

SkySense V2: A Unified Foundation Model for Multi-modal Remote Sensing

Yingying Zhang¹ Lixiang Ru¹ Kang Wu² Lei Yu¹ Lei Liang¹ Yansheng Li² Jingdong Chen¹
¹Ant Group ²Wuhan University
 qichu.zyy@antgroup.com

Abstract

The multi-modal remote sensing foundation model (MM-RSFM) has significantly advanced various Earth observation tasks, such as urban planning, environmental monitoring, and natural disaster management. However, most existing approaches generally require the training of separate backbone networks for each data modality, leading to redundancy and inefficient parameter utilization. Moreover, prevalent pre-training methods typically apply self-supervised learning (SSL) techniques from natural images without adequately accommodating the characteristics of remote sensing (RS) images, such as the complicated semantic distribution within a single RS image. In this work, we present SkySense V2, a unified MM-RSFM that employs a single transformer backbone to handle multiple modalities. This backbone is pre-trained with a novel SSL strategy tailored to the distinct traits of RS data. In particular, SkySense V2 incorporates an innovative adaptive patch merging module and learnable modality prompt tokens to address challenges related to varying resolutions and limited feature diversity across modalities. In addition, we incorporate the mixture of experts (MoE) module to further enhance the performance of the foundation model. SkySense V2 demonstrates impressive generalization abilities through an extensive evaluation involving 16 datasets over 7 tasks, outperforming SkySense by an average of 1.8 points.

1. Introduction

Advancements in multi-modal remote sensing foundation models (MM-RSFM) have significantly enhanced the utilization of remote sensing (RS) data across numerous Earth observation (EO) applications[5, 10, 36, 65], such as environmental monitoring, agriculture management, natural disaster response, land cover (or use) mapping *etc.* Recently, Guo *et al.* introduced SkySense [20], the largest MM-RSFM to date. SkySense has demonstrated impressive generalization capabilities through extensive evaluations across many different tasks. A key factor in its success is the simultaneous training of geographically aligned multi-modal




Figure 1. SkySense V2 has achieved superior performance on 16 datasets over 7 distinct tasks compared with its predecessor SkySense and supports a board range of tasks.

and multi-temporal RS data. To accommodate the varying resolutions of different modalities, SkySense employs the Swin Transformer [31] for processing high-resolution (HR) optical images, while using the Vision Transformer (ViT) [16] for medium-resolution multi-spectral (MS) data and synthetic aperture radar (SAR) data. However, this separate backbone design results in redundancy and inefficient use of parameters. On the other hand, SkySense’s pre-training approach is primarily based on DINOv2 [42], renowned for its efficiency in self-supervised learning (SSL). Although it has demonstrated strong performance with RS images, there remains considerable room for improvement by adequately accommodating the involved semantic distribution within a single RS image.

In this paper, we introduce SkySense V2, an advanced MM-RSFM pre-trained through a novel SSL approach tailored to the distinctive traits of RS data. To be specific, SkySense V2 leverages a unified transformer backbone capable of processing data across diverse modalities and resolutions. It is pre-trained using a new SSL strategy that utilizes query-based attention to gather the similar semantic features distributed in different regions of RS images.

Designing a unified backbone for geographically aligned multi-modal data presents two primary challenges. The first is processing feature resolutions across data with varying ground sample distance (GSD) during simultaneous pre-training of multiple modalities. To address this challenge, we introduce an innovative Adaptive Patch Merging (APM) module, which is integrated after each stage of the unified backbone. This module determines whether to reduce feature resolution according to the specific feature resolution requirements of each modality. For instance, when dealing with HR optical images, the APM module activates resolution reduction at each stage. In contrast, for medium-resolution data, such as MS and SAR data, the feature resolution is preserved throughout the stages. This design enables the backbone to efficiently process data of different resolutions while adapting output feature resolutions according to actual requirements, which is crucial for the joint training and fusion of multi-modal features.

The second challenge is that fully sharing parameters across the different modalities can reduce the feature diversity. To address this, inspired by the work of [14] and [26], we introduce learnable modality prompt tokens for each modality. By interacting with these individual modal prompt tokens through an attention mechanism, the pre-trained model can capture the unique characteristics of each modality better. Compared to SkySense, our unified design significantly enhances the efficiency of parameter utilization. To be specific, SkySense uses three separate backbones, Swin-H for HR optical data, ViT-L for MS data, and ViT-L for SAR data, collectively totaling 1.26 billion parameters. In contrast, SkySense V2’s unified transformer backbone is capable of simultaneously processing all three modalities while utilizing only 665 million parameters. This efficient use of parameters not only optimizes the architecture but also allows for potential model scaling to further boost performance. Therefore, we integrate the mixture of experts module [24], a technique commonly used in various large language models [17, 45, 68] to further enhance the performance of the foundation model.

To effectively adapt SSL techniques developed for natural images to the pre-training of RSFM, it is essential to consider the distinct differences in data characteristics. Natural images typically feature a single, clear subject, such as a cat, dog, or person. In contrast, RS images comprise multiple subjects distributed across various regions of the image. For instance, as illustrated in Figure 2, one section of an RS image may show a **building**, while other areas might show a **forest**, **pond**, or **land**. Current SSL methods perform contrastive learning directly between different views using various augmented crops, which is effective for natural images due to their singular focus. However, when applied to RS images, this approach can result in semantic inaccuracies, as different views may capture different subjects. For

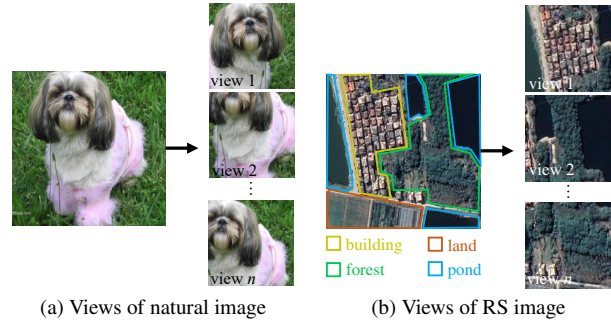


Figure 2. Comparison of different augmented views from natural and RS images in SSL. Natural image views often concentrate on a singular subject, such as a “dog”. RS image views encompass a variety of subjects, including building, forest, pond, and land.

example, in Figure 2b, view 1 of the RS image primarily contains **building**, whereas view 2 shows **forest** and **pond**.

To tackle this challenge, we propose a novel approach called Query-based Semantic Aggregation Contrastive Learning (QSACL). This method employs multiple learnable queries to perform cross-attention with features from different views, generating semantically aggregated features. We then apply contrastive learning to these aggregated features pairs derived from the same query. As demonstrated in Figure 9 from our ablation study, utilizing different queries can aggregate consistent semantic features across multiple views, which enhances the accuracy of contrastive learning.

We evaluated SkySense V2 on a diverse set of 16 datasets [8, 11–13, 15, 19, 27, 48, 51, 53, 59, 61], covering a range of task types, modalities, and spatial scales. As illustrated in Figure 1, SkySense V2 demonstrates substantial improvements in performance over its predecessor, SkySense, achieving state-of-the-art (SOTA) results across various modalities of EO tasks. The experimental results across all test scenarios highlight its competitive advantage over existing RSFMs in a wide array of EO interpretation tasks.

2. Related Work

2.1. Remote Sensing Foundation Model

RSFMs are characterized by their ability to leverage vast amounts of data through self-supervised learning (SSL) technologies, enabling them to learn robust feature representations without the need for extensive annotations. The predominant RSFMs utilize Contrastive Learning (CL) or Masked Image Modeling (MIM). Typical works in CL include RS-BYOL [25], GASSL [64], DINO-MC [58], SeCo [38], and CACo [37]. Other research efforts focus on enhancing MIM framework, such as RingMo [50], S2MAE [29], MA3E [30], SatMAE [13], and SatMAE++ [41]. Additionally, studies like CMID [40] and GFM [39] explore the intersection of CL and MIM through a self-distillation

approach. Recently, Guo *et al.* introduced a comprehensive MM-RSFM known as SkySense [20], which features a factorized multi-modal spatiotemporal encoder. This architecture facilitates independent spatial feature extraction and multi-modal temporal fusion. As the largest MM-RSFM to date, SkySense exhibits exceptional generalization capabilities across a wide range of RS datasets. However, it does face a challenge: the inefficient utilization of parameters due to the separate backbone designs for different modalities. In this work, we propose SkySense V2 to address this issue by employing a unified transformer backbone design.

2.2. Unified Framework for Multi-modal Learning

Multi-modal learning aims to train models that effectively process and relate information from various modalities. Substantial advancements have been made in this field, culminating in numerous studies focused on creating unified networks for the integration and processing of diverse modalities. The main work includes VLMO [46], Meta-Transformer [66], Uni-Perceiver [67], and UniTR [21], *etc.* There are several methods that utilize a unified backbone design for multi-modal RS data. For example, OFA-Net [63] pre-trains a single transformer backbone on a curated multi-modal dataset using MIM. Han *et al.* proposed a RSFM called msGFM[22], which utilizes a shared transformer encoder while employing different decoders for each modality. In contrast to our approach, both OFA-Net and msGFM simply share all of parameters of backbone and learn feature representations for multi-modal data separately, which results in a lack of geographical alignment and fusion between features from different modalities. AnySat [2] employs a joint embedding predictive architecture for a multi-modal model, training a single model on heterogeneous data in a SSL manner. However, the backbone design and SSL pre-training methods of our SkySense V2 differ significantly from AnySat.

3. SkySense V2

3.1. Model Architecture

3.1.1. Unified Transformer Backbone

As illustrated in Figure 3, our unified transformer backbone is a hierarchical encoder structure with four stages. In the first two stages, we employ Swin Transformer V2 Blocks (SwinV2B) [32] to incorporate essential visual priors, such as locality and translation invariance. The window-based self-attention mechanisms in SwinV2B also help reduce computational complexity compared to global self-attention, particularly given the high spatial resolution of features in these initial stages. In the last two stages, we utilize vanilla Transformer Blocks (TB) [16] for two main reasons: First, the spatial resolution of features in these stages is relatively low, making the computational

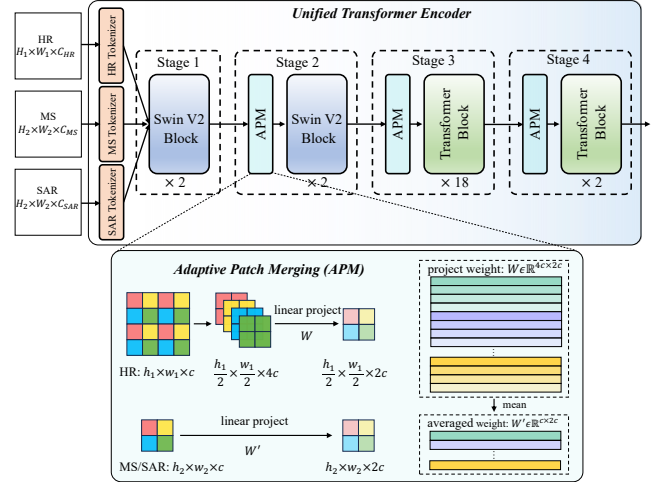


Figure 3. Overview of the unified transformer backbone model in SkySense V2: The entire backbone shares all parameters across different modalities, with the exception of distinct tokenizers.

costs of global self-attention more manageable. Second, our ablation study demonstrates that global self-attention can be complementary to window-based self-attention, enabling the model to achieve stronger representations.

Given a group of multi-modal inputs consisting of a high-resolution optical image x_{HR} , multi-spectral data x_{MS} , and synthetic aperture radar data x_{SAR} where each 'pixel' in the different modalities is naturally aligned by geo-location, we employ three distinct tokenizers to process these inputs into tokens. In each tokenizer, the input data is first divided into non-overlapping 4×4 patches. Subsequently, a linear embedding layer is applied to this raw, patched data to project it into patch tokens. Subsequently, four stages of either SwinV2B or TB, with shared parameters across modalities, are applied to these patch tokens.

Due to the varying GSD of satellite sensors across different modalities, input data corresponding to the same region at specific geo-locations during joint aligned training exhibit differing spatial resolutions. To harmonize these spatial resolutions for tokens from various modalities, we propose Adaptive Patch Merging (APM) and utilize it to selectively reduce spatial resolution at each stage, except for Stage 1. Specifically, for the tokens of optical image data with high spatial resolution, APM concatenates the features from groups of 2×2 neighboring patches and applies a linear layer to the $4c$ -dimensional concatenated features. This process reduces the number of tokens by a factor of $2 \times 2 = 4$ downsampling of spatial resolution, while the output dimension is set to $2c$. In contrast, for the tokens of MS and SAR data with lower spatial resolution, APM maintains the resolution by applying a linear projection with averaging the weights across the input dimensions. By integrating the APM module, our unified backbone can efficiently process

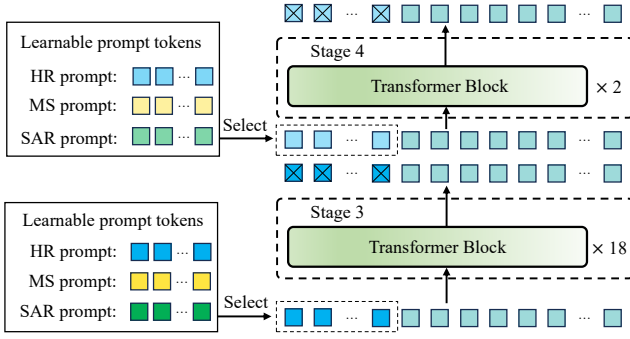


Figure 4. The illustration of modality-specific prompt tokens added to the unified backbone.

multi-modal data with varying resolutions while maintaining the spatial alignment of features across modalities

3.1.2. Modality-specific Prompt Tokens

Our unified transformer backbone utilizes fully shared parameters across different modalities. To enhance feature diversity, we introduce learnable modality-specific prompt tokens for each modality. By interacting with these individual modal prompt tokens through an attention mechanism, the pre-trained model can better capture the unique characteristics of each modality. As illustrated in Figure 4, we incorporate N learnable prompt tokens for each modality in the last two stages. We denote the input tokens of each stage as $E_i^j \in \mathbb{R}^{h_j w_j \times c_j}$, $i \in \{HR, MS, SAR\}$, $j \in \{3, 4\}$, where h_j and w_j are the height and width of the spatial resolution at stage j , while c_j refers to the dimension of the tokens. For each modality i , we insert the modality-specific learnable prompt tokens $P_i^j \in \mathbb{R}^{N \times c_j}$ into the input of the first block in j -th stage \mathcal{F}_j . At the last block of each stage, these prompt tokens are discarded. The entire process can be formulated as:

$$\begin{aligned} [E_{drop}, E_i^4] &= \mathcal{F}_3([P_i^3, E_i^3]), \\ [E_{drop}, E_i^{out}] &= \mathcal{F}_4([P_i^4, E_i^4]). \end{aligned} \quad (1)$$

Here, E_{drop} denotes the discarded tokens, and E_i^{out} represents the final output tokens of backbone. Our design of modality-specific prompts enhances feature diversity while maintaining full parameter sharing by incorporating only a small number of modality-specific parameters.

3.1.3. Scaling Up Model Capacity

Our unified backbone design across various modalities significantly enhances parameter utilization efficiency. To further boost the model’s capabilities, we integrate a Mixture of Experts (MoE) [24] into the backbone network. We choose MoE instead of simply increasing the backbone’s width and depth because MoE enables models to be pre-trained with substantially less computation by leveraging

sparse feed-forward layers (*i.e.*, experts) for individual tokens. Following common approaches [33, 60], we insert MoE modules into the last L transformer blocks, replacing the original feed-forward network (FFN) layers. Each MoE module consists of M experts (denoted as $\mathcal{E}_i(\cdot)$, $i = 1, 2, \dots, M$), which share the same structure as the FFN and operate as independent networks. For the gating network in MoE, we employ a learnable linear layer followed by a Softmax function, $\mathcal{G}(x) = \text{Softmax}(Wx)$, where W is the gating parameter. Finally, the output of the MoE is computed as a linear combination of the outputs from the selected experts, weighted by the corresponding gate values. This process is formulated as follows:

$$MOE(x) = \sum_{i \in \mathcal{T}} \mathcal{G}_i(x) \cdot \mathcal{E}_i(x), \quad (2)$$

\mathcal{T} represent the set of the top- k indices. In our SkySense V2 backbone, we set $L = 6$, $M = 8$, and $k = 1$.

3.2. Overall Pre-training Pipeline

Figure 5 provides an overview of our pre-training procedure. The pre-training framework of SkySense V2 primarily adopts the teacher-student architecture from SkySense [20], where the teacher network’s parameters are updated using an exponential moving average (EMA) [42] of the student network’s parameters. To train SkySense V2, we employ the Multi-Granularity Contrastive Learning (MGCL) loss \mathcal{L}_{MGCL} as proposed in SkySense. For each modality, MGCL utilizes contrastive loss [42] to align the representations of teacher and student networks at pixel, object and image-level granularity, respectively. Additionally, we incorporate the unsupervised Geo-Context Prototype Learning (GCPL), as validated in SkySense, to enhance the learning of complementary regional context clues, aiding downstream tasks. To further enhance dense interpretation capabilities, we introduce an auxiliary supervision loss \mathcal{L}_{ITA} to employ dense Image-Text Alignment (ITA) according to OpenStreetMap (OSM) labels¹. Detailed implementations of MGCL, GCPL, and ITA are included in the Appendix A. In following paragraph, we will focus on proposed query-based semantic aggregation contrastive learning (QSACL).

Query-based Semantic Aggregation Contrastive Learning.

QSACL utilizes learnable queries to aggregate similar semantics across different regions of images, enabling more accurate contrastive learning. Given features g_1, g_2 from two global views and l_1, l_2, \dots, l_n from n local views (where g and l denotes the fused features F_{fus} from global and local view), we employ a transformer decoder layer to perform cross-attention between m learnable queries q_1, q_2, \dots, q_m and the feature sets g and l . This process generates sets of globally and locally aggregated features set

¹<https://www.openstreetmap.org/>

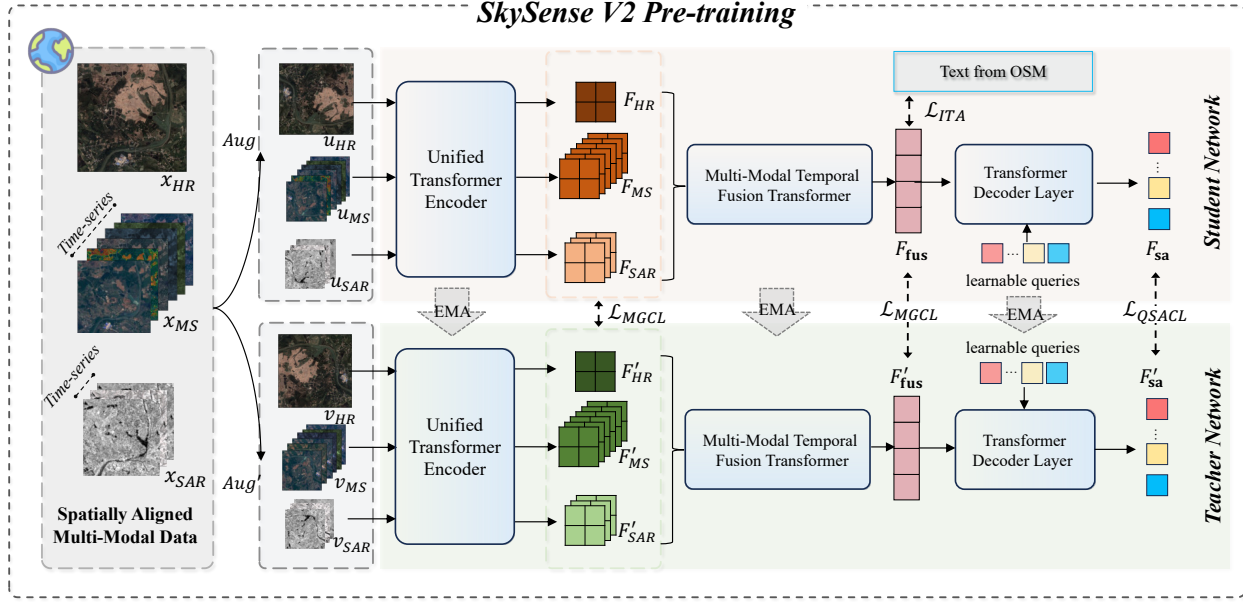


Figure 5. Overview of the SkySense V2 pre-training pipeline. The SkySense V2 employs data augmentations on multi-modal inputs, which are then fed into both the student and teacher networks. The pre-training process incorporates Multi-Granularity Contrastive Learning, Dense Image-Text Alignment, and Query-based Semantic Aggregation Contrastive Learning to effectively train the network. Once pre-training is complete, the parameters from the teacher branch are utilized for downstream applications.

z_i^g and z_i^l for each $i = 1, 2, \dots, m$. As is shown in Figure 6, we use superscript l denotes features from the teacher network, the QSACL loss \mathcal{L}_{QSACL} is computed by averaging the loss \mathcal{L}_{CL} of aggregation features between local and global views for each query:

$$\mathcal{L}_{QSACL} = \frac{1}{2m} \sum_{i=1}^m (\mathcal{L}_{CL}(z_i^g, z_i^l) + \mathcal{L}_{CL}(z_i^l, z_i^g)), \quad (3)$$

$$\mathcal{L}_{CL}(x, x') = -\mathcal{H}(x) \log(\mathcal{H}'(x')). \quad (4)$$

Here, \mathcal{H} and \mathcal{H}' denotes the learning head and corresponding EMA part as defined in [6]. The overall training objective is the weighted sum of \mathcal{L}_{MGCL} , \mathcal{L}_{ITA} , and \mathcal{L}_{QSACL} , i.e.,

$$\mathcal{L} = \lambda_1 \mathcal{L}_{MGCL} + \lambda_2 \mathcal{L}_{ITA} + \lambda_3 \mathcal{L}_{QSACL}. \quad (5)$$




Figure 6. Computation pipelines for QSACL with multi-crop augmentation (2 global and n local views).

4. Experiments

Following SkySense [20], we conducted experiments on 16 datasets across various modalities and tasks to thoroughly evaluate SkySense V2. The applications of pre-trained SkySense V2 to diverse tasks are detailed in Appendix B. Each module is engineered for combined or individual use, providing the flexibility to be frozen or fine-tuned as needed.

4.1. Pre-training Implementation

Training Datasets. We use the same training dataset of SkySense [20] to train SkySense V2. This extensive dataset comprises approximately 21 million multi-modal RS imagery sets. Each set comprises three components: a high-resolution optical RGB image of spatial dimension 2048×2048 , a sequence of Sentinel-2 (S2) multi-spectral (MS) images with 10 spectral bands and spatial dimensions of 64×64 (average sequence length of 65 images), and a sequence of Sentinel-1 (S1) synthetic aperture radar (SAR) images with 2 bands, also having spatial dimensions of 64×64 (average sequence length of 10 images). In each training iteration, we randomly sample 10 S2 images and 5 S1 images of different time series to accommodate the GPU memory constraints.

Setting of Unified Backbone. We initially set the token dimension to $C = 352$. In the first two stages, the win-

dow size for the Swin Transformer V2 Blocks (SwinV2B) is configured to 8. Each head has a query dimension of 32, and each MLP layer consists of two full-connected layers with dimensions expanding by a factor of 4 consistently across all blocks within the unified backbone. Prompt tokens are incorporated during the last two stages, with the number of prompt tokens being the same for each modality. We allocate 4 prompt tokens for stage-3 and stage-4.

Pre-training Settings. SkySense V2 is trained using a batch size of 1024, distributed across 128 H20 GPUs. The model undergoes a total of 600k iterations, utilizing the AdamW optimizer [35]. The learning rate is initially set to 2×10^{-4} and decays to 1×10^{-6} following a cosine annealing schedule [34]. Additional implementation details can be found in the Appendix C.

4.2. Single-Modal Tasks

We evaluate SkySense V2 on four common single-modal RS interpretation tasks: scene classification, object detection, semantic segmentation, and change detection. The experiments are conducted via fine-tuning a pre-trained SkySense V2 backbone along with a task-specific head. Further implementation details can be found in the Appendix D.

Scene Classification. In Table 1, we report the performance of scene classification on four widely used datasets: AID [61] and RESISC-45 [11], both featuring static RGB optical images; BEN-S2 [48], which includes static MS images; and fMoW-S2 [13], which contains temporal MS images. For each dataset, the performance is evaluated using different training ratios (TR) following the methodology [13, 38]. In our experiments, we employ a simple linear classifier as the classification head. For multi-label datasets, we use the mean average precision (mAP) as the evaluation metric, whereas for single-label datasets, we report the overall accuracy (OA). As illustrated in Table 1, SkySense V2 consistently outperforms previous RSFMs across various datasets and training configurations, particularly in low TR settings. This demonstrates that SkySense V2 exhibits a stronger representation capability compared to other RSFMs, including its predecessor, SkySense [20].

Semantic Segmentation. In Table 2a, we report the segmentation results from four representative semantic segmentation datasets: iSAID [59] and Potsdam [47] for high-resolution optical image segmentation, and Dyna.-Pla [53] and Dyna.-S2 [53] for multi-spectral and multi-temporal image segmentation. Following previous work, we use the mean F1-score (mF1) as the evaluation metric for the Potsdam dataset. While for the other datasets, we report the mean intersection over union (mIoU). All segmentation experiments employ the UperNet [62] as the decoder head.

Models	Single-label		Multi-label	Temporal
	AID	RESISC-45	BEN-S2	fMoW-S2
	(TR=20%/50%)	(TR=10%/20%)	(TR=10%/100%)	(TR=100%)
	OA	OA	mAP	Top-1/5 Acc
GASSL [3]	93.55/95.92	90.86/93.06	79.24/87.40	50.69/77.99
SeCo [38]	93.47/95.99	89.64/92.91	82.62/87.81	51.65/77.40
SatMAE [13]	95.02/96.94	91.72/94.10	86.18/89.50	63.84/-
RingMo [50]	96.90/98.34	94.25/95.67	-	-
RVSA [55]	97.03/98.50	93.93/95.69	-	-
DINO-MC [58]	-	-	84.20/88.75	60.16/83.49
TOV [52]	95.16/97.09	90.97/93.79	-	-
SSL4EO [57]	91.06/94.74	87.60/91.27	87.10/91.80	51.70/76.77
CMID [40]	96.11/97.79	94.05/95.53	-	-
CACo [37]	90.88/95.05	88.28/91.94	81.30/87.00	50.72/76.31
CROMA [18]	-	-	88.29/-	63.59/-
SatLas [4]	94.96/97.38	92.16/94.70	82.80/88.37	57.95/79.00
GFM [39]	95.47/97.09	92.73/94.64	86.30/-	-
Scale-MAE [43]	96.44/97.58	92.63/95.04	-	-
MA3E [30]	-/99.04	-/96.23	-	-
SatMAE++ [41]	-	-/97.48	85.11/-	63.23/-
SkySense [20]	97.68/98.60	94.85/96.32	88.67/92.09	64.38/87.27
SkySense V2	98.34/99.05	96.42/97.24	89.13/93.78	66.65/89.32

Table 1. Scene classification results. - means the task is not supported or the value is unavailable in the paper.

As illustrated in Table 2a, our SkySense V2 achieve higher performance than recent RSFMs across all four datasets. Specifically, compared to the previous SOTA method SkySense [20], SkySense V2 achieves 1.5% higher performance on average.

Horizontal & Oriented Object Detection. We utilize the DIOR dataset [27] to evaluate SkySense V2’s performance in horizontal object detection, while utilizing the DIOR-R [12] and FAIR1M [51] datasets for oriented object detection. These datasets only consist of RGB RS imagery. Consistent with previous research [20], we adopt Faster R-CNN [44] and Oriented R-CNN [28] as the basic detectors for horizontal and oriented object detection, respectively. As reported in Table 2b, our SkySense V2 outperforms previous RSFMs by a notable margin. Specifically, SkySense V2 surpass the previous SOTA, SkySense [20], by an average of 1.1% mAP.

Change Detection. For change detection, we conducted experiments on LEVIR-CD [8], OSCD [15] and Dyna.-S2 [53] datasets. Since LEVIR-CD and OSCD datasets only contain binary change detection results, we use the basic change detector in [9] and report the F1 score as the evaluation metric. While Dyna.-S2 focuses on semantic change detection, we use UperNet [62] as the segmentation network and calculate the semantic change detection score (SCS) on validation and test set following [20]. As presented in Table 2c, SkySense V2 consistently achieves higher performance than the previous SOTA method SkySense on all 3 datasets, and surpasses SkySense by 2.7% on average. Particularly, on the OSCD dataset, our SkySense V2 remark-

Models	Publication	Dyna.-Pla.		iSAID	Potsdam	Dyna.-S2		Models	Horizontal		Oriented		Models	LEVIR-CD	OSCD	Dyna.-S2
		mIoU	mIoU	mF1	mIoU	DIOR	DIOR-R		FAIR1M	F1	F1	SCS				
GASSL [3]	ICCV'21	34.0/40.8	65.95	91.27	28.1/41.0	GASSL [3]	67.40	65.65	48.15	GASSL [3]	78.19	46.26	13.6/16.7			
SeCo [38]	ICCV'21	-	57.20	89.03	29.4/39.8	SeCo [38]	70.89	65.66	46.55	SeCo [38]	90.14	47.67	13.9/16.0			
SatMAE [13]	NIPS'22	32.8/39.9	62.97	90.63	30.1/38.7	SatMAE [13]	75.90	-	46.21	SatMAE [13]	87.65	52.76	14.8/16.2			
RingMo [50]	TGRS'22	-	67.20	91.27	-	RingMo [50]	73.22	71.05	47.04	RingMo [50]	91.86	-	-			
RVSA [55]	TGRS'22	34.3/44.4	64.49	-	-	RVSA [55]	-	73.62	-	RVSA [55]	90.86	-	-			
BFM [7]	Arxiv'23	-	-	92.12	-	BFM [7]	70.16	66.33	49.62	SpectralGPT [23]	-	54.29	-			
TOV [52]	JSTARS'23	32.1/37.8	66.24	92.03	-	TOV [52]	64.82	61.23	49.37	MATTER [1]	-	59.37	-			
SSL4EO [57]	GRSM'23	35.3/42.1	64.01	91.54	31.8/42.7	SSL4EO [57]	75.11	66.37	50.58	DINO-MC [58]	-	52.70	14.5/15.6			
CMID [40]	TGRS'23	36.4/43.5	66.21	91.86	-	CMID [40]	66.91	64.10	47.83	SSL4EO [57]	89.05	35.08	12.3/17.5			
CACo [37]	CVPR'23	35.4/42.7	64.32	91.35	30.2/42.5	CACo [37]	74.10	67.59	46.19	CMID [40]	91.72	-	-			
SAMRS [56]	NIPS'23	-	66.26	91.43	-	SatLas [4]	72.84	67.67	49.69	CACo [37]	81.04	52.11	15.3/15.8			
SatLas [4]	ICCV'23	37.4/40.7	68.71	91.28	31.9/43.5	GFM [39]	73.81	66.47	48.31	SatLas [4]	90.62	-	13.3/17.8			
GFM [39]	ICCV'23	36.7/45.6	66.62	91.85	-	Scale-MAE [43]	-	71.82	-	GFM [39]	91.73	59.82	-			
Scale-MAE [43]	ICCV'23	34.0/41.7	65.77	91.54	-	MA3E [30]	78.73	74.27	54.57	Scale-MAE [43]	92.07	-	-			
MA3E [30]	ECCV'24	-	64.06	91.50	-	SkySense [20]	-	-	-	SkySense [20]	92.58	60.06	15.4/18.0			
SkySense [20]	CVPR'24	39.7/46.5	70.91	93.99	33.1/46.2	SkySense V2	79.50	75.29	55.96	SkySense V2	94.83	65.29	16.0/20.7			

(a) Semantic segmentation results.

(b) Object detection results.

(c) Change detection results.

Table 2. Results of semantic segmentation, object detection and change detection.

Task & Dataset	Data Source	SkySense[20]	SkySense V2
(a) Multi-modal Seg: Dyna.-MM	(i) Planet.	46.5	47.6
	(ii) S2	46.2	47.5
	(iii) Planet. + S2	47.3	48.7
	(iv) Planet. + S2 + S1	47.7	48.9
(b) Multi-modal Seg: PASTIS-MM	(i) S2	73.5	75.0
	(ii) S2-Ts	84.6	85.5
	(iii) S2-Ts + S1-Ts	84.8	85.6
	(iv) S2-Ts + GEP	85.8	86.7
(c) Multi-modal Cls: BEN-MM	(i) S1	86.2	86.5
	(ii) S2 + S1	92.2	93.8

Table 3. Fine-tuning results on multi-modal tasks.

ably outperform SkySense by 5.2% on F1-score.

4.3. Multi-modal Tasks

In this section, we evaluate SkySense V2 on two representative multi-modal tasks, *i.e.*, Multi-modal Segmentation and Multi-modal Scene Classification, to demonstrate its generalization capability across data from various modalities.

Multi-modal Segmentation. In Table 3 (a) and (b), we present the performance results of Dyna.-MM [53] and PASTIS-MM [19], respectively. Dyna.-MM contains high-resolution optical imagery from PlanetFusion (Planet.), multi-spectral imagery from Sentinel-2 (S2), and SAR imagery from Sentinel-1 (S1). For the experiments conducted on Dyna.-MM, we employ a basic UperNet [62] as the decoder head and utilize mIoU as the evaluation metric. As depicted in Table 3 (a), SkySense V2 consistently outperforms the SkySense across various modalities, demonstrating a stronger representation ability. Furthermore, integrating multiple modalities yields higher performance compared to using a single modality. This results confirms that the proposed SkySense V2 can effectively extract representations from diverse modalities using a unified backbone.

PASTIS-MM is a crop mapping dataset that contains high-resolution optical imagery from Google Earth Pro (GEP), multi-temporal multi-spectral imagery from Sentinel-2 (S2-Ts), and multi-spectral SAR imagery from Sentinel-1 (S1-Ts). We employ a straightforward FCN head to decode the semantic segmentation results and report the overall accuracy as the evaluation metric in Table 3 (b). The results demonstrate that SkySense V2 outperforms the previous SOTA method SkySense by an average of 1.0%. Additionally, it is worth noting that S2-Ts can remarkably improve the segmentation accuracy compared to S2, underscoring the importance of temporal information for crop mapping.

Multi-modal Scene Classification. We conduct multi-modal scene classification experiments on the BEN-MM dataset [49]. As is shown in Table 3 (c), the proposed SkySense V2 surpasses SkySense, demonstrating a stronger representation ability. Furthermore, by integrating S1 and S2 imagery, SkySense V2 achieves higher improvements, highlighting its enhanced ability to extract superior representations from different modalities compared to SkySense.

4.4. Ablation and Discussion

Ablation Study of Components. To understand the contributions of each component, we conduct a k-nearest neighbors (k-NN) evaluation on the AID and RESISC-45 datasets. Figure 7 illustrates the performance of different component combinations across various pre-training iterations. Notably, the unified backbone design significantly accelerates representation learning. This is primarily due to shared parameters across different modalities, which allow gradients to aggregate, thereby speeding up convergence. Moreover, the unified design allows the backbone to be trained with data from different modalities, thereby enhancing the model’s generalization capabilities (see Ap-




Figure 7. Ablation study using k-NN classification was conducted on the AID and RESISC-45 datasets. In this context, the abbreviations have the following meanings: GA refers to replacing window-based self-attention with global self-attention in the last two stages; MsP involves the addition of modality-specific prompt tokens; MoE indicates scaling up the model through a mixture of experts approach; and QSACL means adding additional query-based semantic aggregation contrastive learning.

pendix F.3 for more details). The results also indicate that global self-attention complements window-based self-attention, facilitating improved representations. Additionally, modality-specific prompt tokens and MoE scaling enhance representation learning by increasing feature diversity and model capacity, respectively. QSACL enhances model performance by improving the semantic precision of features in contrastive learning.

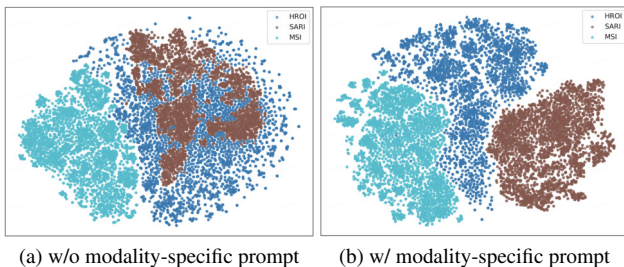


Figure 8. t-SNE visualization of the output feature representation from the final block of the unified transformer backbone. HROI, SARI, and MSI represent the features from high-resolution optical imagery, synthetic aperture radar imagery, and multi-spectral imagery, respectively.

How Do Modality-Specific Prompt Tokens Facilitate Pre-training? To investigate the impact of modality-specific prompt (MsP) tokens on the multi-modal pre-training process within a unified backbone, we employed t-SNE [54] for visualizing the features corresponding to different modalities. These features are extracted from the last block of the unified transformer backbone, and the comparative results are presented in Figure 8. Without the use of MsP, the features of high-resolution optical imagery (HROI) and synthetic aperture radar imagery (SARI) appear closely clustered, leading to overlapping feature distri-

butions. However, with the inclusion of MsP, the features of HROI and SARI become distinctly separable. These visualization outcomes clearly demonstrate that MsP effectively enhances feature diversity and imparts modality-specific characteristics.




Figure 9. Visualization of cross-attention weights for features patch corresponding to different queries in QSACL. Queries can effectively aggregate features with specific semantics.

What Types of Features Does Query-based Semantic Aggregation Contrastive Learning Capture? To provide a clearer understanding of our proposed QSACL, we visualize the attention weights assigned to different features during their interaction with certain query, as illustrated in Figure 9. During pre-training, we use two global augmentation crops and six local augmentation crops. Here, we present two distinct queries, each focusing on different semantic features of image patches. For example, one query focuses on features with building semantics (left part of the first row in Figure 9), while another focuses on features with land semantics (right part of the first row in Figure 9). The attention weights of different queries aggregate features from each patch into distinct feature representations. Subsequently, we perform contrastive learning between aggregated features from the same query across teacher and student models. This approach significantly enhances semantic accuracy compared to traditional contrastive learning applied to entire image features.

5. Conclusion

In this paper, we introduce SkySense V2, a MM-RSFM that utilizes a unified backbone to accommodate various modalities. This unified approach enhances parameter utilization efficiency and improves the model’s generalization ability. SkySense V2 is pre-trained with an innovative QSACL strategy, specifically designed to leverage the unique characteristics of RS images. As a result, SkySense V2 significantly boosts performance while maintaining the multi-modal flexibility benefits of its predecessor, SkySense. Looking ahead, our future work will focus on integrating the language modality and incorporating a geographical knowledge graph to develop a more powerful and versatile MM-RSFM.

References

- [1] Peri Akiva, Matthew Purri, and Matthew Leotta. Self-supervised material and texture representation learning for remote sensing tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8203–8215, 2022. 7
- [2] Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. AnySat: An Earth observation model for any resolutions, scales, and modalities. *arXiv preprint arXiv:2412.14123*, 2024. 3
- [3] Kumar Ayush, Burak Uzket, Chenlin Meng, Kumar Tanmay, Marshall Burke, David Lobell, and Stefano Ermon. Geography-aware self-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10181–10190, 2021. 6, 7
- [4] Favyen Bastani, Piper Wolters, Ritwik Gupta, Joe Ferdinando, and Aniruddha Kembhavi. Satlaspretrain: A large-scale dataset for remote sensing image understanding. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16772–16782, 2023. 6, 7
- [5] Yinxia Cao, Xin Huang, and Qihao Weng. A multi-scale weakly supervised learning method with adaptive on-line noise correction for high-resolution change detection of built-up areas. *Remote Sensing of Environment*, 297:113779, 2023. 1
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 5
- [7] Keumgang Cha, Junghoon Seo, and Taekyung Lee. A billion-scale foundation model for remote sensing images. *arXiv preprint arXiv:2304.05215*, 2023. 7
- [8] Hao Chen and Zhenwei Shi. A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. *Remote Sensing*, 12(10):1662, 2020. 2, 6
- [9] Hao Chen, Zipeng Qi, and Zhenwei Shi. Remote sensing image change detection with transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–14, 2021. 6
- [10] Gong Cheng and Junwei Han. A survey on object detection in optical remote sensing images. *ISPRS journal of photogrammetry and remote sensing*, 117:11–28, 2016. 1
- [11] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 2, 6
- [12] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2022. 6
- [13] Yezhen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David Lobell, and Stefano Ermon. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems*, 35:197–211, 2022. 2, 6, 7
- [14] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [15] Rodrigo Caye Daudt, Bertr Le Saux, Alexandre Boulch, and Yann Gousseau. Urban change detection for multispectral earth observation using convolutional neural networks. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 2115–2118. Ieee, 2018. 2, 6
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1, 3
- [17] William Fedus, Barret Zoph, and Noam M. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *ArXiv*, abs/2101.03961, 2021. 2
- [18] Anthony Fuller, Koreen Millard, and James R. Green. Croma: Remote sensing representations with contrastive radar-optical masked autoencoders. *Advances in Neural Information Processing Systems*, 2023. 6
- [19] Vivien Sainte Fare Garnot, Loic Landrieu, and Nesrine Chehata. Multi-modal temporal attention models for crop mapping from satellite time series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 187:294–305, 2022. 2, 7
- [20] Xin Guo, Jiangwei Lao, Bo Dang, Yingying Zhang, Lei Yu, Lixiang Ru, Liheng Zhong, Ziyuan Huang, Kang Wu, Dingxiang Hu, Huimei He, Jian Wang, Jingdong Chen, Ming Yang, Yongjun Zhang, and Yansheng Li. Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27672–27683, 2024. 1, 3, 4, 5, 6, 7
- [21] Wang Haiyang, Tang Hao, Shi Shaoshuai, Li Aoxue, Li Zhenguo, Schiele Bernt, and Liwei Wang. Unitr: A unified and efficient multi-modal transformer for bird’s-eye-view representation. In *ICCV*, 2023. 3
- [22] Boran Han, Shuai Zhang, Xingjian Shi, and Markus Reichstein. Bridging remote sensors with multisensor geospatial foundation models. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27852–27862, 2024. 3
- [23] Danfeng Hong, Bing Zhang, Xuyang Li, Yuxuan Li, Chenyu Li, Jing Yao, Pedram Ghamisi, Naoto Yokoya, Hao Li, Xiuping Jia, Antonio Plaza, Paolo Gamba, Jon Atli Benediktsson, and Jocelyn Chanussot. Spectralgpt: Spectral remote sensing foundation model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (8):5227–5244, 2024. 7
- [24] Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3:79–87, 1991. 2, 4
- [25] Pallavi Jain, Bianca Schoen-Phelan, and Robert J. Ross. Self-supervised learning for invariant representations from

- multi-spectral and sar images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15: 7797–7808, 2022. 2
- [26] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [27] 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, 6
- [28] Wentong Li, Yijie Chen, Kaixuan Hu, and Jianke Zhu. Oriented reppoints for aerial object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1829–1838, 2022. 6
- [29] Xuyang Li, Danfeng Hong, and Jocelyn Chanussot. S2mae: A spatial-spectral pretraining foundation model for spectral remote sensing data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24088–24097, 2024. 2
- [30] Zhihao Li, Biao Hou, Siteng Ma, Zitong Wu, Xianpeng Guo, Bo Ren, and Licheng Jiao. Masked angle-aware autoencoder for remote sensing images. In *European Conference on Computer Vision*, pages 260–278. Springer, 2024. 2, 6, 7
- [31] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 1
- [32] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12009–12019, 2022. 3
- [33] Zhili Liu, Kai Chen, Jianhua Han, Lanqing Hong, Hang Xu, Zhenguo Li, and James Tin-Yau Kwok. Task-customized masked autoencoder via mixture of cluster-conditional experts. *ArXiv*, abs/2402.05382, 2024. 4
- [34] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv: Learning*, 2016. 6
- [35] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017. 6
- [36] ZhiYong Lv, HaiTao Huang, Xinghua Li, MingHua Zhao, Jon Atli Benediktsson, WeiWei Sun, and Nicola Falco. Land cover change detection with heterogeneous remote sensing images: Review, progress, and perspective. *Proceedings of the IEEE*, 2022. 1
- [37] Utkarsh Mall, Bharath Hariharan, and Kavita Bala. Change-aware sampling and contrastive learning for satellite images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5261–5270, 2023. 2, 6, 7
- [38] Oscar Manas, Alexandre Lacoste, Xavier Giró-i Nieto, David Vazquez, and Pau Rodriguez. Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9414–9423, 2021. 2, 6, 7
- [39] Matías Mendieta, Boran Han, Xingjian Shi, Yi Zhu, Chen Chen, and Mu Li. Towards geospatial foundation models via continual pretraining. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16806–16816, 2023. 2, 6, 7
- [40] Dilxat Muhtar, Xueliang Zhang, Pengfeng Xiao, Zhenshi Li, and Feng Gu. Cmid: A unified self-supervised learning framework for remote sensing image understanding. *IEEE Transactions on Geoscience and Remote Sensing*, 2023. 2, 6, 7
- [41] Mubashir Noman, Muzammal Naseer, Hisham Cholakkal, Rao Muhammad Anwer, Salman Khan, and Fahad Shahbaz Khan. Rethinking transformers pre-training for multi-spectral satellite imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27811–27819, 2024. 2, 6
- [42] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1, 4
- [43] Colorado J Reed, Ritwik Gupta, Shufan Li, Sarah Brockman, Christopher Funk, Brian Clipp, Kurt Keutzer, Salvatore Candido, Matt Uyttendaele, and Trevor Darrell. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4088–4099, 2023. 6, 7
- [44] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 6
- [45] Sheng Shen, Le Hou, Yan-Quan Zhou, Nan Du, S. Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, Tu Vu, Yuexin Wu, Wuyang Chen, Albert Webson, Yunxuan Li, Vincent Zhao, Hongkun Yu, Kurt Keutzer, Trevor Darrell, and Denny Zhou. Mixture-of-experts meets instruction tuning: A winning combination for large language models. In *International Conference on Learning Representations*, 2023. 2
- [46] Sheng Shen, Le Hou, Yanqi Zhou, Nan Du, Shayne Longpre, Jason Wei, Hyung Won Chung, Barret Zoph, William Fedus, Xinyun Chen, Tu Vu, Yuexin Wu, Wuyang Chen, Albert Webson, Yunxuan Li, Vincent Y Zhao, Hongkun Yu, Kurt Keutzer, Trevor Darrell, and Denny Zhou. Mixture-of-experts meets instruction tuning: A winning combination for large language models. In *The Twelfth International Conference on Learning Representations*, 2024. 3
- [47] Jamie Sherrah. Fully convolutional networks for dense semantic labelling of high-resolution aerial imagery. *arXiv preprint arXiv:1606.02585*, 2016. 6
- [48] Gencer Sumbul, Jian Kang, Tristan Kreuziger, Filipe Marcelino, Hugo Costa, Pedro Benevides, Mario Caetano, and Begüm Demir. Bigearthnet: A large-scale benchmark archive for remote sensing image understanding. In *IEEE International Geoscience and Remote Sensing Symposium*, pages 5901–5904, 2019. 2, 6

- [49] Gencer Sumbul, Arne de Wall, Tristan Kreuziger, Filipe Marcelino, Hugo Costa, Pedro Benevides, Mario Caetano, Begüm Demir, and Volkerl Mark. BigEarthNet-MM: A large-scale, multimodal, multilabel benchmark archive for remote sensing image classification and retrieval. *IEEE Geoscience and Remote Sensing Magazine*, 9(3):174–180, 2021. [7](#)
- [50] Xian Sun, Peijin Wang, Wanxuan Lu, Zicong Zhu, Xiaonan Lu, Qibin He, Junxi Li, Xuee Rong, Zhujun Yang, Hao Chang, et al. Ringmo: A remote sensing foundation model with masked image modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 2022. [2](#), [6](#), [7](#)
- [51] 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](#), [6](#)
- [52] Chao Tao, Ji Qi, Guo Zhang, Qing Zhu, Weipeng Lu, and Haifeng Li. Tov: The original vision model for optical remote sensing image understanding via self-supervised learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2023. [6](#), [7](#)
- [53] Aysim Toker, Lukas Kondmann, Mark Weber, Marvin Eisenberger, Andrés Camero, Jingliang Hu, Ariadna Pregel Hoderlein, Çağlar Şenaras, Timothy Davis, Daniel Cremers, et al. Dynamicearthnet: Daily multi-spectral satellite dataset for semantic change segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21158–21167, 2022. [2](#), [6](#), [7](#)
- [54] Laurens van der Maaten and Geoffrey E. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9: 2579–2605, 2008. [8](#)
- [55] Di Wang, Qiming Zhang, Yufei Xu, Jing Zhang, Bo Du, Dacheng Tao, and Liangpei Zhang. Advancing plain vision transformer toward remote sensing foundation model. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–15, 2022. [6](#), [7](#)
- [56] Di Wang, Jing Zhang, Bo Du, Dacheng Tao, and Liangpei Zhang. Scaling-up remote sensing segmentation dataset with segment anything model. *Advances in Neural Information Processing Systems*, 2023. [7](#)
- [57] Yi Wang, Nassim Ait Ali Braham, Zhitong Xiong, Chenying Liu, Conrad M Albrecht, and Xiao Xiang Zhu. Ssl4eo-s12: A large-scale multi-modal, multi-temporal dataset for self-supervised learning in earth observation. *IEEE Geoscience and Remote Sensing Magazine*, 11(3):98–106, 2023. [6](#), [7](#)
- [58] Xinye Wanyan, Sachith Seneviratne, Shuchang Shen, and Michael Kirley. Dino-mc: Self-supervised contrastive learning for remote sensing imagery with multi-sized local crops. *arXiv preprint arXiv:2303.06670*, 2023. [2](#), [6](#), [7](#)
- [59] Syed Waqas Zamir, Aditya Arora, Akshita Gupta, Salman Khan, Guolei Sun, Fahad Shahbaz Khan, Fan Zhu, Ling Shao, Gui-Song Xia, and Xiang Bai. isaid: A large-scale dataset for instance segmentation in aerial images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 28–37, 2019. [2](#), [6](#)
- [60] Lemeng Wu, Mengchen Liu, Yinpeng Chen, Dongdong Chen, Xiyang Dai, and Lu Yuan. Residual mixture of experts. *ArXiv*, abs/2204.09636, 2022. [4](#)
- [61] Gui-Song Xia, Jingwen Hu, Fan Hu, Baoguang Shi, Xiang Bai, Yanfei Zhong, Liangpei Zhang, and Xiaoqiang Lu. Aid: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55(7):3965–3981, 2017. [2](#), [6](#)
- [62] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European conference on computer vision (ECCV)*, pages 418–434, 2018. [6](#), [7](#)
- [63] Zhitong Xiong, Yi Wang, Fahong Zhang, and Xiao Xiang Zhu. One for all: Toward unified foundation models for earth vision. *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, pages 2734–2738, 2024. [3](#)
- [64] Longqi Yang, Liangliang Zhang, and Wenjing Yang. Graph adversarial self-supervised learning. In *Advances in Neural Information Processing Systems*, 2021. [2](#)
- [65] Xiaohui Yuan, Jianfang Shi, and Lichuan Gu. A review of deep learning methods for semantic segmentation of remote sensing imagery. *Expert Systems with Applications*, 169: 114417, 2021. [1](#)
- [66] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Meta-transformer: A unified framework for multimodal learning. *arXiv preprint arXiv:2307.10802*, 2023. [3](#)
- [67] Xizhou Zhu, Jinguo Zhu, Hao Li, Xiaoshi Wu, Xiaogang Wang, Hongsheng Li, Xiaohua Wang, and Jifeng Dai. Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16783–16794, 2021. [3](#)
- [68] Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam M. Shazeer, and William Fedus. Stmoe: Designing stable and transferable sparse expert models. 2022. [2](#)