

# ZoRI: Towards Discriminative Zero-Shot Remote Sensing Instance Segmentation

Shiqi Huang<sup>1\*</sup>, Shuting He<sup>2\*</sup>, Bihan Wen<sup>1†</sup>

<sup>1</sup>Nanyang Technological University

<sup>2</sup>Shanghai University of Finance and Economics

shiqi006@e.ntu.edu.sg, heshuting555@gmail.com, bihan.wen@ntu.edu.sg

## Abstract

Instance segmentation algorithms in remote sensing are typically based on conventional methods, limiting their application to seen scenarios and closed-set predictions. In this work, we propose a novel task called zero-shot remote sensing instance segmentation, aimed at identifying aerial objects that are absent from training data. Challenges arise when classifying aerial categories with high inter-class similarity and intra-class variance. Besides, the domain gap between vision-language models’ pretraining datasets and remote sensing datasets hinders the zero-shot capabilities of the pretrained model when it is directly applied to remote sensing images. To address these challenges, we propose a **Zero-Shot Remote Sensing Instance Segmentation** framework, dubbed **ZoRI**. Our approach features a discrimination-enhanced classifier that uses refined textual embeddings to increase the awareness of class disparities. Instead of direct fine-tuning, we propose a knowledge-maintained adaptation strategy that decouples semantic-related information to preserve the pre-trained vision-language alignment while adjusting features to capture remote sensing domain-specific visual cues. Additionally, we introduce a prior-injected prediction with cache bank of aerial visual prototypes to supplement the semantic richness of text embeddings and seamlessly integrate aerial representations, adapting to the remote sensing domain. We establish new experimental protocols and benchmarks, and extensive experiments convincingly demonstrate that **ZoRI** achieves the state-of-art performance on the zero-shot remote sensing instance segmentation task.

**Code** — <https://github.com/HuangShiqi128/ZoRI>

## Introduction

Instance segmentation for remote sensing images aims to identify aerial objects and locate them with pixel-level masks for each instance (Su et al. 2019, 2020; Carvalho et al. 2020; Zhang et al. 2021b; Xu et al. 2021; Liu et al. 2024b). The task is of great importance for precise interpretation and comprehension of remote sensing images, which are essential for advanced earth observation. It has extensive real-world applications, such as military operations, marine

\*These authors contributed equally.

†Corresponding author

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

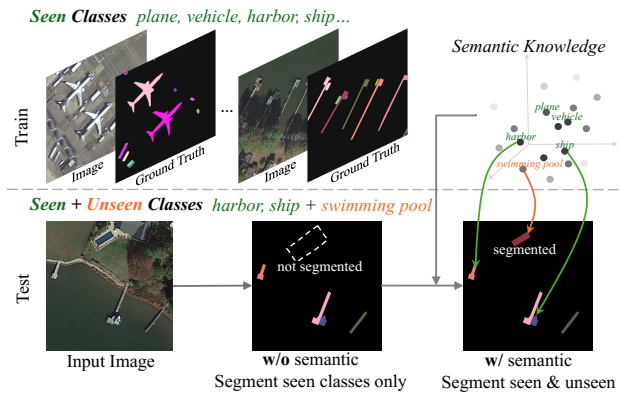


Figure 1: Illustration of zero-shot remote sensing instance segmentation, which transfers the learned semantic knowledge from seen classes, e.g., harbor and ship, to the unseen class, e.g., swimming pool.

monitoring, and urban planning (Liu et al. 2020; de Albuquerque et al. 2021; Wei et al. 2022; Yasir et al. 2023), making it a pivotal research area in earth vision studies.

The standard pipeline for aerial instance segmentation starts with annotating instance masks for remote sensing datasets, followed by training a segmentation model using conventional methods (Zamir et al. 2019; Su et al. 2019). However, due to relatively small object sizes and the large number of instances present in aerial images, labeling large scale datasets becomes time-consuming and cost-intensive. Consequently, applications are often constrained to limited scenarios. Moreover, training with conventional methods restricts the model to close-set predictions, making it unable to identify novel classes absent from training data.

We, for the first time, propose the zero-shot instance segmentation scheme for remote sensing images, which enables models to segment novel aerial object classes without instance mask annotations, as shown in Figure 1. The proposed scheme utilizes Vision-Language Models (VLMs) such as CLIP (Radford et al. 2021) which demonstrated remarkable capabilities in zero-shot tasks due to their highly aligned vision and language modalities (Ding et al. 2022; Xu et al. 2022). However, this alignment does not always extend consistently along channel dimensions (Zhu et al. 2023). Our investigation using aerial images reveals that not

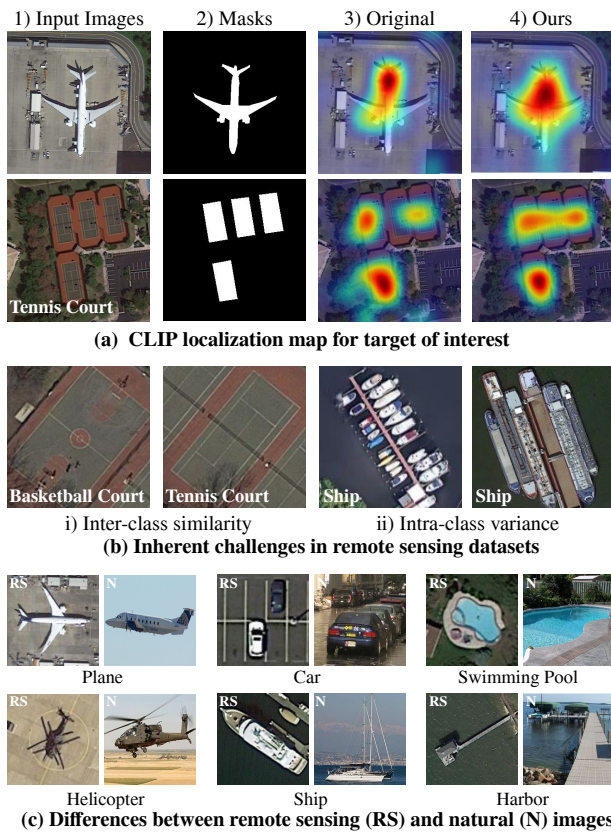


Figure 2: **(a)** After refinement, the head of *plane* is highlighted and the activation map strictly follows its shape; For *tennis court*, the activation map is more focused and the missed one in the middle using original channels is also emphasized. **(b)** Classes such as *basketball court* and *tennis court* share similar color and shape, whereas instances from *ship* can have various appearances. **(c)** Remote sensing images are in bird’s eye view while natural images are from ground-level perspective.

all CLIP-extracted text features attend to the corresponding visual targets across these channels. As shown in Figure 2(a) 3), the activation region deviates from the object of interest when using the original CLIP text features. This indicates the pretrained CLIP struggles to establish a semantic-visual correspondence for remote sensing images. We further contend that this problem can be worsened by the inherent challenges of high inter-class (between-class) similarity and intra-class (within-class) variance (Zhang et al. 2020; Cheng et al. 2020, 2021; Chen et al. 2023; Yang et al. 2022) in aerial objects, present in Figure 2(b). To address this, we propose a Discrimination-Enhanced Classifier (DEC), which refines the feature channels in the CLIP text embedding classifier to select most attentive ones to the target, encapsulating the most discriminative semantic information. After refinement, as shown in Figure 2(a) 4), the activation map accurately depicts the shape of the *plane* and highlights the middle *tennis court* missed in Figure 2(a) 3).

While directly applying pretrained VLMs achieves success with natural image, it fails to produce similar gains due

to the inherent domain gap between VLM pretraining and remote sensing datasets. Specifically, as shown in Figure 2(c), aerial images shot from above have different appearances compared to natural images, which are typically captured from a ground-level perspective and have a central layout. To bridge this gap, we design a Knowledge-Maintained Adaptation (KMA) with proposed decoupling strategy for CLIP visual features to adapt CLIP to the remote sensing domain without drastically forgetting pretrained knowledge or losing generalization ability. After decoupling, cross-modal alignment is preserved by retaining the decoupled semantic-related feature channels, while non-semantic-related knowledge is learned to adapt to the specific visual domain of remote sensing. This design not only maintains the well-trained discrimination power of CLIP but also derives tailored visual representations for aerial images.

Furthermore, although our predictions become more discriminative using DEC, text embeddings in the classifier still represent general concepts for natural images and convey limited information, leading to reduced classification capability with current zero-shot classifier. To inject aerial-specific knowledge, we introduce a Prior-Injected Prediction (PIP), which utilizes visual prototypes to complement the final prediction. By incorporating a cache bank constructed with remote sensing visual prototypes, domain-specific visual information is seamlessly integrated into the model. Text meanings in the original classifier are re-weighted to align with remote sensing characteristic, and visual information is added to supplement various class representations that cannot be delivered by text alone.

Our main contributions are summarized as follows:

- We introduce the problem of zero-shot remote sensing instance segmentation. To overcome inherent challenges in aerial imagery and the domain gap, we develop a **Z**ero-shot **R**emote sensing **I**nstance segmentation framework (**ZoRI**) and establish new experimental protocols and benchmarks. **ZoRI** achieves the state-of-art performance.
- To address the classification ambiguity for aerial objects, we propose DEC, which refines CLIP textual embeddings to enhance the discriminative power, improving the ability to classify different classes more distinctly.
- We design KMA with a decoupled visual feature learning strategy to adapt CLIP to the remote sensing domain while preserving its vision-language alignment.
- We identify the limitation of classifying with text embeddings and introduce PIP using a cache bank consisting of aerial visual prototypes to enhance semantic richness, thereby narrowing the representation gap between general textual expressions and remote sensing objects.

## Related Work

### Instance Segmentation in Remote Sensing

Segmentation frameworks for remote sensing images have seen significant advancements. Various methods have been developed to address challenges such as severe scale variations, foreground-background confusion, and class ambiguity (Liu et al. 2024b; Ye et al. 2023; Su et al. 2022,

2020, 2019; Liu and Di 2022; Zhang et al. 2021b) in remote sensing instance segmentation datasets. Recently, the integration of vision foundation models with remote sensing image segmentation has also been evolving. For instance, RSPrompter (Chen et al. 2024) improves the instance segmentation performance of SAM (Kirillov et al. 2023) on remote sensing images through prompt learning. Additionally, SAMRS (Wang et al. 2023) expands remote sensing segmentation datasets by leveraging SAM. However, current algorithms for remote sensing image segmentation primarily focus on close-set prediction, limiting their ability to recognize a broader array of object categories that are not present in the training set.

### Zero-shot Learning in Remote Sensing

Zero-shot learning has attracted significant research interest in the field of remote sensing. Earlier works primarily utilized language models like Word2vec (Mikolov et al. 2013) and BERT (Devlin et al. 2019) to extract semantic representations for zero-shot scene classification (Li et al. 2017, 2021; Wang, Peng, and De Baets 2021). Recently, with the emergence of vision-language models such as CLIP (Radford et al. 2021), research has shifted towards constructing highly-aligned vision-language embedding space (Li et al. 2023; Liu et al. 2024a; Zavras et al. 2024; Mall et al. 2023). DescReg (Zang et al. 2024) presents a zero-shot aerial object detection framework by introducing visual descriptions to regularize semantic embeddings retrieved from BERT. However, zero-shot remote sensing instance segmentation remains unexplored.

### Zero-shot Segmentation

Zero-shot segmentation aims to identify objects beyond seen classes during training with pixel-wise masks. Owing to the powerful generalization capability of CLIP, many zero-shot segmentation models are developed by harnessing its zero-shot classification ability. Previous methods (Ding et al. 2022; Xu et al. 2022; Liang et al. 2023; He et al. 2024) retrieve class-agnostic masks with a mask generator and classify using a CLIP-based classifier. To avoid duplicate feature extractions for mask generation and CLIP classification, ZegCLIP (Zhou et al. 2023) extends CLIP zero-shot capability from image-level to pixel-level by deep prompt tuning. FC-CLIP (Yu et al. 2023) further demonstrates the efficacy of zero-shot classification using the frozen convolutional CLIP backbone. Zero-shot instance segmentation was first introduced in ZSI (Zheng et al. 2021), followed by works (He, Ding, and Jiang 2023a,b) to handle classification bias and ambiguity issues.

## Method

### Problem Definition

In zero-shot remote sensing instance segmentation (ZSRI), there are  $N$  classes divided into two groups,  $N^s$  seen classes denoted as  $C^s$  and  $N^u$  unseen classes denoted as  $C^u$ , with  $C^s \cap C^u = \emptyset$  and  $N^s + N^u = N$ , where  $C^s = \{c_1^s, c_2^s, \dots, c_{N^s}^s\}$ ,  $C^u = \{c_1^u, c_2^u, \dots, c_{N^u}^u\}$ . The training set comprises images  $x_s$  with seen classes  $C^s$  but no

object of unseen classes  $C^u$ , along with pixel-wise instance mask annotations for each object in  $C^s$ . During inference, two settings are considered: ZSRI and Generalized ZSRI (GZSRI). In ZSRI setting, the goal is to segment objects of only  $N^u$  unseen classes, while in GZSRI, the object is to segment all  $N^u + N^s$  classes, including both seen and unseen, which is more reflective of real-world scenarios.

### Architecture Overview

The architecture of our proposed **ZoRI** is illustrated in Figure 3. We adopt the paradigm of FC-CLIP (Yu et al. 2023), developed based on Mask2Former (Cheng et al. 2022). The CLIP image encoder is partially trained with knowledge-maintained adaptation (KMA) to extract backbone features, which are then fed into a mask generator to produce mask predictions and class embeddings. Text embeddings are encoded by the frozen CLIP text encoder and further refined to construct discriminative-enhanced classifier (DEC), which is used to classify class embeddings. During inference, cache bank is incorporated into CLIP zero-shot predictions to derive prior-injected predictions (PIP). The final classification probability is obtained through an ensemble approach similar in FC-CLIP, by combining scores based on trained class embeddings, and prior-injected predictions based on mask-pooled CLIP backbone features.

### Discrimination-Enhanced Classifier (DEC)

Considering the inherent challenges present in remote sensing classes, such as high inter-class similarity and intra-class variance, a naive classifier constructed from text embeddings struggles to overcome the resulting classification ambiguity. To address this, we propose a discrimination-enhanced classifier, which refines the naive classifier by selecting the most discriminative channels, discarding redundant ones which are less attentive to object targets and may introduce noise during classification. The resulting text embeddings in the classifier will have an improved one-to-one correspondence with the target objects.

For each category  $c_i$  from  $C^s \cup C^u$ , we obtain its category text embedding by placing the category name into prompt templates, *e.g.*, ‘satellite imagery of { }.’, ‘aerial imagery of { }.’, and inputting to the CLIP text encoder. The final text embedding for each class is obtained by averaging the CLIP-extracted features across all templates, denoted as  $x_i \in \mathbb{R}^D$ , where  $i \in \{1, 2, \dots, N\}$ . The zero-shot classifier is then constructed from text embeddings  $[x_1 \ x_2 \ \dots \ x_N] \in \mathbb{R}^{D \times N}$ .

To refine the classifier, we select top-k channels from  $D$  dimensions that minimize similarity and maximize variance across different classes. The similarity among channels across classes is represented as:

$$S = \frac{1}{N \times (N - 1)} \sum_{i=1}^N \sum_{j=1, j \neq i}^N x_i x_j \in \mathbb{R}^D, \quad (1)$$

where  $x_i$  and  $x_j$  are L2-normalized. The variance is formulated as:

$$V = \frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2 \in \mathbb{R}^D, \quad (2)$$

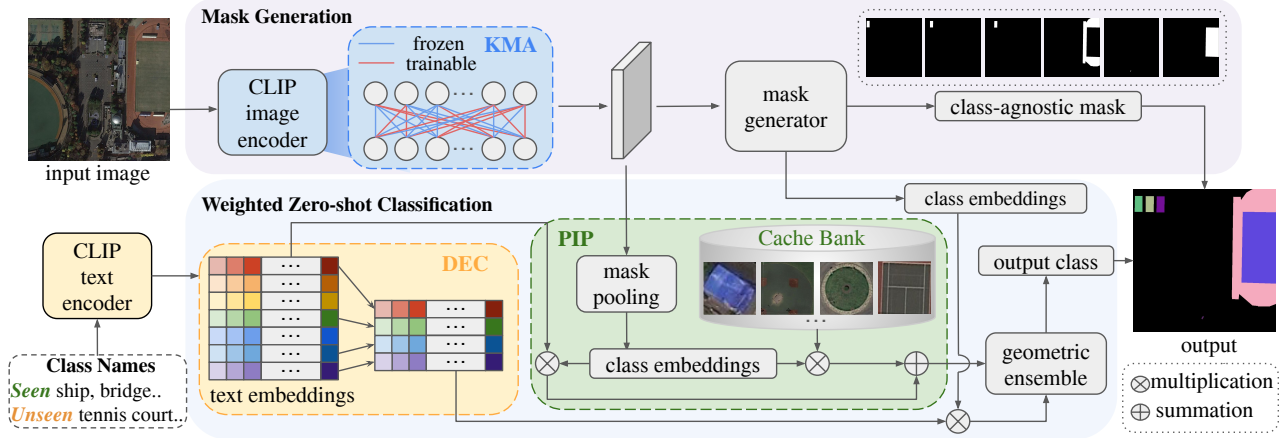


Figure 3: Overview of **ZoRI** framework. The CLIP image encoder is partially trained with knowledge-maintained adaptation (KMA) to extract backbone features, which are then fed into a mask generator to produce mask predictions and class embeddings. Discriminative-enhanced classifier (DEC) constructed by refining the original text embeddings is then used to classify class embeddings. During inference, cache bank is incorporated into CLIP zero-shot predictions to derive prior-injected predictions (PIP). The final classification probability is obtained through an ensemble approach.

where  $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i$  is the average across categories. Taking both similarity and variance into account, we formulate our objective as:

$$O = -\lambda S + (1 - \lambda)V \in \mathbb{R}^D, \quad (3)$$

where  $\lambda$  is a balance factor. The indices corresponding to the top-k largest values in  $O \in \mathbb{R}^D$  are selected as the refined feature channels. These channels possess the highest discriminative power by suppressing resemblance and enhancing distinctiveness among different categories, thereby are most attentive to object targets.

### Knowledge-Maintained Adaptation (KMA)

The shared frozen convolutional CLIP backbone (Yu et al. 2023) offers a favorable balance between accuracy and computational cost by combining feature extraction for mask generation and CLIP zero-shot prediction. However, this frozen pretrained CLIP vision backbone struggles with generating good representations for aerial images. VLMs like CLIP are pretrained on large-scale web-based image-text pairs predominantly featuring natural images captured at ground level, leading to a domain gap when applied to aerial images. Objects viewed from a bird’s-eye view exhibit different appearances compare to those in natural images, which are typically captured from a ground-level perspective and have a central layout. Applying the frozen CLIP backbone directly inevitably overlooks such difference and fails to use training priors to improve feature extraction for aerial images. To adapt the pretrained backbone to remote sensing domain, directly fine-tuning tends to overfit to seen classes, resulting in drastically forgetting pretrained knowledge and compromising the model’s ability to handle new, unseen classes (Zhou, Loy, and Dai 2022; Xu et al. 2023; Zhou et al. 2023). To transfer to remote sensing domain without compromising generalization ability, we propose a decoupling strategy that separates visual features into two groups. The first group includes channels more

closely related to semantics, which exhibit better vision-language alignment and are thus more discriminative in classifying objects. The second group consists of channels less connected to semantic meaning but rich in visual cues.

Object instances for each seen class from  $C^s$  can be cropped from the downstream training dataset with ground truth annotations. Suppose we retrieve  $T$  instances for each class, and extract intermediate backbone features for these instances with the pretrained CLIP vision encoder. To decouple backbone feature channels into two groups, we follows the criterion present in Eq.(3). Similarity and variance of backbone features among categories of  $T$  instances are calculated in a manner similar to text embeddings  $x_i$ . Channels with lowest similarity and highest variance are considered as the first group. Due to its highly aligned vision and language embedding space, we freeze them to keep them intact which helps maintain the pretrained cross-modality alignment. The remaining channels are less related to semantics and have knowledge in visual representation, thus we make them trainable to adapt to remote sensing specific visual domain. After decoupling, visual representations can be tuned to adapt to remote sensing domain while preserving cross-modal alignment. This design not only maintains the well-trained discrimination power of CLIP, but also derives tailored visual representation for aerial images.

### Prior-Injected Prediction (PIP)

Current classification results solely rely on text embeddings extracted from category descriptions using the CLIP text encoder. However, due to the high intra-class variability within remote sensing images, objects in the same category can have diverse appearances. Text features convey limited information and are insufficient to encompass this variability. Consequently, text classifier may struggle to establish a reliable correspondence between various aerial visual features and pretrained textual features, resulting in less discriminative ability. Therefore, we propose to introduce

visual prototypes to enhance classification capability by covering various vision samples through a cache bank. By integrating the cache bank constructed with visual prototypes, the classification score derived from text classifier is re-weighted to align with remote sensing domain. This approach incorporates aerial-specific visual prior to enhance the representation of various classes, addressing the limitations that cannot be delivered through text alone.

**Cache Bank Construction** To construct a cache bank with supplementary visual representations specific to remote sensing domain, we first crop object instances for each seen class  $c_i^s \in \{c_1^s, c_2^s, \dots, c_{N^s}^s\}$  from the downstream training dataset with ground truth annotations. Suppose we retrieve  $K$  sample instances for each class, denoted as  $\{I_n\}_{n=1}^{KN^s}$ , these instances are regarded as representatives for different categories in  $\{c_1^s, c_2^s, \dots, c_{N^s}^s\}$  in aerial images. Visual embeddings for sample instances  $\{I_n\}_{n=1}^{KN^s}$  are then obtained using the pretrained CLIP image encoder  $E_{vis}$ . Then, we concatenate the training-set instance visual features along the sample dimension, and store them as the cache bank in remote sensing domain, formulated as:

$$F = \text{concat}(\{E_{vis}(I_n)\}_{n=1}^{KN^s}) \in \mathbb{R}^{KN^s \times D}, \quad (4)$$

where  $D$  denotes the feature dimension of the vision encoder. The cache bank  $F$  is constructed in a training-free manner similar to works in few-shot learning (Rong et al. 2023; Zhang et al. 2021a, 2023). Such bank comprises diverse representations retrieved from remote sensing visual examples, which augment CLIP with aerial knowledge, different from natural images for its bird’s eye view.

For  $N^u$  unseen classes whose ground truth annotations are assumed unavailable, we use  $P$  predictions with the highest probability as visual samples for unseen categories. Then, visual embeddings are also retrieved using the pre-trained CLIP image encoder and concatenated to the cache bank, which is updated to obtain  $F \in \mathbb{R}^{(KN^s + PN^u) \times D}$ .

**Cache Bank Prediction** After getting mask predictions, mask-pooling is applied to the CLIP-extracted backbone feature for the input image, denoted as  $f \in \mathbb{R}^{1 \times D}$  for each predicted segment. The mask-pooled CLIP backbone feature serves as the query, and the pre-encoded instance features  $F$  can be regarded as keys. The cosine similarity metric between the query and keys is calculated as:

$$M = fF^T \in \mathbb{R}^{1 \times (KN^s + PN^u)}, \quad (5)$$

where mask-pooled feature and cache bank features are L2-normalized. Along with the one-hot class label  $l \in \mathbb{R}^{1 \times N}$  associated with each sample instance in the cache bank, values for query and key are thus constructed by concatenating one-hot labels, denoted as  $L \in \mathbb{R}^{(KN^s + PN^u) \times N}$ . We thus formulate the cache bank prediction as below:

$$\text{logits}_{\text{cb}} = \text{softmax}(M)L \in \mathbb{R}^{1 \times N}. \quad (6)$$

The cache bank prediction is fused into CLIP zero-shot prediction to derive the prior-injected prediction with the formulation:

$$\text{logits}_{\text{pip}} = fW^T + \alpha \text{logits}_{\text{cb}} \in \mathbb{R}^{1 \times N}, \quad (7)$$

where  $W \in \mathbb{R}^{N \times D}$  is CLIP’s classifier,  $fW^T$  calculates the classification score based on pretrained zero-shot prediction from CLIP,  $\alpha$  is a balance factor.

## Experiments

### Experimental Setup

**Datasets** We establish two remote sensing zero-shot instance segmentation benchmarks with iSAID (Zamir et al. 2019) and NWPU-VHR-10 (Cheng et al. 2014; Su et al. 2019) datasets. iSAID is a large-scale dataset for instance segmentation in aerial images containing 15 classes. NWPU-VHR-10 is a very high resolution (VHR) object detection dataset (Cheng et al. 2014) further annotated with instance masks (Su et al. 2019), which contains 10 classes. iSAID dataset is divided into 11 seen classes and 4 unseen classes (‘tennis court’, ‘helicopter’, ‘swimming pool’ and ‘soccer ball field’), which has the same seen/unseen split for DOTA (Zang et al. 2024; Xia et al. 2018), and NWPU-VHR-10 dataset is split into 7 seen classes and 3 unseen classes (‘ship’, ‘basketball court’ and ‘harbor’). For the training set, only images containing seen class objects are selected, while any images with unseen classes are excluded to avoid information leakage. Both ZSRI and GZSRI are evaluated for testing. Dataset details can be found in the supplementary material (Huang, He, and Wen 2024).

**Evaluation Metrics** Referring to previous works (Zheng et al. 2021; He, Ding, and Jiang 2023b), we use Recall@100, i.e., top 100 instances, across different IoU thresholds {0.4, 0.5, 0.6} and mean Average Precision (mAP) with IoU threshold 0.5 as the metrics. For GZSRI setting, harmonic mean (HM) (Xian et al. 2020) of seen and unseen classes is computed to indicate overall performance.

**Implementation Details** The proposed method is developed based on FC-CLIP (Yu et al. 2023). All hyperparameters remain unchanged unless otherwise specified. We use the LAION-2B pretrained ConvNext-Large (Liu et al. 2022) from OpenCLIP (Ilharco et al. 2021) as the feature extractor. The mask generator follows Mask2Former (Cheng et al. 2022) with object query number set to 300. Prompt templates for RESISC45 (Cheng, Han, and Lu 2017) used in CLIP (Radford et al. 2021) are employed to obtain text embeddings with the pretrained CLIP text encoder. We train the model for 50 epochs with training batch size 2. Input images are resized to  $512 \times 512$  during training. Hyperparameters  $\lambda$  and  $\alpha$  are set to 0.7, 0.5, respectively. Instance number  $T$  and trainable channels in KMA is empirically set to 1 and 32. The model is optimized using AdamW optimizer. The learning rate is set to  $1.25 \times 10^{-5}$ . All experiments are conducted with one RTX5000 GPU.

### Component Analysis

We conduct extensive experiments to demonstrate the effectiveness of proposed modules in Table 1 with both iSAID (Zamir et al. 2019) and NWPU-VHR-10 (Cheng et al. 2014) datasets. FC-CLIP (Yu et al. 2023) is used as our baseline with replaced prompt templates RESISC45 for remote sensing images to derive text embeddings for the

Dataset	DEC	KMA	PIP	mAP $\uparrow$			Recall@100 $\uparrow$		
				Seen	Unseen	HM	Seen	Unseen	HM
iSAID	$\times$	$\times$	$\times$	43.47	4.91	8.83	67.01	36.31	47.10
	$\checkmark$	$\times$	$\times$	46.56	8.28	14.05	68.24	40.24	50.62
	$\checkmark$	$\checkmark$	$\times$	45.90	8.64	14.55	68.14	<b>41.52</b>	<b>51.60</b>
	$\checkmark$	$\checkmark$	$\checkmark$	<b>47.05</b>	<b>9.30</b>	<b>15.53</b>	<b>68.89</b>	37.73	48.76
NWPU	$\times$	$\times$	$\times$	80.85	6.86	12.65	94.25	36.16	52.27
	$\checkmark$	$\times$	$\times$	<b>82.90</b>	8.81	15.93	<b>94.81</b>	38.11	54.37
	$\checkmark$	$\checkmark$	$\times$	80.65	9.32	16.70	93.53	<b>46.44</b>	<b>62.07</b>
	$\checkmark$	$\checkmark$	$\checkmark$	80.57	<b>12.26</b>	<b>21.28</b>	93.80	44.01	59.91

Table 1: Component Analysis of **ZoRI** under GZSRI.

#Channel	200	300	400	500	600	700	768	
mAP $\uparrow$	Seen	45.61	<b>46.56</b>	44.99	45.09	44.39	43.69	43.47
	Unseen	6.26	<b>8.28</b>	8.08	7.09	6.59	6.19	4.93
	HM	11.02	<b>14.06</b>	13.69	12.26	11.47	10.86	8.86

Table 2: Ablation study on number of selected channels in text classifier on iSAID dataset.

classifier. From the baseline result, we can see the performance on unseen classes lags behind seen classes by a large margin, which indicates the model does not generalize well on aerial images. In the following sections, we provide both quantitative and qualitative analysis of our proposed modules. By integrating all the modules, we achieve a final result that significantly outperforms the baseline.

**Discrimination-Enhanced Classifier** By selecting the most discriminative channels in CLIP text classifier, we observe enhancements in both seen and unseen classes in Table 1. Due to these universal improvements, there is a notable increase of 5.22% HM-mAP for iSAID and 3.28% HM-mAP for NWPU-VHR-10. The improvements in both seen and unseen classes indicate that the classifier can effectively capture the most deterministic feature for each aerial object. As a result, the weak semantic correspondence and classification ambiguity caused by high inter-class similarity and intra-class variance in remote sensing categories is reduced, allowing for more accurate recognition of class instances.

In Table 2, we conduct an ablation study to evaluate the impact of the number of selected channels on the text classifier’s performance on iSAID dataset. The original text embeddings have 768 channels. The results indicate that there is an optimal range where the classifier’s discriminatory power reaches its peak. Beyond this range, performance declines due to excessive information loss or insufficient disentanglement of the text embeddings. From the result, we thereby select the top-300 channels with the highest score formulated in Eq.(3).

**Knowledge-Maintained Adaptation** By integrating visual feature decoupling for knowledge-maintained adaptation with DEC, the mAP for unseen classes sees a cumulative improvement of 3.73% on iSAID dataset. The HM-recall increases by nearly 10% for NWPU dataset. The improvement demonstrates that the CLIP vision encoder

Method	mAP $\uparrow$			Recall@100 $\uparrow$		
	Seen	Unseen	HM	Seen	Unseen	HM
baseline	43.47	4.91	8.83	67.01	36.31	47.10
baseline+FT	43.97	5.71	10.10	<b>67.33</b>	36.69	47.50
baseline+VPT	40.93	7.18	12.22	62.46	33.46	43.57
baseline+KMA	<b>44.08</b>	<b>7.19</b>	<b>12.36</b>	67.19	<b>40.96</b>	<b>50.90</b>

Table 3: Comparison with other finetuning methods on iSAID dataset. The baseline method uses a frozen vision encoder, while FT denotes full finetuning of vision encoder.

Cache Size	0	1	4	8	16	32	
mAP $\uparrow$	Seen	45.90	46.07	47.05	47.09	47.09	<b>47.11</b>
	Unseen	8.64	8.69	<b>9.30</b>	9.27	9.29	9.28
	HM	14.55	14.62	<b>15.53</b>	15.49	15.52	15.51

Table 4: Ablation study on cache size on iSAID dataset.

has successfully adapted to the remote sensing domain and can generate tailored representations for aerial objects. It is worth noting that the unseen mAP shows a significant improvement for both datasets, indicating the well-pretrained vision-language alignment is preserved during the adaptation process, maintaining the model’s zero-shot ability to handle novel classes.

To validate our adaptation method, we report the results using other finetuning techniques in Table 3. Compared with full finetuning, we achieve a larger seen mAP while only train a few parameters and our unseen mAP is about 1.5% higher. This indicates our method is more effective in adaptation while preserving pretrained knowledge. We also report the result using VPT (Jia et al. 2022) following their implementation for ConvNext. From the results, we can see that KMA outperforms VPT in both seen and unseen classes, with a 7.33% increase in HM-recall. Our advantage demonstrates the efficacy of our proposed module KMA.

**Prior-Injected Prediction** Incorporating a cache bank of visual priors into the final class prediction leads to performance improvements on both iSAID and NWPU datasets, as seen from Table 1. Specifically, on the iSAID dataset, there is an overall 6.7% gain in HM-mAP when combined with DEC and KMA compared to the baseline. Similarly, on the NWPU dataset, there is a 8.63% improvement in HM-mAP. This enhancement demonstrates that integrating visual samples through the cache bank boosts the model’s accuracy by exposing it to a wider variety of remote sensing objects. The visual prototypes effectively address the limitations of text descriptions by considering the intra-class variability of aerial objects, thus facilitating better adaptation to the remote sensing domain.

To examine how cache size influences prediction accuracy, we create cache bank with varying numbers of visual prototypes, and report the results in Table 4. The results indicate that a moderate number of prototypes from each class is enough to improve the performance, and only 1 sample cannot capture sufficient variability. To balance computational cost and accuracy, we choose 4 visual prototypes.

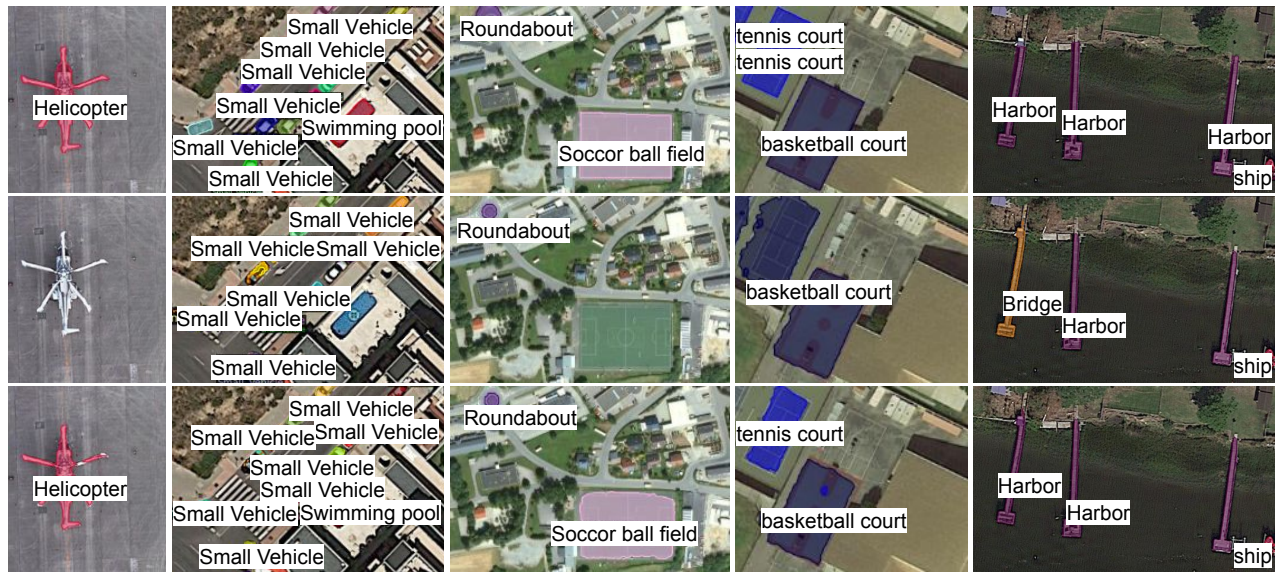


Figure 4: Comparison of GZSRI results: (top) ground truth, (middle) FC-CLIP (Yu et al. 2023) and (bottom) our results.

Dataset	Method	mAP $\uparrow$			Recall@100 $\uparrow$		
		Seen	Unseen	HM	Seen	Unseen	HM
iSAID	ZSI (Zheng et al. 2021)	45.83	0.64	1.26	61.13	16.66	26.19
	D <sup>2</sup> Zero (He et al. 2023)	39.80	0.39	0.77	63.45	10.16	17.52
	FC-CLIP (Yu et al. 2023)	43.47	4.91	8.83	67.01	36.31	47.10
	<b>ZoRI</b>	<b>47.05</b>	<b>9.30</b>	<b>15.53</b>	<b>68.89</b>	<b>37.73</b>	<b>48.76</b>
	Supervised	52.40	0.0	-	72.89	0.0	-
NWPU	ZSI (Zheng et al. 2021)	<b>83.01</b>	0.13	0.25	89.90	11.10	19.76
	D <sup>2</sup> Zero (He et al. 2023)	77.62	0.41	0.81	<b>95.41</b>	9.78	17.75
	FC-CLIP (Yu et al. 2023)	80.85	6.86	12.65	94.25	36.16	52.27
	<b>ZoRI</b>	80.57	<b>12.26</b>	<b>21.28</b>	93.80	<b>44.01</b>	<b>59.91</b>
	Supervised	83.87	0.0	-	95.58	0.0	-

Table 5: Comparison under GZSRI setting.

## Comparison with State-of-the-Art Methods

In Table 5 and Table 6, we compare with SOTA methods on iSAID and NWPU-VHR-10 datasets, evaluated under the GZSRI and ZSRI settings respectively. We also provide the supervised results by replacing CLIP text classifier to the conventional trainable classifier for reference. As we can see from the tables, our proposed **ZoRI** surpasses all compared methods. Under the GZSRI setting, our approach achieve a notable improvement of 14.27% HM-mAP over ZSI (Zheng et al. 2021) on the iSAID dataset and 21.03% HM-mAP on the NWPU-VHR-10 dataset. Additionally, we outperform D<sup>2</sup>Zero by 14.76% and 20.47% HM-mAP on each dataset. Compared with FC-CLIP, we achieve 6.7% and 8.63% gain on HM-mAP on iSAID and NWPU respectively. Under the ZSRI setting, **ZoRI** obtains more than 5% gain in mAP and significant increases in recall across various thresholds compared to the best-reported method on both datasets.

Figure 4 shows a qualitative comparison with FC-CLIP on the iSAID dataset for both seen and unseen classes. Our method successfully segments the unseen objects missed out by FC-CLIP, such as *helicopter*, *swimming pool*, and

Dataset	Method	mAP $\uparrow$		Recall@100 $\uparrow$	
		0.5	0.4	0.5	0.6
iSAID	ZSI (Zheng et al. 2021)	1.5	22.6	16.7	10.6
	D <sup>2</sup> Zero (He et al. 2023)	0.5	36.6	28.1	19.4
	FC-CLIP (Yu et al. 2023)	3.0	51.4	42.2	30.8
	<b>ZoRI</b>	<b>8.6</b>	<b>55.2</b>	<b>46.0</b>	<b>36.6</b>
NWPU	ZSI (Zheng et al. 2021)	0.6	13.9	11.1	7.9
	D <sup>2</sup> Zero (He et al. 2023)	2.7	23.9	19.2	14.3
	FC-CLIP (Yu et al. 2023)	4.5	46.2	37.1	30.8
	<b>ZoRI</b>	<b>9.6</b>	<b>55.9</b>	<b>47.3</b>	<b>33.8</b>

Table 6: Comparison under ZSRI setting.

*soccer ball field* in the first three columns. This is because our model adapts to remote sensing domain and has better knowledge for aerial objects. Notably, we correctly identify similar categories like *tennis court* and *basketball court*, *harbor* and *bridge*, which are misclassified by FC-CLIP. This highlights the effectiveness of our discrimination-enhanced classifier. The proposed method **ZoRI** shows much better results by constructing a more discriminative model adapting to remote sensing domain.

## Conclusion

We propose a new zero-shot remote sensing instance segmentation task and present a novel framework **ZoRI**. Discrimination-enhanced classifier is developed to tackle the classification ambiguity in aerial objects. To adapt CLIP to remote sensing domain, we introduce knowledge-maintained adaptation to derive tailored representations for aerial objects. Additionally, our prior-injected prediction supports text embedding classifier with diverse aerial objects. With new benchmarks established, **ZoRI** achieves SOTA performance in zero-shot remote sensing instance segmentation. This advancement sets the stage for future research in zero-shot segmentation of aerial imagery.

## Acknowledgments

This work was supported in part by the Economic Development Board, Republic of Singapore, through its Space Technology Development Programme Grant under Grant S22-02004-STDP.

## References

- Carvalho, O. L. F. d.; de Carvalho Junior, O. A.; Albuquerque, A. O. d.; Bem, P. P. d.; Silva, C. R.; Ferreira, P. H. G.; Moura, R. d. S. d.; Gomes, R. A. T.; Guimaraes, R. F.; and Borges, D. L. 2020. Instance segmentation for large, multi-channel remote sensing imagery using mask-RCNN and a mosaicking approach. *Remote Sensing*.
- Chen, K.; Liu, C.; Chen, H.; Zhang, H.; Li, W.; Zou, Z.; and Shi, Z. 2024. RSPrompter: Learning to prompt for remote sensing instance segmentation based on visual foundation model. *TGRS*.
- Chen, Y.; Li, Y.; Mao, H.; Liu, G.; Chai, X.; and Jiao, L. 2023. A Novel Discriminative Enhancement Method for Few-Shot Remote Sensing Image Scene Classification. *Remote Sensing*.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshick, R. 2022. Masked-attention Mask Transformer for Universal Image Segmentation. In *CVPR*.
- Cheng, G.; Han, J.; and Lu, X. 2017. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*.
- Cheng, G.; Han, J.; Zhou, P.; and Guo, L. 2014. Multi-class geospatial object detection and geographic image classification based on collection of part detectors. *ISPRS Journal of Photogrammetry and Remote Sensing*.
- Cheng, G.; Si, Y.; Hong, H.; Yao, X.; and Guo, L. 2021. Cross-Scale Feature Fusion for Object Detection in Optical Remote Sensing Images. *GRSL*.
- Cheng, G.; Xie, X.; Han, J.; Guo, L.; and Xia, G.-S. 2020. Remote sensing image scene classification meets deep learning: Challenges, methods, benchmarks, and opportunities. *J-STARs*.
- de Albuquerque, A. O.; de Carvalho, O. L. F.; e Silva, C. R.; de Bem, P. P.; Gomes, R. A. T.; Borges, D. L.; Guimarães, R. F.; Pimentel, C. M. M.; and de Carvalho Júnior, O. A. 2021. Instance segmentation of center pivot irrigation systems using multi-temporal SENTINEL-1 SAR images. *Remote Sensing Applications: Society and Environment*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805.
- Ding, J.; Xue, N.; Xia, G.-S.; and Dai, D. 2022. Decoupling Zero-Shot Semantic Segmentation. In *CVPR*.
- He, S.; Ding, H.; and Jiang, W. 2023a. Primitive Generation and Semantic-related Alignment for Universal Zero-Shot Segmentation. In *CVPR*.
- He, S.; Ding, H.; and Jiang, W. 2023b. Semantic-Promoted Debiasing and Background Disambiguation for Zero-Shot Instance Segmentation. In *CVPR*.
- He, S.; Ding, H.; Jiang, X.; and Wen, B. 2024. SegPoint: Segment Any Point Cloud via Large Language Model. In *ECCV*.
- Huang, S.; He, S.; and Wen, B. 2024. ZoRI: Towards Discriminative Zero-Shot Remote Sensing Instance Segmentation. arXiv:2412.12798.
- Ilharco, G.; Wortsman, M.; Wightman, R.; Gordon, C.; Carlini, N.; Taori, R.; Dave, A.; Shankar, V.; Namkoong, H.; Miller, J.; Hajishirzi, H.; Farhadi, A.; and Schmidt, L. 2021. OpenCLIP.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual Prompt Tuning. In *ECCV*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. arXiv:2304.02643.
- Li, A.; Lu, Z.; Wang, L.; Xiang, T.; and Wen, J.-R. 2017. Zero-shot scene classification for high spatial resolution remote sensing images. *TGRS*.
- Li, X.; Wen, C.; Hu, Y.; and Zhou, N. 2023. RS-CLIP: Zero shot remote sensing scene classification via contrastive vision-language supervision. *International Journal of Applied Earth Observation and Geoinformation*.
- Li, Y.; Zhu, Z.; Yu, J.-G.; and Zhang, Y. 2021. Learning deep cross-modal embedding networks for zero-shot remote sensing image scene classification. *TGRS*.
- 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, X.; and Di, X. 2022. Global Context Parallel Attention for Anchor-Free Instance Segmentation in Remote Sensing Images. *GRSL*.
- Liu, Y.; Chen, D.; Ma, A.; Zhong, Y.; Fang, F.; and Xu, K. 2020. Multiscale U-Shaped CNN Building Instance Extraction Framework With Edge Constraint for High-Spatial-Resolution Remote Sensing Imagery. *TGRS*.
- Liu, Y.; Li, H.; Hu, C.; Luo, S.; Luo, Y.; and Chen, C. W. 2024b. Learning to aggregate multi-scale context for instance segmentation in remote sensing images. *IEEE Transactions on Neural Networks and Learning Systems*.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A ConvNet for the 2020s. arXiv:2201.03545.
- Mall, U.; Phoo, C. P.; Liu, M. K.; Vondrick, C.; Hariharan, B.; and Bala, K. 2023. Remote Sensing Vision-Language Foundation Models without Annotations via Ground Remote Alignment.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient Estimation of Word Representations in Vector Space. arXiv:1301.3781.

- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020*.
- Rong, J.; Chen, H.; Chen, T.; Ou, L.; Yu, X.; and Liu, Y. 2023. Retrieval-enhanced visual prompt learning for few-shot classification. *arXiv preprint arXiv:2306.02243*.
- Su, H.; Huang, P.; Yin, J.; and Zhang, X. 2022. Faster and better instance segmentation for large scene remote sensing imagery. In *IGARSS*. IEEE.
- Su, H.; Wei, S.; Liu, S.; Liang, J.; Wang, C.; Shi, J.; and Zhang, X. 2020. HQ-ISNet: High-quality instance segmentation for remote sensing imagery. *Remote Sensing*.
- Su, H.; Wei, S.; Yan, M.; Wang, C.; Shi, J.; and Zhang, X. 2019. Object detection and instance segmentation in remote sensing imagery based on precise mask R-CNN. In *IGARSS*. IEEE.
- Wang, C.; Peng, G.; and De Baets, B. 2021. A distance-constrained semantic autoencoder for zero-shot remote sensing scene classification. *IEEE J-STARS*.
- Wang, D.; Zhang, J.; Du, B.; Xu, M.; Liu, L.; Tao, D.; and Zhang, L. 2023. SAMRS: Scaling-up Remote Sensing Segmentation Dataset with Segment Anything Model. In *NeurIPS*.
- Wei, S.; Zeng, X.; Zhang, H.; Zhou, Z.; Shi, J.; and Zhang, X. 2022. LFG-Net: Low-level Feature Guided Network for Precise Ship Instance Segmentation in SAR Images. *TGRS*.
- 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*.
- Xian, Y.; Lampert, C. H.; Schiele, B.; and Akata, Z. 2020. Zero-Shot Learning – A Comprehensive Evaluation of the Good, the Bad and the Ugly. *arXiv:1707.00600*.
- Xu, M.; Zhang, Z.; Wei, F.; Hu, H.; and Bai, X. 2023. Side Adapter Network for Open-Vocabulary Semantic Segmentation. In *CVPR*.
- 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*.
- Xu, X.; Feng, Z.; Cao, C.; Li, M.; Wu, J.; Wu, Z.; Shang, Y.; and Ye, S. 2021. An improved swin transformer-based model for remote sensing object detection and instance segmentation. *Remote Sensing*.
- Yang, X.; Yan, J.; Liao, W.; Yang, X.; Tang, J.; and He, T. 2022. Scrdet++: Detecting small, cluttered and rotated objects via instance-level feature denoising and rotation loss smoothing. *IEEE TPAMI*.
- Yasir, M.; Zhan, L.; Liu, S.; Wan, J.; Hossain, M. S.; Isiacik Colak, A. T.; Liu, M.; Islam, Q. U.; Raza Mehdi, S.; and Yang, Q. 2023. Instance segmentation ship detection based on improved Yolov7 using complex background SAR images. *Frontiers in Marine Science*.
- Ye, W.; Zhang, W.; Lei, W.; Zhang, W.; Chen, X.; and Wang, Y. 2023. Remote sensing image instance segmentation network with transformer and multi-scale feature representation. *Expert Systems with applications*.
- Yu, Q.; He, J.; Deng, X.; Shen, X.; and Chen, L.-C. 2023. Convolutions Die Hard: Open-Vocabulary Segmentation with Single Frozen Convolutional CLIP. In *NeurIPS*.
- Zamir, S. W.; Arora, A.; Gupta, A.; Khan, S.; Sun, G.; Khan, F. S.; Zhu, F.; Shao, L.; Xia, G.-S.; and Bai, X. 2019. iSAID: A Large-scale Dataset for Instance Segmentation in Aerial Images. *arXiv:1905.12886*.
- Zang, Z.; Lin, C.; Tang, C.; Wang, T.; and Lv, J. 2024. Zero-Shot Aerial Object Detection with Visual Description Regularization. *arXiv:2402.18233*.
- Zavras, A.; Michail, D.; Demir, B.; and Papoutsis, I. 2024. Mind the Modality Gap: Towards a Remote Sensing Vision-Language Model via Cross-modal Alignment.
- Zhang, R.; Fang, R.; Gao, P.; Zhang, W.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2021a. Tip-Adapter: Training-free CLIP-Adapter for Better Vision-Language Modeling. *arXiv preprint arXiv:2111.03930*.
- Zhang, R.; Hu, X.; Li, B.; Huang, S.; Deng, H.; Li, H.; Qiao, Y.; and Gao, P. 2023. Prompt, Generate, then Cache: Cascade of Foundation Models makes Strong Few-shot Learners. *arXiv preprint arXiv:2303.02151*.
- Zhang, T.; Zhang, X.; Zhu, P.; Tang, X.; Li, C.; Jiao, L.; and Zhou, H. 2021b. Semantic attention and scale complementary network for instance segmentation in remote sensing images. *IEEE Transactions on Cybernetics*.
- Zhang, X.; Han, L.; Han, L.; and Zhu, L. 2020. How well do deep learning-based methods for land cover classification and object detection perform on high resolution remote sensing imagery? *Remote Sensing*.
- Zheng, Y.; Wu, J.; Qin, Y.; Zhang, F.; and Cui, L. 2021. Zero-Shot Instance Segmentation. *arXiv:2104.06601*.
- Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract Free Dense Labels from CLIP. In *ECCV*.
- Zhou, Z.; Lei, Y.; Zhang, B.; Liu, L.; and Liu, Y. 2023. ZegCLIP: Towards adapting CLIP for zero-shot semantic segmentation. In *CVPR*.
- Zhu, X.; Zhang, R.; He, B.; Zhou, A.; Wang, D.; Zhao, B.; and Gao, P. 2023. Not All Features Matter: Enhancing Few-shot CLIP with Adaptive Prior Refinement. *arXiv:2304.01195*.