the image.			
<b>Interactive (Point):</b>	<b>Question:</b> Please segment the region/target corresponding to the points [0.48,0.377].			
<b>Interactive (Box):</b>	<b>Question:</b> Please segment the region/target corresponding to the box $x_0,y_0=[0.328,0.705]$ , $x_1,y_1=[0.998,0.998]$ .			

Figure 13. Additional samples of model prediction by UniGeoSeg.

## References

- [1] MR Bayanlou and M Khoshboresh-Masouleh. Multi-task learning from fixed-wing uav images for 2d/3d city modelling. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 44:1–5, 2021. 1, 8
- [2] Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024. 10
- [3] Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*, 2024. 11
- [4] Javiera Castillo-Navarro, Bertrand Le Saux, Alexandre Boulch, Nicolas Audebert, and Sébastien Lefèvre. Semi-supervised semantic segmentation in earth observation: The minifrance suite, dataset analysis and multi-task network study. *Machine Learning*, 111(9):3125–3160, 2022. 1, 8
- [5] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*, pages 24185–24198, 2024. 11
- [6] Wei-Lin Chiang, Zhuohan Li, Ziqing Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6, 2023. 11
- [7] Ilke Demir, Krzysztof Koperski, David Lindenbaum, Guan Pang, Jing Huang, Saikat Basu, Forest Hughes, Devis Tuia, and Ramesh Raskar. Deepglobe 2018: A challenge to parse the earth through satellite images. In *CVPR workshops*, pages 172–181, 2018. 1, 8
- [8] Anatol Garioud, Nicolas Gonthier, Loic Landrieu, Apolline De Wit, Marion Valette, Marc Poupée, Sébastien Giordano, et al. Flair: a country-scale land cover semantic segmentation dataset from multi-source optical imagery. *NeurIPS*, 36: 16456–16482, 2023. 1, 8
- [9] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In *CVPR*, pages 9579–9589, 2024. 11
- [10] Ke Li, Gang Wan, Gong Cheng, Liqiu Meng, and Junwei Han. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS journal of photogrammetry and remote sensing*, 159:296–307, 2020. 2, 8
- [11] Kaiyu Li, Zepeng Xin, Li Pang, Chao Pang, Yupeng Deng, Jing Yao, Guisong Xia, Deyu Meng, Zhi Wang, and Xiangyong Cao. Seearth-r1: Geospatial pixel reasoning via large language model. *arXiv preprint arXiv:2504.09644*, 2025. 2, 8, 10, 11
- [12] Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. Textbooks are all you need ii: phi-1.5 technical report. *arXiv preprint arXiv:2309.05463*, 2023. 10, 11
- [13] Fan Liu, DeLong Chen, Zhangqingyun Guan, Xiaocong Zhou, Jiale Zhu, Qiaolin Ye, Liyong Fu, and Jun Zhou. Remoteclip: A vision language foundation model for remote sensing. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–16, 2024. 11
- [14] Sihan Liu, Yiwei Ma, Xiaoqing Zhang, Haowei Wang, Jiayi Ji, Xiaoshuai Sun, and Rongrong Ji. Rotated multi-scale interaction network for referring remote sensing image segmentation. In *CVPR*, pages 26658–26668, 2024. 2, 8, 10, 11
- [15] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer:

- Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 10, 11
- [16] Ruizhe Ou, Yuan Hu, Fan Zhang, Jiabin Chen, and Yu Liu. Geopix: A multimodal large language model for pixel-level image understanding in remote sensing. *IEEE Geoscience and Remote Sensing Magazine*, 2025. 11
- [17] Jerome Quenum, Wen-Han Hsieh, Tsung-Han Wu, Ritwik Gupta, Trevor Darrell, and David M Chan. Lisat: Language-instructed segmentation assistant for satellite imagery. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*. 11
- [18] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PmlR, 2021. 11
- [19] Zhongwei Ren, Zhicheng Huang, Yunchao Wei, Yao Zhao, Dongmei Fu, Jiashi Feng, and Xiaojie Jin. Pixellm: Pixel reasoning with large multimodal model. In *CVPR*, pages 26374–26383, 2024. 11
- [20] Caleb Robinson, Le Hou, Kolya Malkin, Rachel Soobitsky, Jacob Czawlytko, Bistra Dilkina, and Nebojsa Jojic. Large scale high-resolution land cover mapping with multi-resolution data. In *CVPR*, pages 12726–12735, 2019. 1, 8
- [21] Akashah Shabbir, Mohammed Zumri, Mohammed Benamoun, Fahad Shahbaz Khan, and Salman Khan. Geopixel: Pixel grounding large multimodal model in remote sensing. In *ICML*, 2025. 11
- [22] Qian Shi, Da He, Zhengyu Liu, Xiaoping Liu, and Jingqian Xue. Globe230k: A benchmark dense-pixel annotation dataset for global land cover mapping. *Journal of Remote Sensing*, 3:0078, 2023. 1, 8
- [23] Yan Shu, Bin Ren, Zhitong Xiong, Danda Pani Paudel, Luc Van Gool, Begum Demir, Nicu Sebe, and Paolo Rota. Earthmind: Towards multi-granular and multi-sensor earth observation with large multimodal models. *arXiv preprint arXiv:2506.01667*, 2025. 11
- [24] Xian Sun, Peijin Wang, Zhiyuan Yan, Feng Xu, Ruiping Wang, Wenhui Diao, Jin Chen, Jihao Li, Yingchao Feng, Tao Xu, et al. Fair1m: A benchmark dataset for fine-grained object recognition in high-resolution remote sensing imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 184:116–130, 2022. 2, 8
- [25] Xin-Yi Tong, Gui-Song Xia, and Xiao Xiang Zhu. Enabling country-scale land cover mapping with meter-resolution satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 196:178–196, 2023. 1, 8
- [26] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 11
- [27] Di Wang, Jing Zhang, Bo Du, Minqiang Xu, Lin Liu, Dacheng Tao, and Liangpei Zhang. Samrs: Scaling-up remote sensing segmentation dataset with segment anything model. *NeurIPS*, 36:8815–8827, 2023. 2, 9
- [28] Junjue Wang, Zhuo Zheng, Xiaoyan Lu, Yanfei Zhong, et al. Loveda: A remote sensing land-cover dataset for domain adaptive semantic segmentation. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. 1, 8
- [29] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *CVPR*, pages 3974–3983, 2018. 2, 8
- [30] Qikai Lu Huangfeng Shen Shengyang Li Shucheng You Liangpei Zhang Xin-Yi Tong, Gui-Song Xia. Land-cover classification with high-resolution remote sensing images using transferable deep models. *Remote Sensing of Environment*, doi: 10.1016/j.rse.2019.111322, 2020. 1, 8
- [31] Liang Yao, Fan Liu, DeLong Chen, Chuanyi Zhang, Yijun Wang, Ziyun Chen, Wei Xu, Shimin Di, and Yuhui Zheng. Remotesam: Towards segment anything for earth observation. In *ACM MM*, pages 3027–3036, 2025. 2, 8, 11
- [32] Yang Zhan, Zhitong Xiong, and Yuan Yuan. Rsvg: Exploring data and models for visual grounding on remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 61: 1–13, 2023. 9
- [33] Zheng Zhang, Yeyao Ma, Enming Zhang, and Xiang Bai. Psalm: Pixelwise segmentation with large multi-modal model. In *ECCV*, pages 74–91. Springer, 2024. 11