

Towards Open-Vocabulary Remote Sensing Image Semantic Segmentation

Chengyang Ye, Yunzhi Zhuge, Pingping Zhang*

School of Future Technology, School of Artificial Intelligence, Dalian University of Technology
yecy@mail.dlut.edu.cn, {zgyz,zhpp}@dlut.edu.cn

Abstract

Recently, deep learning based methods have revolutionized remote sensing image segmentation. However, these methods usually rely on a predefined semantic class set, thus needing additional image annotation and model training when adapting to new classes. More importantly, they are unable to segment arbitrary semantic classes. In this work, we introduce Open-Vocabulary Remote Sensing Image Semantic Segmentation (OVRSISS), which aims to segment arbitrary semantic classes in remote sensing images. To address the lack of OVRSISS datasets, we develop LandDiscover50K, a comprehensive dataset of 51,846 images covering 40 diverse semantic classes. In addition, we propose a novel framework named GSNet that integrates domain priors from special remote sensing models and versatile capabilities of general vision-language models. Technically, GSNet consists of a Dual-Stream Image Encoder (DSIE), a Query-Guided Feature Fusion (QGFF), and a Residual Information Preservation Decoder (RIPD). DSIE first captures comprehensive features from both special models and general models in dual streams. Then, with the guidance of variable vocabularies, QGFF integrates specialist and generalist features, enabling them to complement each other. Finally, RIPD is proposed to aggregate multi-source features for more accurate mask predictions. Experiments show that our method outperforms other methods by a large margin, and our proposed LandDiscover50K improves the performance of OVRSISS methods. The dataset and method will be publicly available.

Code & Datasets — <https://github.com/yecy749/GSNet>

Introduction

Remote sensing image analysis aims at processing and interpreting remote sensing images to provide insights into natural environments and human activities. It serves as a pivotal tool in improving human welfare. As a central technique, Remote Sensing Image Semantic Segmentation (RSISS) facilitates various real-world applications, including enhancing agricultural yields (FAO 2018), mitigating natural disasters (Van Westen 2000), and managing land cover changes (Yan, Wan, and Zhang 2022; Yan et al. 2023).

*Corresponding author (zhpp@dlut.edu.cn).

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

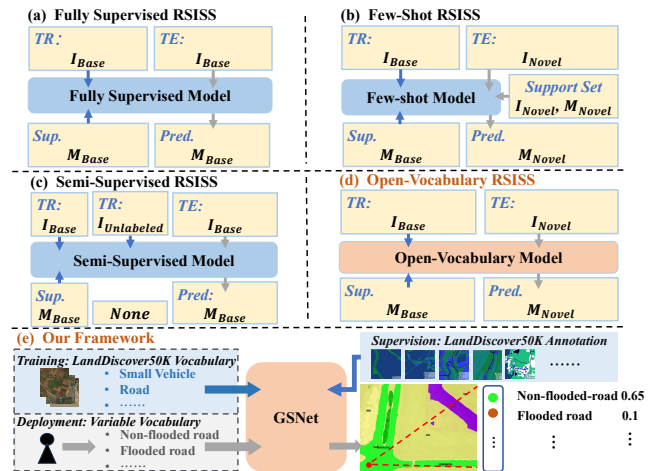


Figure 1: Comparison of different learning paradigms for RSISS. (a) Fully supervised methods train and test on the same dataset. (b) Few-shot methods train on large annotated base classes and test on novel classes using a small support set. (c) Semi-supervised methods use large-scale unlabeled data with small-scale labeled base classes for training, then test on base classes. (d) Open-Vocabulary methods train on large-scale labeled data and test on arbitrary semantic classes. (e) Our framework is illustrated in brief.

Recently, deep learning has revolutionized RSISS by enabling automatic segmentation methods. With the advantages of different networks, researchers have leveraged Fully Convolutional Networks (FCN) (Long, Shelhamer, and Darrell 2015), UNet (Ronneberger, Fischer, and Brox 2015), and Vision Transformers (ViT) (Dosovitskiy et al. 2020) to improve the RSISS performance. Furthermore, there has been a growing interest in exploring partially supervised RSISS methods, such as few-shot methods (Jiang, Zhou, and Li 2022), weakly supervised methods (Hua et al. 2021), and semi-supervised methods (Li et al. 2021). However, existing RSIS methods cannot segment arbitrary semantic classes, since they train and test on a predefined set of classes.

With the advancement of RSI technologies, the excessive burden for data annotation highlights the need for a generalist model capable of adapting to diverse natural environ-

ments. Different from natural images, RSI often incurs significantly higher annotation costs due to its high resolution and inherent semantic ambiguity (Yao et al. 2016). However, state-of-the-art RSSIS methods are typically trained and tested on a predefined set of classes using annotated data, as shown in Fig. 1 (a), (b), and (c). Considering the need for generalization, the above defects become even more significant. In fact, current RSISS approaches are not only unable to segment classes outside the predefined set but also struggle to generalize across different domains.

Furthermore, the importance of a generalist model in RSISS is amplified in scenarios requiring rapid responses, such as natural disasters, since there is insufficient time for extensive data annotation and model training. More specifically, due to the temporal and spatial changes in remote sensing data, great domain gaps exist in RSI. With such domain gaps, if a model can only perform accurate RSISS under specific conditions, its utility becomes severely restricted. Therefore, we propose Open-Vocabulary Remote Sensing Image Semantic Segmentation (OVRSSIS) to address these challenges. OVRSSIS aims at segmenting arbitrary semantic classes in RSI. Without the limitation of the pre-defined class set, OVRSSIS allows users to flexibly switch between desired class sets based on their needs. Thus, OVRSSIS not only reduces costs but also enables quicker responses in critical situations. We show the comparison between OVRSSIS and existing RSISS learning paradigms in Fig. 1.

To the best of our knowledge, this is the first study to address the challenges of OVRSSIS. Recognizing the absence of datasets specifically designed for OVRSSIS, we developed LandDiscover50K, a dataset comprising 51,846 remote sensing images spanning 40 distinct classes. Alongside this dataset, we have formulated a comprehensive benchmark to facilitate the robust evaluation of OVRSSIS methodologies.

For OVRSSIS methods, there are two kinds of intuitive methods. First, one can simply train existing Open-Vocabulary Natural Image Semantic Segmentation (OVNISS) methods on LandDiscover50K. However, this kind of methods face notable performance limitations, primarily due to the absence of tailored designs for the RSI domain. Second, one can enhance existing OVNISS methods for RSI domain by replacing the generic CLIP (Radford et al. 2021) with domain-specific RemoteCLIP (Liu et al. 2024a). However, performance degradation is observed, primarily due to the limited generalization of RemoteCLIP. Both failures highlight the challenge of balancing domain-specific knowledge with generalization. Thus, how to effectively integrate RSI domain priors while maintaining a strong generalization ability remains an unsolved problem.

To address the above challenges, we propose GSNet, a novel framework tailored to effectively integrate RSI specialist domain priors with generalist CLIP. It employs a Dual-Stream Image Encoder (DSIE) to simultaneously extract generalist features from CLIP in parallel with RSI domain-specific features from a RSI backbone. A Query-Guided Feature Fusion (QGFF) is further introduced to integrate special RSI features and general features, enabling them to complement each other with the guidance of variable vocabularies. We also design a Residual Information

Preservation Decoder (RIPD) to aggregate multi-source features for more accurate mask predictions.

To summarize, our contributions are as followed:

- We put forward Open-Vocabulary Remote Sensing Image Semantic Segmentation along with a tailored dataset named LandDiscover50K.
- We propose a novel framework GSNet for OVRSSIS which first extracts both generalist and specialist features with DSIE, followed by QGFF for multi-source feature fusion, and finally employs RIPD for information preservation and detail refinement.
- We conduct extensive experiments to demonstrate that our GSNet outperforms other state-of-the-art OVNISS methods by a large margin, and our LandDiscover50K markedly boosts the performance of OVRSSIS methods.

Related Work

Remote Sensing Image Semantic Segmentation

Existing RSISS methods mostly focus on close-set performance evaluated on certain benchmarks. Datasets such as LoveDA (Wang et al. 2021), iSAID (Waqas Zamir et al. 2019) are proposed, enabling robust model training and evaluation. Based on these datasets, works such as FarSeg (Zheng et al. 2020) and AerialFormer (Yamazaki et al. 2023) have successfully extended UNet and ViT to RSISS, boosting the performance on those datasets.

More recently, some work has addressed the data scarcity in RSI domain. For instance, Li et al. (Li et al. 2021) delve into semi-supervised RSISS with consistency self-training. Jiang et al. (Jiang, Zhou, and Li 2022) use prototype-based semantic matching and a non-parametric metric learning loss to address few-shot RSISS. Hua et al. (Hua et al. 2021) propose a feature and spatial relational regularization to boost the performance of weakly-supervised RSISS. Zhu et al. (Zhu et al. 2023) delve into universal domain adaptation in RSISS, addressing domain distribution discrepancies through adversarial learning. However, none of them can segment arbitrary classes.

As for RSISS datasets, we refer the readers to Tab. 1 for more details. To our knowledge, there exists no prior dataset tailored for OVRSSIS. Except for the recently proposed FLAIR (Garioud et al. 2024) and SAMRS, most existing datasets are too small for OVRSSIS. As for FLAIR, it only annotates regular land cover types, overlooking small objects, making it hard to smoothly adapt to OVRSSIS settings. SAMRS (Wang et al. 2024) has successfully adapted several large-scale bounding-box annotated RSI object detection datasets to pixel-wisely annotated RSISS datasets. However, it has a limited generalization ability to land-cover segmentation tasks, which is widely used in multiple applications of RSI analysis.

Open-Vocabulary Natural Image Semantic Segmentation

OVNISS aims at segmenting arbitrary semantic class without needing explicit training examples for every possible

Datasets	Year	Images	Category	GSD (m)	Image Width	Total Area (km ²)
ISPRS Potsdam* (ISPRS 2013a)	2013	38	6	0.05	6,000	3.42
ISPRS Vaihingen (ISPRS 2013b)	2013	33	6	0.09	2,494	1.4
Zuric Summer (Volpi and Ferrari 2015)	2015	20	8	0.62	1,150	9.37
UAVid (Lyu et al. 2020)	2018	420	8	-	3,840-4,096	-
Deep Globe Land Cover [†] (Demir et al. 2018)	2018	1,146	7	0.5	2,448	1,716.9
iSAID (Waqas Zamir et al. 2019)	2019	2,806	15	-	800-13,000	-
FloodNet* (Rahnemoonfar et al. 2021)	2020	3,200	9	0.02	3,000	-
GID (Tong et al. 2018)	2020	150	5	1-4	7,200	75,900
Landcover.ai (Boguszewski et al. 2021)	2020	41	3	0.25-0.5	800-13,000	216.27
LoveDA [†] (Wang et al. 2021)	2021	5,987	7	0.3	1,024	536.15
FLAIR* (Garioud et al. 2024)	2022	77,762	19	0.1-0.2	512	817
Open Earth Map [†] (Xia et al. 2023)	2022	5,000	8	0.25-0.5	1,024	799
SIOR [†] (Wang et al. 2024)	2023	23,463	20	0.5-30	800	-
SOTA [†] (Wang et al. 2024)	2023	17,480	18	-	1,024	-
FAST* (Wang et al. 2024)	2023	64,147	37	0.3-0.8	600	-
LandDiscover50K	2024	51,846	40	-	400-1,200	-

Table 1: Comparison of LandDiscover50K with other RSISS Datasets. Datasets marked with [†] contributed to the compilation of LandDiscover50K, while those marked with * were selected for evaluating OVRSSIS.

category. Recently, OVNISS has seen a significant performance boost thanks to the emergence of large-scale pre-trained Vision-Language Models (VLMs) such as CLIP. There are mainly two lines of work in OVNISS, i.e., single-stage methods and two-stage methods. For single-stage methods, LSeg (Li et al. 2022) leverages CLIP’s patch embeddings, in line with CLIP’s text embeddings to build the correlation of patches and texts. SAN (Xu et al. 2023) introduces an additional backbone alongside a frozen CLIP, to directly generate and classify region proposals. CAT-SEG (Cho et al. 2024) computes a patch-level cost volume map between image and text embeddings, and refines the cost volume map to obtain the final prediction. Furthermore, SED (Xie et al. 2024) enhances the performance by introducing a hierarchical encoder-decoder framework. For two-stage methods, OpenSeg (Ghiasi et al. 2022) decouples the OVNISS task into class-agnostic proposal generation and proposal classification. OVSeg (Liang et al. 2023) refines CLIP with mask-adapted image-text pairs to better classify mask proposals. SCAN (Liu et al. 2024b) incorporates CLIP’s generalized semantic prior into proposal embeddings to prevent collapsing on known categories and applies a contextual shift strategy to address global context gaps.

LandDiscover50K Dataset

To address the lack of generalizable datasets on OVRSSIS, we present LandDiscover50K. This dataset is designed to overcome several limitations in existing RSISS datasets. LandDiscover50K includes 51,846 high-resolution remote sensing images annotated across 40 object classes. Meanwhile, LandDiscover50K addresses domain shifts in RSI semantic segmentation by incorporating diverse sensors, resolutions, classes, and ground sample distances. Moreover, LandDiscover50K enhances model robustness and generalizability through the integration of fine-grained object detection datasets and large-scale land cover datasets.

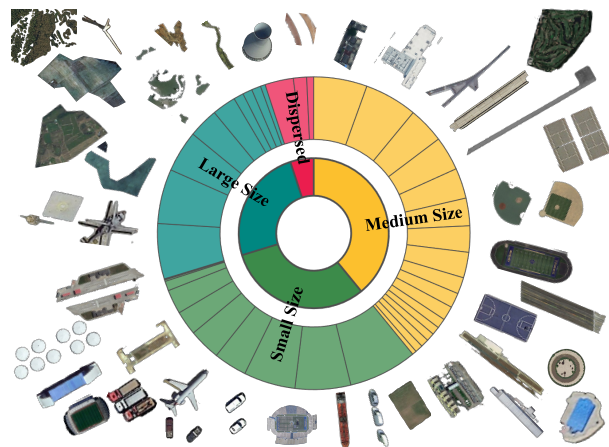


Figure 2: Illustration of the semantic class distribution and visual samples from LandDiscover50K.

Data Acquisition and Annotation

To build the LandDiscover50K dataset, we collect images from established RSISS datasets that provide pixel-wise annotations. These datasets include Open Earth Map (OEM) (Xia et al. 2023), LoveDA (Wang et al. 2021), Deep Globe Land Cover (Demir et al. 2018), SIOR (Wang et al. 2024), and SOTA (Wang et al. 2024), with a detailed overview provided in Tab. 1. LandDiscover50K comprises 51,846 image-annotation pairs across 40 diverse classes, drawn from sources such as DOTA (Xia et al. 2018), DIOR (Li et al. 2020), xBD (Gupta et al. 2019), Inria (Maggiore et al. 2017), OpenCities AI (Open Cities AI Challenge), SpaceNet (Van Etten, Lindenbaum, and Bacastow 2018), LandCover.ai (Boguszewski et al. 2021), AIRS (Chen et al. 2019), GeoNRW (Baier et al. 2021) and HTCD (Shao et al. 2021) etc. To ensure the comprehensive evaluation,

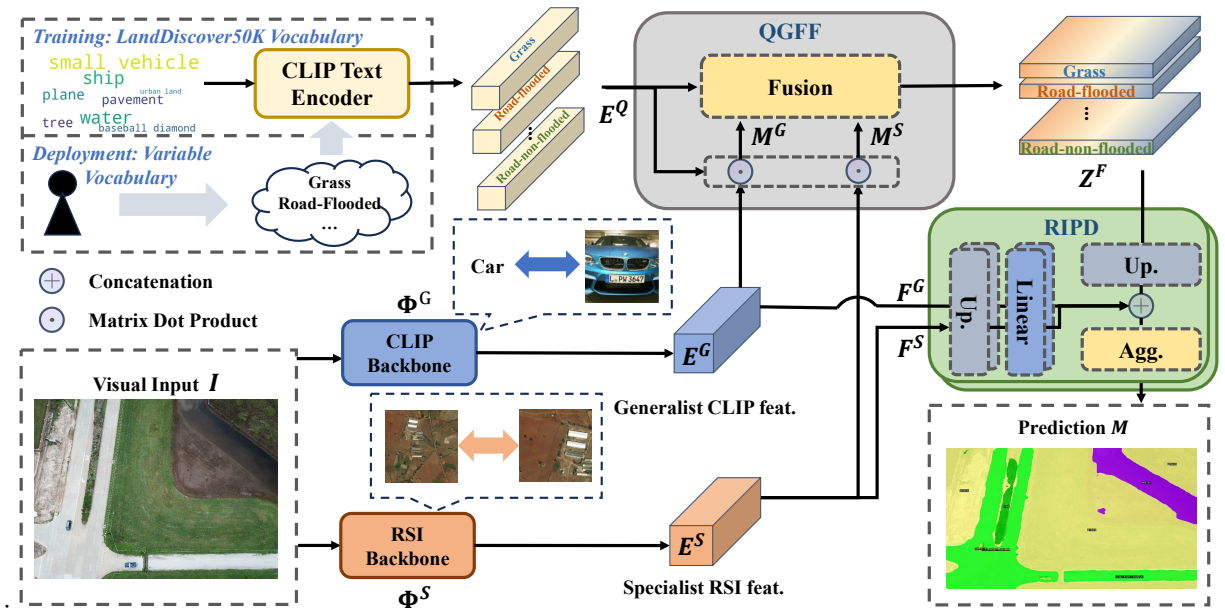


Figure 3: The overall architecture of GSNet. DSIE consists of a generalist CLIP backbone and a specialist RSI backbone. The specialist RSI backbone is pre-trained on RSI using self-supervised learning paradigm, while CLIP is pre-trained on image-text datasets using contrastive learning paradigm. QGFF enables dual stream features to complement each other under the guidance of variable vocabularies. RIPD further aggregates the multi-source features for more accurate mask predictions.

fair comparison, and standardized preprocessing, we chose to utilize only the RGB modality, which is broadly accessible and aligns with the needs of open-vocabulary tasks.

For annotations, our approach emphasizes diversity and scale. We merge identical classes and preserve unique fine-grained ones. In addition, we consolidate generic “background” labels into a single “unlabeled” category to mitigate over-fitting risks associated with the semantic bias introduced by the varying purposes of the source datasets.

Statistics and Analysis

As shown in Fig. 2, the LandDiscover50K dataset covers a diverse range of semantic classes from expansive land covers to salient objects. This diversity is crucial for modeling general variations inherent in real-world remote sensing tasks. In addition, LandDiscover50K provides a balanced spatial coverage of segments within images, facilitating robust model training and reducing positional biases. Visualization of the attributes of LandDiscover50K are detailed in the Supplementary Materials.

Our Method

In this section, we introduce the Generalist and Specialist Network (GSNet) for OVRSISS. The overall architecture of GSNet is depicted in Fig. 3. It employs a DSIE, integrating a domain-specific RSI specialist image encoder and a CLIP-based image-text aligned image encoder. DSIE produce two kinds of complementary feature maps that are then harmonized using a Query-Guided Feature Fusion to leverage both generic and domain-specific strengths. The fused features are further denoised and upsampled through the

RIPD, which ensures the preservation of critical information from both the specialist and generalist streams to generate the final segmentation results. We will delve into the detailed designs of GSNet in the following sections.

Dual-Stream Image Encoder

Directly applying CLIP to the RSI domain or using a RSI-specific model like RemoteCLIP has been found to underperform on OVRSISS tasks. This is mainly because RemoteCLIP is pre-trained on a much smaller dataset compared to the general CLIP model. In addition, the general CLIP lacks the specialized knowledge for OVRSISS. Thus, we propose a Dual-Stream Image Encoder that synergistically extracts specialist features and generalist features.

Generalist CLIP Backbone. Following the most common practices, We use the ViT/B-16 architecture of the pre-trained CLIP. More specifically, we employ the image encoder of CLIP excluding the final projection heads, denoted as Φ^G . Given an image $I \in \mathbb{R}^{H \times W \times 3}$, we extract the final feature $E^G \in \mathbb{R}^{H' \times W' \times D}$, where D denotes the feature dimension and $\{H', W'\} = \{H/16, W/16\}$. With the query set C_N , we extract query embeddings $E^Q \in \mathbb{R}^{N \times D}$ with CLIP text encoder Φ^Q , as depicted in Fig. 3.

Specialist RSI Backbone. Although CLIP excels in recognizing novel objects, its effectiveness diminishes in dense segmentation tasks which are typical in complex remote sensing images. This is primarily due to its lack of RSI domain priors. Thus, we propose a specialist RSI backbone to enhance the RSI domain priors complementing the general CLIP model. Specifically, it incorporates self-

supervised pre-trained DINO (Caron et al. 2021), which conducts contrastive learning utilizing local and global image views to capture spatial hierarchies effectively. It effectively introduces RSI domain priors, while mitigating potential over-fitting that can be harmful to the generalization ability of CLIP. Moreover, the specialist RSI backbone is trained solely on the image set of LandDiscover50K without labels, minimizing its latent impact on the generalization ability of GSNet. Note that, DINO adopts the same ViT architecture as CLIP, ensuring a seamless integration. With the specialist RSI backbone Φ^S , we can extract the specific feature $E^S \in \mathbb{R}^{H' \times W' \times D}$.

Query-Guided Feature Fusion

One of the biggest challenges in OVRSISS is establishing robust associations between texts and images. To address this issue, CAT-SEG (Cho et al. 2024) simply utilizes the cosine similarity of CLIP’s text and image embeddings. Differently, our approach does not rely solely on CLIP’s image embeddings. Instead, we decouple the scene features into specialist features and generalist features. Then, we integrate text embeddings with both generalist CLIP’s embedding and specialist RSI embedding. Hence, our approach offers advantages over CAT-SEG in RSI domain, which will be further evaluated in the experiment section. In addition, since the RSIB remains fixed during training, we do not introduce substantial computational overhead involved in training and testing. Specifically, after obtaining E^G and E^S , we normalize them and respectively compute matrix dot products with the embedded text-based query features, then we obtain the cost volume map of dual streams: M^G and M^S , as follows:

$$M_{H',W',N}^{\{G,S\}} = \sum_D \text{norm}(E_{H',W',D}^{\{G,S\}} \cdot \text{norm}((E_{N,D}^Q)^\top)), \quad (1)$$

where $\text{norm}(X)$ denotes the L_2 normalization in channel dimension. This ensures that each query benefits from dual-stream image embeddings. Afterwards, we further embed the volume map to a latent space:

$$Z_{H' \times W' \times N \times D}^{\{G,S\}} = \sigma(\varphi_{7 \times 7}(M_{H' \times W' \times N}^{\{G,S\}})), \quad (2)$$

where $\varphi_{7 \times 7}$ denotes the convolution operation with a kernel size of 7×7 , and σ denotes the sigmoid activation function. The processed feature maps then undergo concatenation and aggregation, producing a query-guided fused feature map. To preserve the generalist integrity of CLIP and prevent feature degradation, a residual connection is applied to the fused feature map. The process is formed as follows:

$$Z^F = \sigma(\varphi_{7 \times 7}(Z^G \oplus Z^S)) + Z^G, \quad (3)$$

where \oplus denotes the concatenation operation in channel dimension, $+$ denotes the element-wise addition. Note that, in the above procedure, we use a simple structure design to validate our ideas. More complex structures can be used to further improve the performances.

Residual Information Preservation Decoder

Since OVRSISS is a pixel-wise labeling task, prediction noises are inevitably introduced when using a versatile backbone. Firstly, we utilize the ViT-B/16 as the image encoder,

which downsamples images by a factor of 16, unavoidably resulting in some loss of details. Secondly, the matrix dot product between image embeddings and text embeddings compresses information within the hidden space, further introducing semantic ambiguities. To address the above challenges, we introduce RIPD for the backbone regularization and detail refinement, reducing prediction noises. RIPD aggregates multi-source features for more accurate mask predictions. Specifically, after obtaining the query-guided fused feature maps Z^F from QGFF, we first adopt deconvolution to obtain \tilde{Z}_0^F . Then, the upsampled feature map \tilde{Z}_0^F is concatenated with the projected intermediate feature maps $\tilde{F}_n^{\{G,S\}}$ from the CLIP backbone and RSI backbone. Here, n is the layer number. This is followed by a linear projection for dimension reduction. Then, we concatenate Z^F , \tilde{F}_n^G , and \tilde{F}_n^S in channel dimension. Finally, we aggregate the concatenated feature to obtain \tilde{Z}_1^F . The whole RIPD is build by stacking the above block as follows:

$$\tilde{F}_{n_i}^{\{G,S\}} = \text{Linear}(\mathcal{D}(F_{n_i}^{\{G,S\}})), \quad (4)$$

$$\tilde{Z}_{i+1} = \text{AGR}(\text{AGR}(\varphi_{3 \times 3}(\mathcal{D}(\tilde{Z}_i) \oplus \tilde{F}_{n_i}^G \oplus \tilde{F}_{n_i}^S))), \quad (5)$$

$$\text{AGR}(X) = \text{ReLU}(\text{GN}(\varphi_{3 \times 3}(X))), \quad (6)$$

where \mathcal{D} is the deconvolution with stride 2. GN is the group normalization. In practice, we stack two blocks for 4x up-sampling, followed by a single convolution to obtain the final prediction M .

Experiments

Datasets and Metrics

Datasets. Since there is no prior work, we select four representative datasets to assess the performance of OVRSISS: FLAIR (Garioud et al. 2024), FAST (Wang et al. 2024), ISPRS Potsdam (ISPRS 2013a) and FloodNet (Rahmehoonfar et al. 2021). Detailed information about the testing datasets is presented in Tab. 1 marked by *. Importantly, we do not restrict our evaluation to novel classes as a prerequisite for OVRSISS. We argue that cross-dataset validation is the most effective way to simulate real-world scenarios, as it reflects the complexity of real-world challenges while efficiently utilizing existing RSIS datasets. Among them, ISPRS Potsdam emphasizes in-vocabulary performance in OVRSISS with its high category similarity to the training set. FloodNet focuses on post-flood analysis, FLAIR covers large-scale land cover types, and FAST specializes in fine-grained object segmentation. These diverse datasets enable a comprehensive evaluation of the OVRSISS methods.

Metrics. Following previous works (Xu et al. 2022; Liang et al. 2023), we employ the mean Intersection-over-Union (mIoU) as the segmentation metric for all experiments. As for computation, we adopt GFLOPs and parameters. More details can be found in the Supplementary Materials.

Implementation Details

We train the GSNet with a per-pixel binary cross-entropy loss. Our implementation is based on PyTorch (Paszke et al. 2019) and Detectron2 (Wu et al. 2019). We adopt

Method	VLM	FLAIR	FAST	Potsdam	FloodNet	Avg.
<i>Fine-tuned on LandDiscover50K with the weight pre-trained on COCO</i>						
CAT-SEG (Cho et al. 2024)	CLIP-ViT-B/16	18.94	15.23	36.01	<u>40.64</u>	<u>27.71</u>
SCAN (Xu et al. 2023)	CLIP-ViT-B/16	19.19	10.02	5.68	36.61	17.87
OV-SEG (Liang et al. 2023)	CLIP-ViT-B/16	10.32	9.10	4.93	27.78	13.03
SAN (Xu et al. 2023)	CLIP-ViT-B/16	19.91	15.30	19.91	35.75	22.71
SED (Xie et al. 2024)	CLIP-ConvNeXt-B	18.59	12.90	28.16	31.65	22.83
EBSeg (Shan et al. 2024)	CLIP-ViT-B/16	<u>21.26</u>	18.53	5.68	35.26	20.18
<i>Trained on LandDiscover50K</i>						
CAT-SEG (Cho et al. 2024)	CLIP-ViT-B/16	19.99	13.90	<u>38.79</u>	37.89	27.64
CAT-SEG (Cho et al. 2024)	RemoteCLIP-ViT-B/32	14.36	10.79	24.54	29.83	19.88
SCAN (Liu et al. 2024b)	CLIP-ViT-B/16	18.49	8.56	5.60	39.23	17.97
SCAN (Liu et al. 2024b)	RemoteCLIP-ViT-B/32	15.38	9.48	10.15	19.05	13.51
OV-SEG (Liang et al. 2023)	CLIP-ViT-B/16	13.65	7.47	5.02	15.17	10.33
SAN (Xu et al. 2023)	CLIP-ViT-B/16	22.48	<u>16.21</u>	10.01	30.28	19.74
SED (Xie et al. 2024)	CLIP-ConvNeXt-B	14.65	12.63	28.64	22.57	19.62
EBSeg (Shan et al. 2024)	CLIP-ViT-B/16	20.24	15.15	8.37	37.93	20.42
GSNet (Ours)	CLIP-ViT-B/16	20.00	16.61	45.75	42.63	31.25

Table 2: Performance comparison with other methods. The best and second-best results are in bold and underlined, respectively.

AdamW (Loshchilov and Hutter 2017) as the optimizer, setting the learning rate to 2×10^{-6} for the CLIP, while keeping the DINO fixed during the training process. The remaining parts of our model are randomly initialized and trained with a learning rate of 2×10^{-4} . The batch size is set to 4, and we use two NVIDIA RTX 3090 GPUs for training. We set the total number of training iterations to 30,000.

Comparison with Other Methods

Here, we compare our methods with recent state-of-the-art OVNISS methods (Xu et al. 2023; Liang et al. 2023; Cho et al. 2024; Xie et al. 2024). We re-evaluate all the methods to show their performance on OVRSISS. There are two settings: 1) fine-tuning on LandDiscover50K with the weight pre-trained on COCO-Stuff (Caesar, Uijlings, and Ferrari 2018) and 2) training exclusively on LandDiscover50K. For all methods, we deploy base models for a fair comparison.

Quantitative Evaluation. Tab. 2 presents the results of different methods on four datasets. It also shows the corresponding VLMs. For a fair comparison, we further improve the state-of-the-art methods SCAN and CAT-SEG by replacing CLIP with RemoteCLIP, which is pretrained on RSI domain. However, performance degradation occurs, further addressing the effectiveness of our proposed GSNet. Among these methods, our proposed method achieved the best averaged mIoU, outperforming the second best model by a large margin of 3.54% mIoU.

Qualitative Evaluation. Fig. 4 shows the qualitative results of our proposed method versus CAT-SEG. Our method clearly exhibits superior boundary awareness and enhanced semantic recognition of target objects in these challenging remote sensing images.

Ablation Studies

Ablation study on DSIE. The row A to D in Tab. 3 present the ablation results on DSIE. We respectively adopt

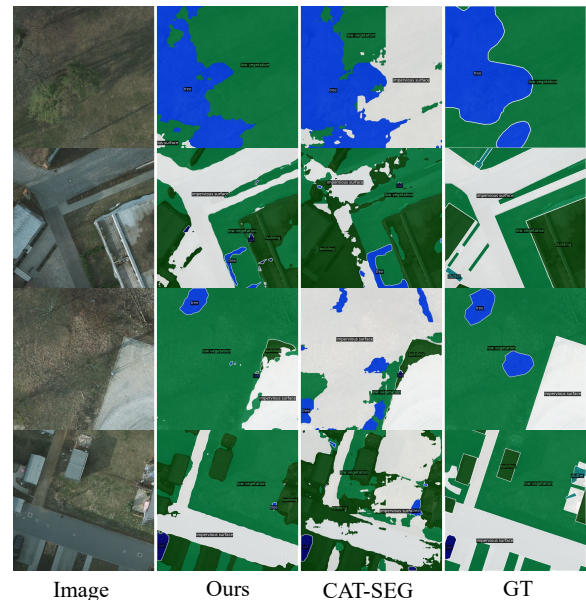


Figure 4: Qualitative evaluation of GSNet. Our method outperforms CAT-SEG in both semantic understanding and edge prediction.

CLIP, Remote-CLIP, and RSIB (Specialist RSI Backbone) alone. In addition, we use RSI-CLIP to complement RSIB. Our proposed DSIE has outperformed previous methods by 12%, 57%, 348%, and 59% measured by average mIoU on four datasets. Though pre-training on RSI image-caption pairs, RemoteCLIP’s smaller annotated dataset leads to a weaker generalization on OVRSISS. Our RSIB alone underperforms, as it has not been exposed to any annotated data. Our proposed GSNet complements CLIP with RSIB’s specialist priors, achieving a significant performance boost.

	CLIP	RSI-CLIP	RSIB	Cat	Fusion Q	QGFF	VLMs	RSIB	FLAIR	FAST	Potsdam	FloodNet	Avg.
A	✓	✗	✗	✗	✗	✗	✓	✗	18.62	14.29	38.59	40.10	27.90
B	✗	✓	✗	✗	✗	✗	✓	✗	14.36	10.79	24.54	29.83	19.88
C	✗	✗	✓	✗	✗	✗	✗	✓	3.87	2.63	14.65	6.74	6.98
D	✗	✓	✓	✗	✗	✓	✓	✓	13.25	11.30	24.89	28.99	19.61
E	✓	✗	✓	✓	✗	✗	✓	✓	14.47	10.18	18.57	32.37	18.90
F	✓	✗	✓	✗	✓	✗	✓	✓	8.84	10.12	24.32	35.92	19.81
G	✓	✗	✓	✗	✗	✓	✗	✗	17.70	13.57	38.39	35.91	26.39
H	✓	✗	✓	✗	✗	✓	✓	✗	17.59	15.06	35.93	41.31	27.47
I	✓	✗	✓	✗	✗	✓	✗	✓	17.06	15.30	36.50	33.09	25.49
Ours	✓	✗	✓	✗	✗	✓	✓	✓	20.00	16.61	45.75	42.63	31.25

Table 3: Ablation results on model components. Our GSNet is in the last row. Except for the ablated component, the rest of the architecture remains consistent with the full version. The red, yellow, blue parts respectively ablate the DSIE, QGFF, RIPD.

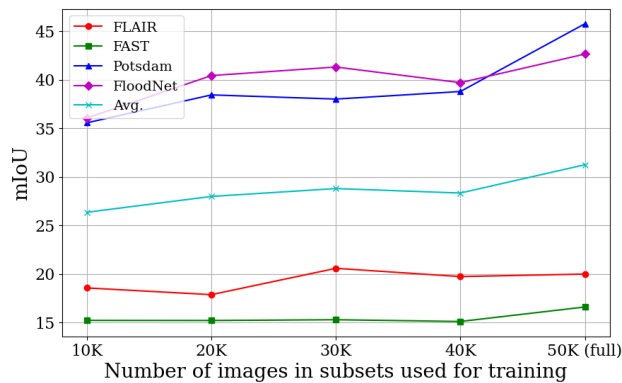


Figure 5: Performance of GSNet trained on different sizes of subsets of LandDiscover50K.

Ablation study on QGFF. The row E to F in Tab. 3 present ablation results on different designs of QGFF. The simple concatenation method, denoted as ‘Cat’, first expands the text embeddings to the scale of image embeddings. Then, it conducts feature concatenation on the two image embeddings with expanded text embeddings, followed by a linear projection to reduce the dimension. This simple concatenation method lags behind our full version by a large margin of 65%, showing the efficacy of QGFF. We further adopt an intuitive fusion and query design, denoted as ‘Fusion Q’, which fuses the feature from two image encoders and then conducts the matrix dot product with query embeddings. This design also lags behind our proposed QGFF by 58%.

Ablation study on RIPD. The row G to I in Tab. 3 present ablation results on various configurations of the RIPD. Compared with the non-integrated decoder, we achieved an improvement of 18%. When integrating only CLIP or RSIB, the model further achieves improvements of 14% and 23%, respectively. The results further confirm the capability of RIPD to synergistically leverage both generalist features from CLIP and RSI domain priors from RSIB for more accurate mask predictions.

Importance of LandDiscover50K. Our proposed LandDiscover50K paves the way of OVRSSS by introducing large-scale, multi-domain, multi-granularity RSI images

Training Set	FLAIR	FAST	Potsdam	FloodNet	Avg.
LoveDA	8.92	10.32	18.87	7.67	11.45
DeepGlobe	6.11	12.38	<u>21.26</u>	11.56	12.82
SIOR	13.26	13.93	13.80	14.47	13.87
SOTA	12.91	18.70	14.94	14.33	<u>15.22</u>
OEM	<u>13.85</u>	7.41	17.76	<u>20.09</u>	14.78
Ours	20.00	<u>16.61</u>	45.75	42.63	31.25

Table 4: Ablation results with different training datasets.

with various class annotations. Here, we evaluate the importance of LandDiscover50K. First, we randomly sample 10,000, 20,000, 30,000, and 40,000 image-mask pairs from LandDiscover50K to create subsets. Then, we train our proposed GSNet on the subsets and other existing RSIS datasets, respectively. The results are presented in Fig. 5. As the subset’s scale increase, the performance is significantly improved. Compared with the smallest subset, the model achieves over 18% performance improvement. We also train our model on typical RSIS datasets, as presented in Tab. 4. Though the model trained some of RSIS datasets can achieve outstanding performance on certain test datasets, they generally lag significantly behind our dataset. It further confirms the generalization and inclusiveness of LandDiscover50K.

Conclusion

In this work, we put forward the OVRSSS task and provide a large-scale dataset named LandDiscover50K for model training and testing. In addition, we develop GSNet, a novel framework that integrates RSI specialist domain priors with generalist pre-trained VLMs. Extensive experiments have evaluated the effectiveness of our methods and the dataset. We believe that this work will advance the current state of remote sensing research.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (No.62101092).

References

- Baier, G.; Deschamps, A.; Schmitt, M.; and Yokoya, N. 2021. Synthesizing optical and SAR imagery from land cover maps and auxiliary raster data. *TGRS*, 60: 1–12.
- Boguszewski, A.; Batorski, D.; Ziemba-Jankowska, N.; Dziedzic, T.; and Zambrzycka, A. 2021. LandCover.ai: Dataset for automatic mapping of buildings, woodlands, water and roads from aerial imagery. In *CVPR*, 1102–1110.
- Caesar, H.; Uijlings, J.; and Ferrari, V. 2018. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 1209–1218.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *ICCV*, 9650–9660.
- Chen, Q.; Wang, L.; Wu, Y.; Wu, G.; Guo, Z.; and Waslander, S. L. 2019. TEMPORARY REMOVAL: Aerial imagery for roof segmentation: A large-scale dataset towards automatic mapping of buildings. *ISPRS JPRS*, 147: 42–55.
- Cho, S.; Shin, H.; Hong, S.; Arnab, A.; Seo, P. H.; and Kim, S. 2024. Cat-seg: Cost aggregation for open-vocabulary semantic segmentation. In *CVPR*, 4113–4123.
- Demir, I.; Koperski, K.; Lindenbaum, D.; Pang, G.; Huang, J.; Basu, S.; Hughes, F.; Tuia, D.; and Raskar, R. 2018. Deepglobe 2018: A challenge to parse the earth through satellite images. In *CVPRW*, 172–181.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- FAO. 2018. *Guidelines on the Use of Remote Sensing Products to Improve Agricultural Crop Production Forecast Statistics in Sub-Saharan African Countries*. UN.
- Garioud, A.; Gonthier, N.; Landrieu, L.; De Wit, A.; Valette, M.; Poupée, M.; Giordano, S.; et al. 2024. FLAIR: a country-scale land cover semantic segmentation dataset from multi-source optical imagery. *NeurIPS*, 36.
- Ghiasi, G.; Gu, X.; Cui, Y.; and Lin, T.-Y. 2022. Scaling open-vocabulary image segmentation with image-level labels. In *ECCV*, 540–557.
- Gupta, R.; Goodman, B.; Patel, N.; Hosfelt, R.; Sajeev, S.; Heim, E.; Doshi, J.; Lucas, K.; Choset, H.; and Gaston, M. 2019. Creating xBD: A dataset for assessing building damage from satellite imagery. In *CVPRW*, 10–17.
- Hua, Y.; Marcos, D.; Mou, L.; Zhu, X. X.; and Tuia, D. 2021. Semantic segmentation of remote sensing images with sparse annotations. *GRSL*, 19: 1–5.
- ISPRS. 2013a. 2D Semantic Labeling Potsdam Dataset. <https://www.isprs.org/education/benchmarks/UrbanSemLab/2d-sem-label-potsdam.aspx>. Accessed: 2024-08-11.
- ISPRS. 2013b. 2D Semantic Labeling Vaihingen Dataset. <https://www.isprs.org/education/benchmarks/UrbanSemLab/2d-sem-label-vaihingen.aspx>. Accessed: 2024-08-11.
- Jiang, X.; Zhou, N.; and Li, X. 2022. Few-Shot Segmentation of Remote Sensing Images Using Deep Metric Learning. *GRSL*, 19: 1–5.
- Li, B.; Weinberger, K. Q.; Belongie, S.; Koltun, V.; and Ranftl, R. 2022. Language-driven Semantic Segmentation. In *ICLR*.
- Li, J.; Sun, B.; Li, S.; and Kang, X. 2021. Semisupervised semantic segmentation of remote sensing images with consistency self-training. *TGRS*, 60: 1–11.
- Li, K.; Wan, G.; Cheng, G.; Meng, L.; and Han, J. 2020. Object detection in optical remote sensing images: A survey and a new benchmark. *ISPRS JPRS*, 159: 296–307.
- Liang, F.; Wu, B.; Dai, X.; Li, K.; Zhao, Y.; Zhang, H.; Zhang, P.; Vajda, P.; and Marculescu, D. 2023. Open-vocabulary semantic segmentation with mask-adapted clip. In *CVPR*, 7061–7070.
- Liu, F.; Chen, D.; Guan, Z.; Zhou, X.; Zhu, J.; Ye, Q.; Fu, L.; and Zhou, J. 2024a. Remoteclip: A vision language foundation model for remote sensing. *TGRS*.
- Liu, Y.; Bai, S.; Li, G.; Wang, Y.; and Tang, Y. 2024b. Open-vocabulary segmentation with semantic-assisted calibration. In *CVPR*, 3491–3500.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*, 3431–3440.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lyu, Y.; Vosselman, G.; Xia, G.-S.; Yilmaz, A.; and Yang, M. Y. 2020. UAVid: A semantic segmentation dataset for UAV imagery. *ISPRS JPRS*, 165: 108–119.
- Maggiori, E.; Tarabalka, Y.; Charpiat, G.; and Alliez, P. 2017. Can semantic labeling methods generalize to any city? the inria aerial image labeling benchmark. In *IGRSS*, 3226–3229.
- Open Cities AI Challenge. 2022. Open Cities AI Competition: Segmenting Buildings for Disaster Resilience. <https://www.drivendata.org/competitions/60/buildingsegmentation-disaster-resilience/>. Accessed: 2024-08-11.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Rahnemoonfar, M.; Chowdhury, T.; Sarkar, A.; Varshney, D.; Yari, M.; and Murphy, R. R. 2021. Floodnet: A high resolution aerial imagery dataset for post flood scene understanding. *IEEE Access*, 9: 89644–89654.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 234–241.

- Shan, X.; Wu, D.; Zhu, G.; Shao, Y.; Sang, N.; and Gao, C. 2024. Open-Vocabulary Semantic Segmentation with Image Embedding Balancing. In *CVPR*, 28412–28421.
- Shao, R.; Du, C.; Chen, H.; and Li, J. 2021. SUNet: Change detection for heterogeneous remote sensing images from satellite and UAV using a dual-channel fully convolution network. *RS*, 13(18): 3750.
- Tong, X.-Y.; Lu, Q.; Xia, G.-S.; and Zhang, L. 2018. Large-scale land cover classification in Gaofen-2 satellite imagery. In *IGRSS*, 3599–3602. IEEE.
- Van Etten, A.; Lindenbaum, D.; and Bacastow, T. M. 2018. Spacenet: A remote sensing dataset and challenge series. *arXiv preprint arXiv:1807.01232*.
- Van Westen, C. 2000. Remote sensing for natural disaster management. *International archives of photogrammetry and remote sensing*, 33(B7/4; PART 7): 1609–1617.
- Volpi, M.; and Ferrari, V. 2015. Semantic segmentation of urban scenes by learning local class interactions. In *CVPRW*, 1–9.
- Wang, D.; Zhang, J.; Du, B.; Xu, M.; Liu, L.; Tao, D.; and Zhang, L. 2024. Samrs: Scaling-up remote sensing segmentation dataset with segment anything model. *NeurIPS*, 36.
- Wang, J.; Zheng, Z.; Ma, A.; Lu, X.; and Zhong, Y. 2021. LoveDA: A remote sensing land-cover dataset for domain adaptive semantic segmentation. *arXiv preprint arXiv:2110.08733*.
- Waqas Zamir, S.; Arora, A.; Gupta, A.; Khan, S.; Sun, G.; Shahbaz Khan, F.; Zhu, F.; Shao, L.; Xia, G.-S.; and Bai, X. 2019. isaid: A large-scale dataset for instance segmentation in aerial images. In *CVPRW*, 28–37.
- Wu, Y.; Kirillov, A.; Massa, F.; Lo, W.-Y.; and Girshick, R. 2019. Detectron2. <https://github.com/facebookresearch/detectron2>.
- Xia, G.-S.; Bai, X.; Ding, J.; Zhu, Z.; Belongie, S.; Luo, J.; Datcu, M.; Pelillo, M.; and Zhang, L. 2018. DOTA: A large-scale dataset for object detection in aerial images. In *CVPR*, 3974–3983.
- Xia, J.; Yokoya, N.; Adriano, B.; and Broni-Bediako, C. 2023. Openearthmap: A benchmark dataset for global high-resolution land cover mapping. In *WACV*, 6254–6264.
- Xie, B.; Cao, J.; Xie, J.; Khan, F. S.; and Pang, Y. 2024. Sed: A simple encoder-decoder for open-vocabulary semantic segmentation. In *CVPR*, 3426–3436.
- Xu, M.; Zhang, Z.; Wei, F.; Hu, H.; and Bai, X. 2023. Side adapter network for open-vocabulary semantic segmentation. In *CVPR*, 2945–2954.
- Xu, M.; Zhang, Z.; Wei, F.; Lin, Y.; Cao, Y.; Hu, H.; and Bai, X. 2022. A simple baseline for open-vocabulary semantic segmentation with pre-trained vision-language model. In *ECCV*, 736–753.
- Yamazaki, K.; Hanyu, T.; Tran, M.; Garcia, A.; Tran, A.; McCann, R.; Liao, H.; Rainwater, C.; Adkins, M.; Molthan, A.; et al. 2023. AerialFormer: Multi-resolution Transformer for Aerial Image Segmentation. *arXiv preprint arXiv:2306.06842*.
- Yan, T.; Wan, Z.; and Zhang, P. 2022. Fully transformer network for change detection of remote sensing images. In *ACCV*, 1691–1708.
- Yan, T.; Wan, Z.; Zhang, P.; Cheng, G.; and Lu, H. 2023. Transy-net: Learning fully transformer networks for change detection of remote sensing images. *TGRS*.
- Yao, X.; Han, J.; Cheng, G.; Qian, X.; and Guo, L. 2016. Semantic annotation of high-resolution satellite images via weakly supervised learning. *TGRS*, 54(6): 3660–3671.
- Zheng, Z.; Zhong, Y.; Wang, J.; and Ma, A. 2020. Foreground-aware relation network for geospatial object segmentation in high spatial resolution remote sensing imagery. In *CVPR*, 4096–4105.
- Zhu, J.; Guo, Y.; Sun, G.; Yang, L.; Deng, M.; and Chen, J. 2023. Unsupervised domain adaptation semantic segmentation of high-resolution remote sensing imagery with invariant domain-level prototype memory. *TGRS*, 61: 1–18.