

SkySense: A Multi-Modal Remote Sensing Foundation Model Towards Universal Interpretation for Earth Observation Imagery

Supplementary Material

1. Overview

We provide the following materials to supplement our paper and divide them into 12 sections.

- To demonstrate the exceptional generalization capability of SkySense’s pre-trained features, we assess the model’s performance on multiple downstream datasets using frozen backbone features, as detailed in Sec. 2. Additionally, we present results from datasets acquired using different satellite sensors compared to our pre-training data in the same section, Sec. 2.
- In Sec. 3, we execute a comparative analysis to assess the convergence rates and verify the efficacy of the SkySense in leveraging the features acquired through pre-training.
- In Sec. 4, we conduct a comparison on the SkySense with different numbers of parameters to explore the influence of model scale.
- We compare our model with randomly initialized counterpart and the representative Vision Foundation Model, DINOv2, as described in Sec. 5.
- We perform comparative experiments to establish the superior performance of the fundamental unit of SkySense, Factorized Spatiotemporal Encoder in Sec. 6.
- In Sec. 7, we illustrate the feature visualization results to further analyze the effectiveness of Cross-Modal Alignment, providing the evidence for improved multi-modal feature fusion. Moreover, we provide the experiment result on pre-training with MAE to complement our ablation study, validating our design choice for NOT including MAE.
- Sec. 8 presents qualitative results showcasing a range of visual comparisons, thereby providing further evidence to support the superiority of our method.
- In Sec. 9, we introduce the details of our pre-training dataset construction and showcase illustrative remote sensing tile examples for better understanding.
- In Sec. 10, we elaborate on the details of our SkySense pre-training implementation and its pre-training cost.
- The exposition of the downstream datasets and the corresponding fine-tuning implementation settings is elucidated in Sec. 11.
- Finally, we illustrate the procedure on how to use our SkySense pre-trained weights, and provide a simple scene classification example in Sec. 12.

2. Experimental results of the frozen backbone tuning and various satellite sensors

To validate the superiority of the features learned through our pre-training, we carry out experiments using **frozen** backbone features. Specifically, we tune task-specific heads while keeping the backbone parameters fixed for three downstream tasks: scene classification on the AID dataset [48], object detection on the DIOR dataset [17], and semantic segmentation on the iSAID dataset [46]. The heads chosen for these tasks are the same as our paper’s main experiment (the configuration outlined in Sec. 11). We compare SkySense with four recent works, *i.e.*, CMID [27], SatLas [1], GFM [26], and Scale-MAE [29], the results are shown in Tab. 1. On the AID dataset, SkySense shows impressive advantages over other methods. Especially on the setting of less training data (*i.e.*, 20% data for training and the rest for testing), SkySense achieves notable improvements in overall accuracy (OA) compared to other approaches. Specifically, compared to SatLas, SkySense achieves a 28.09% increase in OA. Similarly, when compared to Scale-MAE, SkySense surpasses it by 17.64%. Furthermore, SkySense exhibits a significant OA improvement of 14.65% over GFM and a noteworthy enhancement of 6.27% over CMID. The results obtained from the DIOR and iSAID datasets further confirm that SkySense outperforms other methods. Remarkably, our method shows a substantial average improvement of 9.69% over other methods when evaluated on the iSAID dataset. The aforementioned experimental results provide compelling evidence that our SkySense achieves superior performance with the fixed backbone. This finding reinforces the notion that our model has successfully acquired more generalized features through pre-training.

To further validate SkySense’s generalizability across multiple satellite sensors other than the data used for pre-training, we conduct experiments on three additional datasets from Gaofen and Landsat satellites, as shown in Tab. 2. SkySense greatly outperforms the other RSFMs.

3. Comparison of convergence rates

The convergence rate in downstream tasks serves as a pivotal metric in comprehensively evaluating a foundational model. In essence, a well-learned and robust feature representation during the pre-training phase facilitates swift convergence and enhances the model’s overall efficacy in downstream tasks [3].

To facilitate a comprehensive comparison between the

Model	Publication	AID	DIOR	iSAID
		OA (TR=20% / 50%)	mAP ₅₀	mIoU
CMID [27]	TGRS'23	87.80/90.92	66.08	59.40
SatLas [1]	ICCV'23	65.98/75.02	60.46	56.03
GFM [26]	ICCV'23	79.42/87.37	67.34	60.86
Scale-MAE [29]	ICCV'23	76.43/87.81	70.55	46.53
SkySense	-	94.07/95.85	72.54	65.40

Table 1. Frozen backbone tuning results for downstream tasks.

Dataset	Sensor	Previous Best RSFM	SkySense
Five-Billion-Pixels[40]	Gaofen-2	69.31 (Scale-MAE)	74.46
SPARCS[16]	Landsat-8	66.84 (GFM)	72.88
AIR-PolSAR-Seg[44]	Gaofen-3 (SAR)	53.90 (CROMA)	56.04

Table 2. Results on datasets built from various sensors. (mIoU)

SkySense and other Remote Sensing Foundation Models (RSFMs), we conduct a comparison of their convergence rates in different tasks. Specifically, we perform comparative experiments between SkySense and recent approaches, such as Scale-MAE, Satlas, CMID and GFM, in three downstream tasks: scene classification, object detection, and semantic segmentation. The convergence curves of these models are depicted in Fig. 1. The results indicate that under the same experimental settings, SkySense exhibits the fastest convergence rate in all three downstream tasks. Particularly noteworthy is its performance in the scene classification task on the AID dataset (TR=1%), where SkySense achieves desirable results with only 20 epochs. Other models require at least 140 epochs to converge to a stable but lower result. This remarkable advantage in convergence rates serves as the evidence that SkySense has successfully captured and encoded valuable information in its feature representations through effective pre-training, enabling it to rapidly adapt and generalize to downstream tasks.

4. Experimental results of SkySense with fewer parameters

Swin Transformer, with its extensive design, offers advanced modeling capabilities, allowing for strong generalization across various tasks and datasets [19]. In this paper, we employ its huge version (Swin-H) as the spatial feature extractor for HSROI in our SkySense. To validate the impact of parameter size on model performance, we conduct experiments by replacing the Swin-H (654M) with Swin-L (197M), a variant with a smaller parameter size. We perform experiments on three representative tasks (*i.e.*, scene classification on the RESISC-45 dataset [6], object detection on the DIOR dataset [17], and semantic segmentation

on the Postdam dataset [33]), as shown in Tab. 3. Comparing with Swin-H, the adoption of Swin-L results in a reduction of the model’s parameter count by 69.8%. The experimental results of downstream tasks show a slight decrease in performance when using Swin-L instead of Swin-H, thereby substantiating the effectiveness and necessity of employing the model with a larger parameter size, *i.e.*, Swin-H. We further compare our method equipping Swin-L with the two representative RSFMs: SatMAE [9] and Scale-MAE [29]. Importantly, the SkySense with Swin-L exhibits a significantly smaller parameter size (197M) compared to SatMAE (307M) and Scale-MAE (307M), while demonstrating a remarkably better performance in three downstream tasks. In particular, on the RESISC-45 dataset with TR=10%, SkySense with Swin-L outperforms SatMAE by 2.62% and Scale-MAE by 1.71% in OA.

Achieving the better performance with the smaller parameter size, our method demonstrates that its exceptional effectiveness is not simply scaling up model parameters. Instead, factors such as innovative network architecture design and advanced pre-training strategies also play vital roles. These factors, beyond the parameter size, significantly contribute to the exceptional performance of SkySense.

5. Comparison with random initialization and Vision Foundation Model

In this section, we employ both SkySense pre-trained weights and randomly initialized weights to fine-tune the same networks on 5 datasets over 4 different tasks. These tasks encompass scene classification using the AID dataset [48], object detection employing the DIOR dataset [17], semantic segmentation utilizing the Dyna.-S2 dataset [39] and

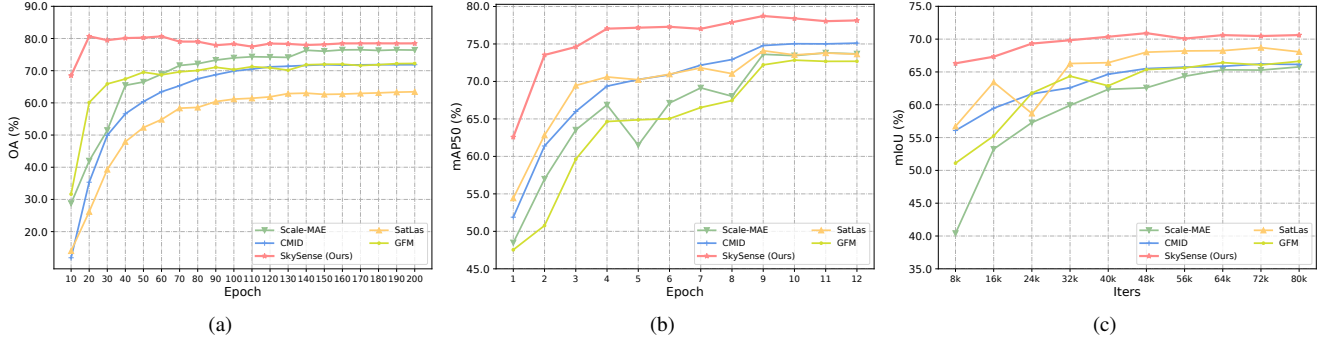


Figure 1. Convergence curves of different methods on downstream tasks: (a) scene classification on the AID dataset (TR=1%), (b) object detection on the DIOR dataset, and (c) semantic segmentation on the iSAID dataset.

Model	Publication	# Parameters	RESISC-45	DIOR	Postdam
			OA (TR=10% / 20%)	mAP ₅₀	mIoU
SatMAE [9]	NIPS'22	307M	91.72 / 94.10	70.89	90.63
Scale-MAE [29]	ICCV'23	307M	92.63 / 95.04	73.81	91.54
SkySense (Swin-L)	-	197M	94.34 / 95.92	76.74	92.86
SkySense (Swin-H)	-	654M	94.85 / 96.32	78.73	93.99

Table 3. Results of SkySense with smaller parameter size.

iSAID dataset [46], and change detection with the Dyna.-S2 dataset. The experimental results are presented in Tab. 4. The results across all 5 datasets demonstrate a substantial performance advantage of the our pre-trained model over the model learned from scratch.

We further conduct experiments between SkySense and the DINOv2 [28], a well-established Vision Foundation Model, on Earth Observation interpretation tasks. Notably, DINOv2 is pre-trained on a large volume of natural images. The obtained experimental results manifest the conspicuous superiority of our method in all 5 datasets, surpassing the performance of DINOv2. We posit that the unsatisfactory performance of DINOv2 in remote sensing scenarios arises from two primary factors. Firstly, a discernible domain gap exists between remote sensing imagery and natural images, rendering it arduous for the DINOv2 model, pre-trained solely on natural images, to effectively generalize across the diverse data sources and modalities inherent in remote sensing interpretation tasks. Secondly, DINOv2 lacks specialized design considerations geared towards remote sensing’s characteristics, particularly in terms of spatiotemporal knowledge of remote sensing imagery. Consequently, the model does not sufficiently leverage the abundant spatiotemporal knowledge from remote sensing imagery in the downstream tasks. Conversely, SkySense is purposefully formulated for remote sensing tasks, with tailored pre-training data, methods, and model structures that align harmoniously with downstream remote sensing inter-

pretation tasks. As a result, it exhibits a markedly superior performance.

6. Effectiveness of Factorized Spatiotemporal Encoder

We conduct a validation of SkySense’s fundamental unit, the Factorized Spatiotemporal (F-ST) encoder, and the results are presented in Tab. 5. In this validation, we compare F-ST encoder with two alternative options: the 3D model (UNet3D [23]) and the factorized temporospatial model (TSViT [38]). All three models are trained from scratch using Sentinel-2 data from the PASTIS-R dataset [13], with the purpose of evaluating the structure design only. To ensure fairness, we use the same training settings and a similar number of model parameters as described in [38]. Results show our F-ST encoder exhibits superior performance with substantially fewer parameters than 3D structure. Moreover, compared with TSViT, our design enables much better flexibility, as the fusion component is flexible enough to accommodate dimensions beyond time, such as modality and knowledge.

7. Effectiveness of Cross-Modal Alignment

The Cross-Modal Alignment of SkySense plays a role in explicitly aligning features from different modalities, facilitating cross-modal interactions. To intuitively validate the effectiveness of Cross-Modal Alignment, we conduct

Model	Scene Classification	Object Detection	Semantic Segmentation		Change Detection
	AID	DIOR	Dyna.-S2	iSAID	Dyna.-S2
	OA (TR=20%/ 50%)	mAP ₅₀	mIoU	mIoU	SCS
Random Init. (from scratch)	65.56/90.57	55.16	25.7/36.0	47.89	14.0/15.9
DINOv2 [28]	96.16/97.69	68.91	30.9/40.9	58.69	13.8/16.6
SkySense	97.68/98.60	78.73	33.1/46.2	70.91	15.4/18.0

Table 4. Results of the model learned from scratch, DINOv2 (ViT-L/14), and SkySense.

Architecture	# Parameters	PASTIS-R
		OA
UNet3D [23]	6.2M	82.3
TSViT [38]	1.7M	83.4
*F-ST encoder	2.3M	83.7

Table 5. Results of different space-time data processing architectures. Three models are trained from scratch using Sentinel-2 data from the PASTIS-R dataset. * denotes the adoption of an identical F-ST architecture as employed by our SkySense.

a feature visualization experiment. Specifically, we calculate the attention map of the output tokens of each Transformer layer in the Multi-modal Temporal Fusion Transformer. Two cases of visualization experiments are conducted: one with the Cross-Modal Alignment before the Multi-modal Temporal Fusion Transformer, and the other without it. The visualization results of different intermediate layers are shown in Fig. 2. The results in Fig. 2 (a) suggest that the model without Cross-Modal Alignment focuses their attention on the extra token and the token corresponding to HSROI. The unnecessary attention to the token corresponding to HSROI indicates that the information from the HSROI is not adequately integrated into the extra token, resulting in ineffective semantic information fusion. After incorporating the Cross-Modal Alignment, the model predominantly concentrates on particular extra token as shown in in Fig. 2 (b), indicating successful integration of the tokens associated with HSROI, TMsI and TSARI within the model. This indicates that the Cross-Modal Alignment indeed facilitates the comprehensive fusion of multi-modal RSI, which is crucial for constructing a Multi-Modal RSFM.

In addition to the results in Table 5(b) from our paper, we conduct an extra experiment on adding MAE into pre-training and a minor decrease (-0.7% mIoU) is observed on the DynamicEarthNet test set. We argue that the pixel-level modeling of MAE is already addressed by our Multi-Granularity Contrastive Learning.

8. Qualitative results of downstream tasks

We show qualitative results of SkySense and other representative RSFMs on various downstream tasks in this section. Specifically, we provide the visualizations of results and features on tasks including semantic segmentation, object detection, scene classification, *etc.* These visualizations are utilized to provide a more comprehensive and intuitive assessment of SkySense’s performance.

Semantic Segmentation. We visualize the semantic segmentation results on the commonly used iSAID dataset [46], as depicted in Fig. 3. In the first row, the scene depicts a port area. In comparison to two recent methods, Scale-MAE [29] and SatLas [1], SkySense achieves the most accurate segmentation results especially at `harbors` and `ships` across different scales. We speculate that our model enhances its ability to segment multi-scale objects through Multi-Granularity Contrastive Learning. The second row depicts an airport where planes are relatively small in terms of GSD. Other methods, particularly Scale-MAE, exhibit rough outlines in the regions of `plane wings`. Furthermore, both SatLas and Scale-MAE erroneously identify the containers in the right of the image as `large vehicles`. Our SkySense consistently achieves the most outstanding results, demonstrating excellence both in overall segmentation and fine details. Specifically, our method performs exceptionally well in segmenting `plane`, particularly in capturing the boundaries. The scene in the third row features a sports stadium that includes a `ground track field`, a `swimming pool`, and a `soccer ball field`. The `soccer ball field` on the left is not complete, posing challenges for accurate segmentation. Among various methods, SkySense successfully detects a large area of the `soccer ball field` and maintains a high consistency with the ground truth. In contrast, methods like SatLas and Scale-MAE even completely disregard the `soccer ball field` and provide incomplete predictions. Throughout all the obtained results, our method consistently delivers superb visual results, which further demonstrates the effectiveness of our approach in downstream semantic segmentation tasks.

Object Detection. The object detection results on the

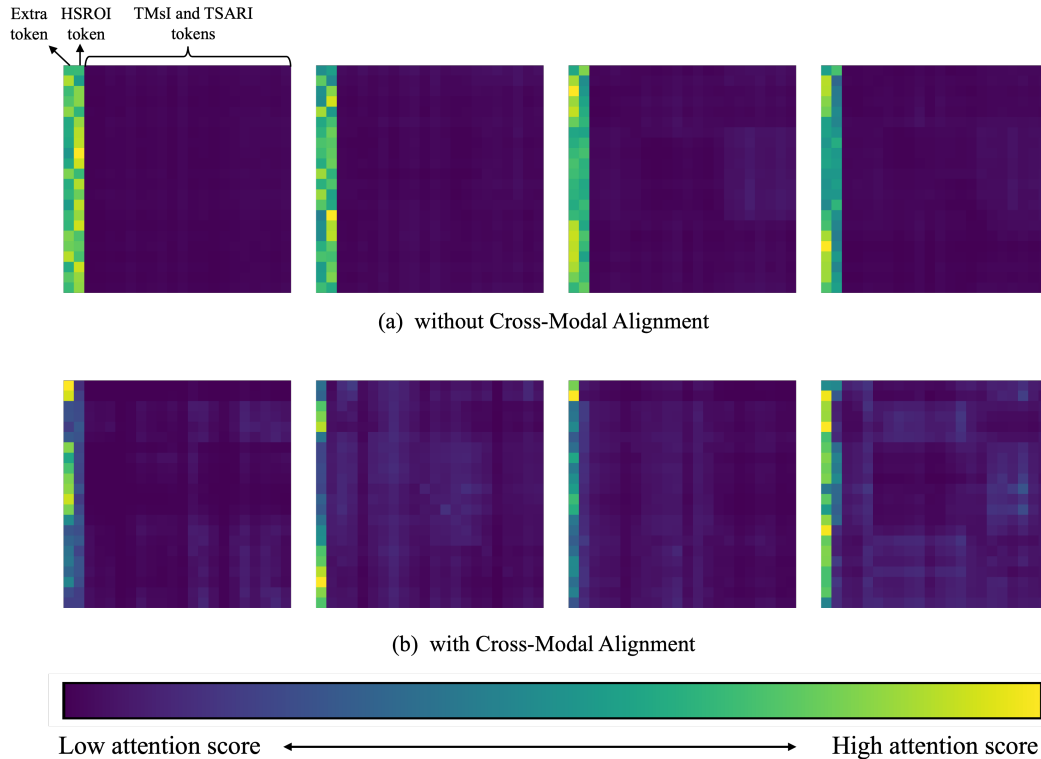


Figure 2. Attention map of the output tokens of Transformer layer in the Multi-modal Temporal Fusion Transformer. In each attention map, the first column corresponds to the extra token described in our paper, the second column represents the HSROI token, and the subsequent columns indicate TMSI and TSARI tokens. (a) Attention maps without the Cross-Modal Alignment; (b) correspondent attention maps with the Cross-Modal Alignment. The absence of Cross-Modal Alignment results in an unwarranted concentration of attention on the HSROI token, highlighting the inadequate integration for HSROI.

FAIR1M dataset [37] are visualized in Fig. 4. It is noteworthy that the FAIR1M dataset does not disclose labels associated with the test set. Quantitative results are obtained through online evaluation¹. The first two rows of the visualized results feature airport scenes. In the first row, SatLas mistakenly identifies the airplane runway lines as trucks. Although SatLas and Scale-MAE can detect the plane, their bounding boxes do not accurately match the plane. Parts of the airplanes are even outside the bounding boxes. Our method not only detects planes but also provides accurate bounding boxes for them. The scene in the second row is noticeably more complex, with multiple intersections and airplanes that visually overlap. Although SatLas captures all the airplanes, similar to the result in the first row, its bounding boxes do not accurately match the airplanes. Furthermore, in this case, SatLas mistakenly identifies an airplane runway line as a truck and completely overlooks intersections. Scale-MAE performs well in detecting airplanes, but it misses one intersection

¹<https://www.gaofen-challenge.com/benchmark>

in this scene and incorrectly detects the shadows on the airplane runway as small cars. The object detection results presented above clearly indicate that SkySense surpasses other recent methods and achieves superior overall performance. It demonstrates exceptional generalization capabilities for downstream object detection task.

Change Detection. We present the change detection visualization results obtained from the LEVIR-CD dataset [4] as shown in Fig. 5. In the first row, the main changes between the bi-temporal images are the newly built buildings surrounding the existing ones. SatLas and Scale-MAE exhibit overly rough change boundaries. SkySense demonstrates the best visual results that matches the ground truth well. The second row shows a significant increase in the number of buildings in the second temporal image. All three methods detect all the changed areas, and SkySense provides the best results regarding the boundaries. The results mentioned above demonstrate that our SkySense, when used for change detection task, not only accurately identifies the changed objects but also provides precise and detailed boundaries.

Scene Classification. In Fig. 6, we visualize the Gradient-

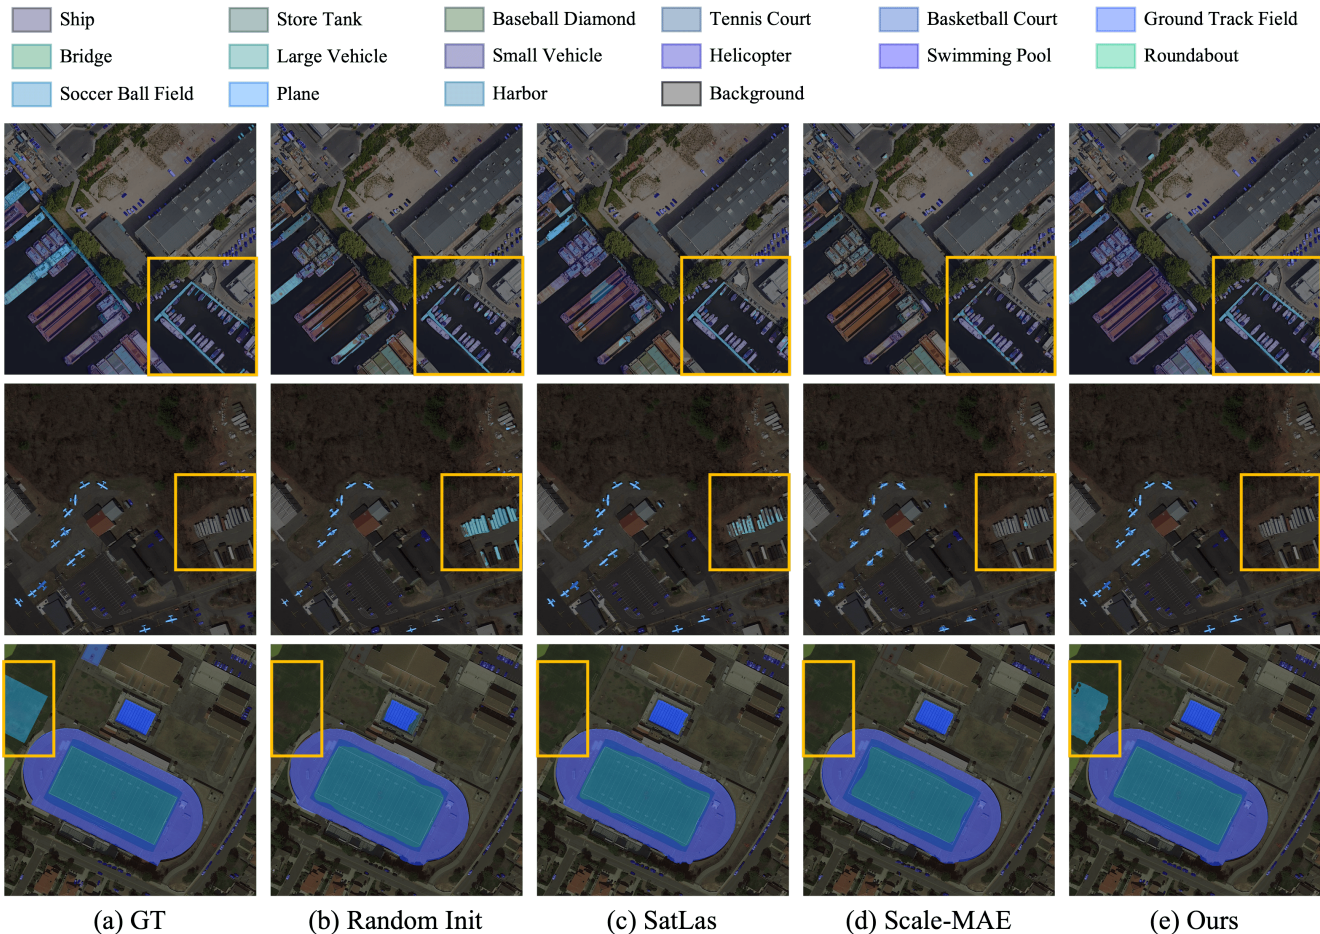


Figure 3. Visualization of semantic segmentation results on the iSAID dataset.

weighted Class Activation Mapping (Grad CAM) [32] on the AID dataset [48], illustrating the superior feature learning capabilities of SkySense. For the `school` category in the first row, SatLas exhibits an excessive emphasis on vegetation-covered areas, while Scale-MAE mainly focuses on some surrounding tiny buildings. SkySense effectively focuses on the main teaching building in the center of the scene and other associated areas, demonstrating its superior feature for this example. The second row represents a `viaduct` scene. SatLas inadequately assigns attention scores to some of the `viaduct` region. Scale-MAE only focuses on certain local areas, significantly lacking focus on the main subjects. Our method effectively and accurately prioritizes all viaducts in the image, assigning them high attention scores. In the third row, representing a `port` area, the observations bear a resemblance to those of the second row. Scale-MAE exhibits an excessive inclination towards smaller objects, while SatLas fails to assign an adequately high attention score to the primary subjects of interest. Our method shows the most compelling visualization results.

From the above feature visualizations, it is evident that SkySense demonstrates better performance compared to Scale-MAE and SatLas. Scale-MAE tends to overly emphasize tiny objects due to its scale-aware design, while SatLas exhibits inadequate attention towards the primary targets, diminishing their performance in this classification task. In contrast, our SkySense’s activation map effectively captures targets and their relevant features.

Multi-Modal Semantic Segmentation. One of the key advantages of SkySense is its support for multi-modal remote sensing imagery. To intuitively validate the robustness and effectiveness of SkySense in multi-modal Earth Observation interpretation tasks, we conduct a visual analysis of the results of multi-modal semantic segmentation on the DynamicEarthNet-MM dataset [39]. This dataset provides diverse and comprehensive multi-modal remote sensing imagery, including PlanetFusion, Sentinel-1, and Sentinel-2 imagery, allowing us to examine the performance of our method across different modalities. It is worth noting that the labels for the validation and test sets have not




Figure 4. Visualization of object detection results on the FAIR1M dataset. It is worth noting that the labels for the test set have not been made public, therefore the visualization results do not contain ground truth.

been made public, therefore the visualization results do not contain ground truth. The qualitative results are illustrated in Fig. 7. The input images corresponding to the first case are shown in the first row of Fig. 7. This particular scene provides a broad field of view, encompassing multiple land cover categories such as agriculture, soil, and water. To provide a comprehensive comparison with SatLas, which does not support multi-modal data input, we also contrast the results of single-modal segmentation on Sentinel-2 and PlanetFusion images. Regarding the single-modal input of Sentinel-2 imagery, SkySense exhibits a higher accuracy in detecting small water bodies within the soil, primarily attributed to the incorporation of Multi-Granularity Contrastive Learning. In the case of single-modal input from PlanetFusion imagery, SkySense successfully segments the rectangular water body area in the lower right corner, demonstrating better single-modal segmentation capabilities. SatLas fails to recognize those specific areas. When the training data includes three sources:

Sentinel-2, Sentinel-1, and PlanetFusion, SkySense exhibits segmentation performance that surpasses the single-modal segmentation. Specifically, with the advantage of multi-modal remote sensing imagery training, SkySense can successfully segment certain agricultural areas that are difficult to detect under single-modal conditions as shown in Fig. 7 (h). The second case is shown in the third row of Fig. 7. The input scene contains a river that runs through the entire image. SatLas, regardless of whether trained on Sentinel-2 or PlanetFusion data, fails to segment the continuous river. Even when training on single-modal data, SkySense is capable of segmenting relatively continuous and complete rivers. However, there are some minor flaws in the details. For example, in the lower right portion of the image, soil is mistakenly detected as water, and the segmentation of wetlands in the upper right portion is incomplete. When training on multiple modalities, SkySense shows significant visual improvement in this case. The segmentation flaws in the single-modal cases mentioned above are no longer present. Thanks to the design of the Factorized Multi-modal Spatiotemporal Encoder that separates spatial feature extraction from multi-modal temporal fusion, SkySense can flexibly utilize either a single modality or multiple modalities for training, which gives it a notable advantage over existing methods. Additionally, in the aforementioned visualized results, even trained and tested with a single modality, SkySense’s performance still surpasses other methods, demonstrating the strong generalization capability and learning ability. With the training data of multiple modalities, the effectiveness of SkySense is further improved. This not only highlights the significant enhancement in semantic segmentation of remote sensing imagery by incorporating multiple modalities, but also underscores the importance of supporting multi-modal remote sensing imagery for RSFM.

After conducting a comprehensive analysis of the results, it becomes apparent that SkySense consistently outperforms other methods in a range of downstream tasks. Notably, it demonstrates remarkable accuracy in executing segmentation, detection, and classification processes. These findings effectively highlight the exceptional learning capacity and generalization capability of SkySense for downstream tasks.

9. Pre-training dataset

Existing remote sensing datasets lack the numerous amounts of multi-modal time-series Remote Sensing Imagery required for building SkySense, thus we develop a comprehensive multi-modal remote sensing dataset with temporal sequences specifically for SkySense pre-training.

Data Acquisition and Preprocessing. This dataset comprises diverse sources of remote sensing imagery collected globally (see Tab. 6), including HSROIs from WorldView-3,




Figure 5. Visualization of change detection results on the LEVIR dataset.

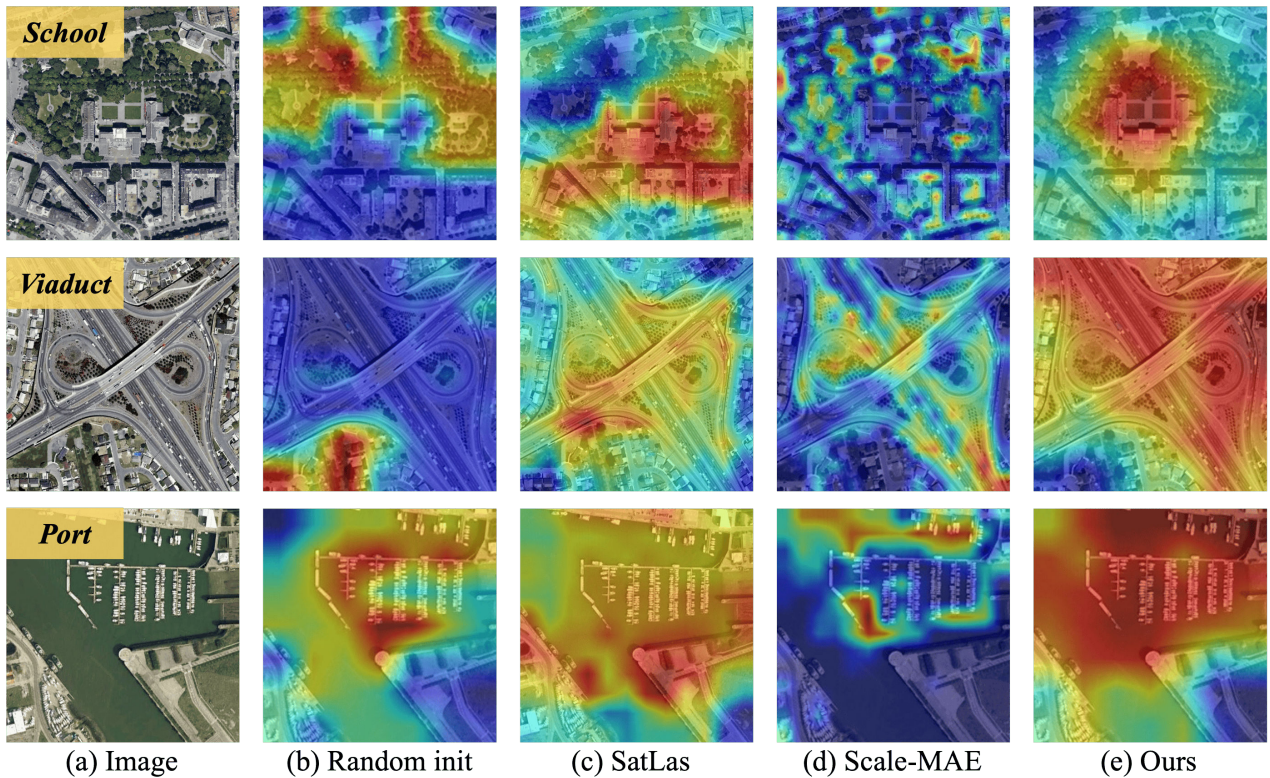


Figure 6. Visualization of Grad CAM on the AID dataset (TR=20%).

4, etc., TMSI from Sentinel-2, and TSARI from Sentinel-1.

- *HSROIs*. We collect high-resolution optical RGB images from a third-party platform, with an average Ground Sample Distance (GSD) of 0.3 meter.
- *TMSI*. We collect the freely available Sentinel-2 level-2A atmospherically corrected surface reflectance sequence images as another significant data source. The bands with a resolution of 10m (visible and NIR) and resam-

pled bands with a resolution of 20m (Vegetation Red Edge and SWIR) are merged to form a multispectral image (10 bands). In addition, cloudy imagery filtering is implemented to obtain higher quality multispectral data. Specifically, according to the Scene Classification Map² provided by the European Space Agency,

²<https://custom-scripts.sentinel-hub.com/>

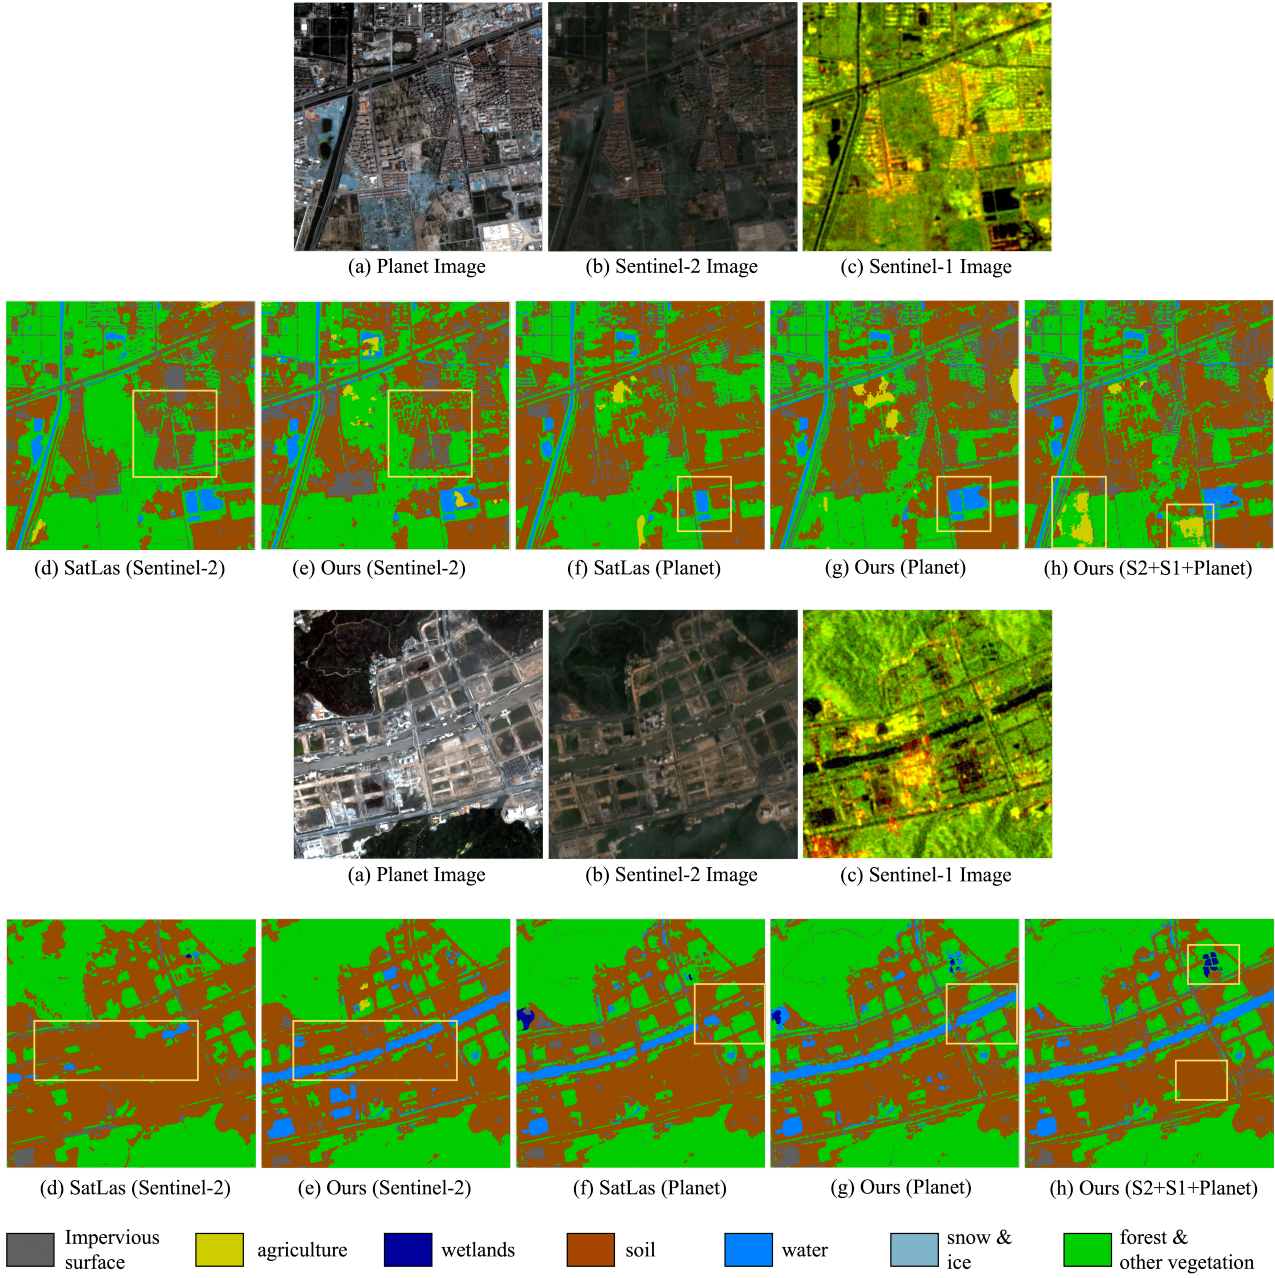


Figure 7. Visualization of multi-modal semantic segmentation results on the DynamicEarthNet-MM dataset. It is worth noting that the labels for the validation and test sets have not been made public, therefore the visualization results do not contain ground truth.

images with a cloud coverage ratio exceeding 1% (calculated as the sum of the proportion of pixels belonging to the categories of Thin cirrus, Cloud medium probability, Cloud high probability, and Cloud shadows categories) are filtered out.

- *TSARI*. We obtain the easily accessible Sentinel-1 ground-range-detected (GRD) products with both VV and

VH polarization, which are acquired in the interferometric wide swath (IW) mode. And the standard calibration³ process is employed to obtain the spatially-aligned SAR images. These images contain the backscatter coefficient (σ°) in decibels (dB).

Dataset statistics. As shown in Fig. 8, our dataset spans

custom-scripts/sentinel-2/scene-classification/

³<https://sentinels.copernicus.eu/web/sentinel/toolboxes/sentinel-1>



Figure 8. Geographical distribution of pre-training data. The green areas represent the countries or regions covered by the pre-training data. The basic world map is obtained from <https://www.resdc.cn/data.aspx?DATAID=205>.




Figure 9. Visualization of three training samples from our pre-training dataset. Each sample contains a HSROI, TMsIs and TSARIs. The first column represents HSROI. The second to fourth columns stand for multispectral false-color images captured at different times. The fifth to seventh columns are SAR grayscale images captured at different times.

over 8.78 million square kilometers across 40 countries and areas. It covers 6 continents, *i.e.*, Asia, Europe, Africa, North America, South America and Oceania. The dataset contains 21.5 million training sequences, each consisting of 1 static HSROI, a randomly-sampled TMsI of sequence

length 20 and a randomly-sampled TSARI of sequence length 10. In total it occupies a storage space of around 300 Terabytes. Our dataset exhibits significant complementarity in terms of temporal information, spatial resolution, and imaging mechanisms, given the three modalities.

Modality	Sensor	Band / Polarization	GSD (m)	Image Size (px)	Avg. seq Length
Optical (HSROI)	WorldView-3,4, etc.	RGB	0.3	2048 × 2048	1
Optical (TMsI)	Sentinel-2 (Level-2A)	B2-8, B8A, B11-12	10	64 × 64	65
SAR (TSARI)	Sentinel-1 (Level-1 IW, GRD)	VV, VH	10	64 × 64	13

Table 6. Statistics of our pre-training dataset. Ground sample distance (GSD) is the distance between the center of one pixel to the center of an adjacent pixel in a remote sensing image.

Example visualization. Fig. 9 illustrates some examples from our pre-training dataset. Each example consists of a HSROI, TMsIs and TSARIs.

10. Pre-training implementation details

SkySense is pre-trained with a batch size of 240 samples for a total of 875k steps, distributed over 80 A100-80GB GPUs with the AdamW optimizer [22]. We adopt a learning rate warmup [14], followed by a decay using a cosine schedule [21] from 0.04 to 0.2. For HSROIs, we apply augmentations including multi-crop [2], Gaussian blur, solarization [15], *etc.* As for TMsI and TSARI, we randomly select two fixed-sized sequences (*i.e.* 20 for TMsI and 10 for TSARI) from the original ones and perform random disturbances on the RSI acquisition date. We employ the huge version⁴ of the Swin Transformer (Swin-H) [19] as the spatial encoder for HSROIs, chosen for its design efficiency in minimizing computational costs for high-resolution imagery [50]. RSI from TMsI or TSARI is equipped with a ViT-L [11]. The Multi-Modal Temporal Fusion Transformer contains 24 Naive Transformer encoder layers. For Geo-Context Prototype Learning, we divide the globe into 4096 regions. One region covers approximately 4294 square kilometers area and contains 100 prototypes. SkySense comprises a total of 2.06 billion parameters, specifically 654 million from the Swin-H and 302 million from a single ViT-L, as shown in Tab. 7.

The pixel size of our pre-training data is shown in Tab. 6, which is up to 2048 × 2048. The pre-training takes 24600 A100 GPU hours. Its computational complexity is 4488.69 GFLOPs.

11. Dataset and implementation details of downstream tasks

In this section, we introduce the experimental datasets and implementation details used in downstream task.

Semantic Segmentation. Semantic segmentation serves as a prevalent application in remote sensing, facilitating

⁴<https://github.com/open-mmlab/mmpretrain>

Model Modules	Architecture	# Parameters
Spatial Encoder-HSROI	Swin-Huge	654M
Spatial Encoder-TMsI	ViT-Large	302M
Spatial Encoder-TSARI	ViT-Large	302M
Multi-modal Temporal Fusion Transformer	Transformer Encoder	398M
Geo-Context Prototype	-	215M
Others	-	189M

Table 7. Parameter breakdown for each module of SkySense.

the automatic extraction of land use classes and ground instances. Considering factors such as spatial resolution, spectrum and number of categories, we select four popular datasets for the semantic segmentation task:

- 1) *DynamicEarthNet-PlanetFusion (Dyna.-Pla.)* [39]. The dataset comprises a collection of images from 75 global locations captured from the PlanetFusion satellite platform. The image acquisition period spans from January 2018 to December 2019. Each location has 24 images and corresponding annotations of 7 land use and land cover semantic classes. Each image contain four bands, namely Red, Green, Blue and Near-Infrared, with a GSD of 3 meters and an image size of 1024×1024. Based on the official leaderboard⁵, these locations are divided into 55 for training, 10 for validation, and 10 for testing. In the experiment, we use the official validation and test sets for evaluation. It is worth noting that ground truth labels for both the validation and test sets are unavailable for local users. Therefore, we submit the predictions to online leaderboard and report the obtained scores.
- 2) *iSAID* [46]. This dataset comprises 2806 remote sensing images with dense annotations from multiple satellite sensors. The images vary in size from 800×800 to 4000×13000 pixels. It includes pixel-level annotations of 655451 instances across 15 object categories, while the remaining non-object pixels are labeled as background. It has been divided into training, validation, and test sets, consisting of 1411, 458, and 937 samples, respectively. Following [36, 41], the performance evaluation of RSFMs is conducted on the validation set.
- 3) *Potsdam* [33]. A total of 38 aerial images with a GSD of 0.05 meters are collected to form the Potsdam dataset. These images are divided into 24 training images and 14 testing images. Each image has a fixed pixel size of 6000×6000. Following [36], we utilize images composed of Near-Infrared, Red, and Green spectral bands. The evaluation is conducted on the test set, focusing on five categories: impervious surfaces, buildings, low veg-

⁵<https://codalab.lisn.upsaclay.fr/competitions/2882#results>

Task	(i) Semantic Segmentation				(ii) Object Detection		
Dataset	Dyna.-Pla.	iSAID	Potsdam	Dyna.-S2	DIOR	DIOR-R	FAIR1M
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
Input Size	1024×1024	896×896	512×512	256×256	800×800	800×800	512×512
Input channel	RGBNIR	RGB	NIRRG	B02-08, B8A, B11-12	RGB	RGB	RGB
Base learning rate	6e-5	6e-5	6e-5	6e-5	1e-4	1e-4	1e-4
Learning rate scheduler	poly	poly	poly	poly	multistep	multistep	multistep
Weight decay	0.01	0.01	0.01	0.01	0.05	0.05	0.05
Batch size	8	16	16	8	4	2	12
Max iteration/epoch	8k iters	80k iters	80k iters	80k iters	12 epoch	12 epoch	8 epoch
Warmup	linear	linear	linear	linear	linear	linear	linear
Warmup iteration/epoch	1.5k iters	1.5k iters	1.5k iters	1.5k iters	1k Iters	1k iters	500 iters
Warmup ratio	1e-6	1e-6	1e-6	1e-6	1e-3	1e-3	1e-3
Drop path rate	0.3	0.3	0.3	0.3	0.3	0.3	0.3
Augmentation	RandomCrop, RandomFlip	RandomScaling (0.5 to 2.0), RandomCrop, RandomFlip	RandomScaling (0.5 to 2.0), RandomCrop, RandomFlip	RandomCrop, RandomFlip	RandomFlip	RandomFlip	RandomFlip, RandomRotate, Multi-Scale
Head/Detector	UperNet	UperNet	UperNet	UperNet	Faster RCNN	Oriented RCNN	Oriented RCNN
Loss function	CrossEntropy	CrossEntropy	CrossEntropy	CrossEntropy	CrossEntropy, SmoothL1	CrossEntropy, SmoothL1	CrossEntropy, SmoothL1

Task	(iii) Scene Classification				(iv) Change Detection		
Dataset	AID	RESISC-45	BEN-S2	fMoW-S2	LEVIR-CD	OSCD	Dyna.-S2
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW	Adam	AdamW
Input Size	224×224	224×224	128×128	96×96	256×256	96×96	256×256
Input channel	RGB	RGB	B02-08, B8A, B11-12	B02-08, B8A, B11-12	RGB	B02-08, B8A, B11-12	B02-08, B8A, B11-12
Base learning rate	6.25e-5	6.25e-5	5e-5	8e-4	6e-5	6e-4	6e-5
Learning rate scheduler	Cosine Annealing	Cosine Annealing	MultiStepLR	Cosine Annealing	LambdaLR	ExponentialLR	poly
Weight decay	0.05	0.05	0.01	0.05	0.01	1e-4	0.05
Batch size	64	64	256	1024	8	32	8
Max iteration/epoch	200 epoch	200 epoch	100 epoch	30 epoch	200 epoch	100 epoch	80k iters
Warmup	linear	linear	-	linear	-	-	linear
Warmup iteration/epoch	5 epoch	5 epoch	-	5 epoch	-	-	1.5k iters
Warmup ratio	0.01	0.01	-	0.2	-	-	1e-6
Drop path rate	0.2	0.2	-	0.2	-	-	0.3
Augmentation	RandomCrop, RandomErasing	RandomCrop, RandomErasing	RandomFlip	RandomCrop, RandomFlip, Mixup, CutMix	RandomCrop, RandomFlip, RandomBlur	RandomFlip, RandomRotate	RandomCrop, RandomFlip
Head/Detector	Linear Classifier	Linear Classifier	Linear Classifier	Linear Classifier	BIT	U-Net	UperNet
Loss function	CrossEntropy	CrossEntropy	MultiLabel SoftMargin	SoftTarget CrossEntropy	CrossEntropy	BCE	CrossEntropy

Table 8. The finetuning setting in single-modal downstream tasks.

etation, trees, and cars. It is important to note that the clutter category is not included in the evaluation.

4) *DynamicEarthNet-Sentinel2 (Dyna.-S2)* [39]. This dataset can be viewed as the Sentinel-2 data version of the abovementioned DynamicEarthNet-PlanetFusion dataset. Specifically, the DynamicEarthNet-Sentinel2

dataset offers monthly Sentinel-2 multispectral images, acquired between January 2018 and December 2019, that are spatially aligned with the corresponding PlanetFusion images. Each image in the dataset comprises 12 spectral channels and is uniformly resampled to match the sizes of PlanetFusion images, which are 1024×1024

Task	(i) Multi-Modal Segmentation: Time-insensitive LandCover Mapping	(ii) Multi-Modal Segmentation: Time-sensitive Crop Mapping	(iii) Multi-Modal Classification
Dataset	Dyna.-MM	PASTIS-MM	BEN-MM
Optimizer	AdamW	AdamW	AdamW
Input Size	planet: 1024×1024 sentinel2: 1024×1024 sentinel1: 1024×1024	gep: 4096×4096 sentinel2: 128×128 sentinel1: 128×128	sentinel2: 128×128 sentinel1: 128×128
Input channel	planet: RGBNIR sentinel2: B02-08, B8A, B11-12 sentinel1: VV, VH	gep: RGB sentinel2: B02-08, B8A, B11-12 sentinel1: VV, VH	sentinel2: B02-08, B8A, B11-12 sentinel1: VV, VH
Base learning rate	6e-05	6e-05	5e-05
Learning rate scheduler	linear	linear	MultiStepLR
Weight decay	0.01	0.01	0.01
Batch size	8	8	256
Max iteration/epoch	6k iters	20k iters	100 epoch
Warmup	linear	linear	-
Warmup iteration/epoch	150 iters	1500 iters	-
Warmup ratio	1e-6	1e-6	-
Drop path rate	0.3	0.3	-
Augmentation	RandomFlip	RandomFlip	RandomFlip
Head/Detector	UperNet	FCN	Linear Classifier
Loss function	CrossEntropy	CrossEntropy	MultiLabel SoftMargin

Table 9. The finetuning setting in multi-modal downstream tasks.

pixels. Following the same official data split protocol, we report the mIoU metric evaluated on the leaderboard and -test in our paper.

We employ the UperNet [49] as the unified segmentation head based on MMsegmentation⁶, following [3, 36, 41]. The detailed fine-tuning setting can be found in Tab. 8 (i).

Horizontal & Oriented Object Detection. We employ the DIOR dataset to evaluate the performance of SkySense and other RSFMs for horizontal object detection task. Following [36], we utilize Faster RCNN [30] as the detector, and other details are presented on Tab. 8 (ii).

1) *DIOR* [17]. This dataset includes 23463 visible remote sensing images and 192472 object instances, which are manually annotated with horizontal bounding boxes and categorized into 20 common object classes. The image size in the dataset is of 800×800 pixels, with GSD ranging from 0.5 meters to 30 meters. The dataset is divided into a training set consisting of 5862 patches, a validation set comprising 5863 patches, and a test set totaling 11738 patches. Following [36], we mix the training set and validation set during training, while the test set is reserved for evaluation. In particular, its high inter-class similarity and intra-class diversity make precise localization and classification extremely challenging.

Remote sensing images encompass a wide range of objects, such as buildings, vehicles, bridges and so on. These objects are densely distributed and display variations in

terms of size, scale and orientation. Consequently, detecting and identifying these objects presents a significant challenge, commonly referred to as oriented object detection [47]. To assess the performance of RSFMs on this task, we utilize the DIOR-R and FAIR1M datasets and employ the Oriented RCNN [18] as the detector, following [3, 36, 41]. The specific implementation details can be found in Tab. 8 (ii) as well.

2) *DIOR-R* [7]. The DIOR-R dataset shares the same images as the abovementioned DIOR dataset. However, the difference lies in the annotated oriented bounding boxes, which make it suitable for oriented object detection task. Similar to the implementation on the DIOR dataset, we follow [41] by merging the training and validation sets during the training process, while the test set is reserved for evaluation.

3) *FAIR1M* [37]. FAIR1M is a large-scale benchmark dataset for fine-grained oriented object detection, including over 40000 high-resolution optical remote sensing images and more than 1 million instances collected from various regions worldwide. It has been annotated with oriented bounding boxes for 5 categories and 37 fine-grained subcategories. We test all models on the official leaderboard-v2.0⁷ and report the mean average precision (mAP) metric score.

Change Detection. Change detection aims to find pixel-level regional changes via bi-temporal or multi-temporal

⁶<https://github.com/open-mmlab/msegmentation>

⁷<https://www.gaofen-challenge.com/benchmark>

images. Based on [36], we incorporate the backbones of different RSFMs into the BIT framework [5] to evaluate their performance on the LEVIR-CD dataset. Following [24, 25], we use U-Net [31] as the segmentation head to evaluate the effectiveness of RSFMs on the bi-temporal change detection task using the OSCD dataset with multispectral imagery. Additionally, we employ the DynamicEarthNet-Sentinel2 dataset to assess the performance of the models on the semantic change detection task, maintaining the same setup as segmentation task. Other setting is shown in Tab. 8 (iv).

- 1) *LEVIR-CD* [4]. LEVIR-CD is a dataset focused on building change detection, containing of 637 pairs of visible images with a GSD of 0.5m. Each image has a size of 1024×1024 pixels, and the image acquisition time span ranges from 2002 to 2018. The images from different time periods exhibit significant building changes, particularly in the area with rapid population growth. In addition, binary labels (1 indicating change, 0 indicating no change) are provided to indicate the buildings' change status in these bi-temporal images. We report the F1-score on the test set using the same split as [36].
- 2) *OSCD* [10]. This dataset contains 24 pairs of multispectral images obtained from the Sentinel-2 sensor. Following the setup in [25], 14 pairs are used for training, and the rest pairs are used for testing. By dividing the original images into non-overlapping patches of size 96×96 pixels, we obtain 827 patches for training and 285 patches for testing.
- 3) *DynamicEarthNet-Sentinel2 (Dyna.-S2)* [39]. The DynamicEarthNet-Sentinel2 dataset can be used to assess the performance of RSFMs on the semantic change detection task as well. Unlike semantic segmentation task, we report the semantic change segmentation (SCS) score, which consists of two components: a class-agnostic binary change score (BC) and a semantic segmentation score among changed pixels (SC). The BC score quantifies the accuracy in identifying changing pixels, while the SC score reflects the ability of methods to accurately predict the category where the changed pixels belong to.

Scene Classification. We select two commonly-used single-label scene classification datasets (*i.e.* AID and NWPU-RESISC45 datasets), a multi-label multispectral scene classification dataset (*i.e.* BigEarthNet-Sentinel2 dataset), and a temporal multispectral scene classification dataset (*i.e.* fMoW-Sentinel2 dataset). We conduct scene classification experiments using a standard linear classifier. The implementation details are provided in Tab. 8 (iii).

- 1) *AID* [48]. The AID dataset comprises 10000 images, each with a size of 600×600 pixels and a GSD ranging from 0.5 meters to 8 meters. The images in the

dataset are divided into 30 categories, with each category containing approximately 220 to 400 images. In experiments, we use $x\%$ of the data for training and the rest $(1 - x\%)$ for testing, following common protocols in [36, 41], where $x \in \{20, 50\}$.

- 2) *NWPU-RESISC45 (RESISC-45)* [6]. This dataset includes 31500 images, each with a size of 256×256 pixels and a GSD ranging from 0.5 meters to 30 meters. It is divided into 45 categories, with each category containing 700 images. Similar to previous works [36, 41], we use 10% and 20% of the data for training and the remaining 90% and 80% for testing respectively.
- 3) *BigEarthNet-Sentinel2 (BEN-S2)* [34]. The BigEarthNet-Sentinel2 dataset is a large-scale multispectral image dataset used for multi-label land cover scene classification. This dataset consists of a total of 590326 multispectral images, each of which provides multiple land use category annotations. In line with previous studies [25, 45], we adopt the new scheme of 19 classes proposed in [34], and exclude approximately 12% of the patches that are completely masked by seasonal snow, clouds, or cloud shadows in the experiment. In addition, we employ the same data splits as [25, 43, 45], where 311667 samples are allocated for training and 103944 images are reserved for validation.
- 4) *fMoW-Sentinel2 (fMoW-S2)* [9]. The fMoW-Sentinel2 dataset is an extension of the fMoW-RGB dataset [8], focusing on temporal multispectral scene classification. For each location, time-series Sentinel-2 images and their corresponding labels are provided. The dataset consists of images with 13 spectral bands provided by Sentinel-2 (B1-12 and B8A), which are captured at different points in time. Some of the time points are the same as the original fMoW-RGB and others are additional points to form a decent time series. In total, the dataset contains 712874 training images, 84939 validation images, and 84966 test images. Following previous works [9, 12], we fine-tune the models using the complete training set and report the Top-1/5 Accuracy metrics on the validation set.

Multi-Modal Semantic Segmentation. By integrating multi-modal data from different sensors, imaging mechanisms, resolutions, and spectral bands, we can obtain a more diverse and discriminative features. These features enhance the understanding and interpretation of the shape, size, and relationships among ground objects. Thus, we employ the DynamicEarthNet-MM dataset and the PASTIS-MM dataset to evaluate the tasks of Time-insensitive Land Cover Mapping and Time-sensitive Crop Mapping, respectively.

- 1) *DynamicEarthNet-MM (Dyna.-MM)* [39]. This dataset contains spatially and temporally aligned multi-modal data, including PlanetFusion imagery

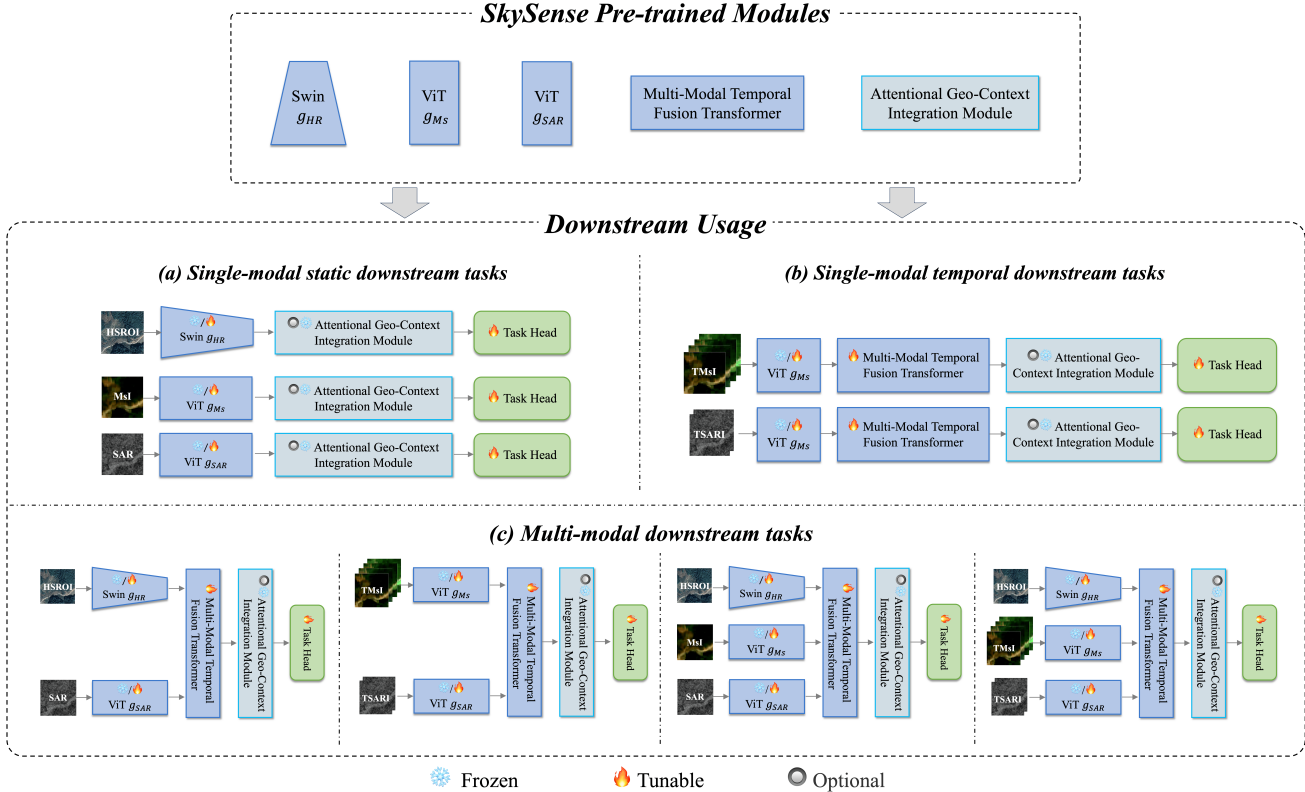


Figure 10. A detailed illustration on combination of pre-trained modules to accommodate different tasks.

(i.e. DynamicEarthNet-PlanetFusion dataset), Sentinel-2 multispectral imagery (i.e. DynamicEarthNet-Sentinel2 dataset), and Sentinel-1 SAR imagery. In case of SAR data, we collect standard-calibrated Sentinel-1 GRD data (VV, VH Polarization) according to the geographical coordinates of the optical imagery, thereby fulfilling the requirements of the multi-modal experiment. To ensure consistency, we employ UperNet as the segmentation head and present the mIoU metric derived from the official Leaderboard-test evaluation hosted online. More implementation details can be seen in Tab. 9 (i).

- 2) *PASTIS-MM* [13]. We develop the PASTIS-MM dataset for the task of fine-grained time-sensitive crop mapping. This dataset is an extension of the PASTIS-R dataset [13], incorporating spatially aligned high-resolution RGB images. It aims to investigate the impact of the combined usage of high-resolution optical imagery, medium-resolution temporal multispectral data, and temporal SAR data in the field of time-sensitive crop mapping. To create the PASTIS-MM dataset, we extract the geo-coordinates and acquisition dates from the image tiles of PASTIS-R dataset. Then we match every image tile with its corresponding static high-resolution optical image, whose GSD is about 0.3 me-

ter. The PASTIS-MM dataset consists of 2433 Sentinel-2 TMsI, each having an image size of 128×128 pixels, 10 spectral bands, and a GSD of 10 meters. It also includes Sentinel-1 GRD SAR images (with VV, VH, and VV/VH channels) and static high-resolution RGB images that we added. For each image tile, the dataset provides all available Sentinel-2 and Sentinel-1 acquisition data between September 2018 and November 2019, along with the additional high-resolution visible acquisition. Based on statistical analysis, each time series includes approximately 33 to 61 multispectral acquisitions, 70 radar acquisitions, and one high-resolution visible acquisition. The dataset encompasses 18 crop categories and covers a geographical area exceeding 4000 square kilometers. We intend to release this extended dataset publicly to facilitate the advancement of agriculture-vision research. In our experiments, the cloud coverage ratio of the Sentinel-2 images is obtained by utilizing Sentinel Hub’s cloud detector⁸. In addition, we use a naive FCN head [20] and report the Overall Accuracy (OA) from the official five-fold validation on this dataset. More implementation information

⁸<https://github.com/sentinel-hub/sentinel2-cloud-detector>

```

1 # Based on MMPretrain codebase
2 # Scene Classification on HSRIOIs.
3 checkpoint = 'skysense_swin_huge_ghr.pth'
4
5 model = dict(
6     type='ImageClassifier',
7     backbone=dict(
8         type='SwinTransformer',
9         img_size=224,
10        arch = dict({'embed_dims': 352,
11                    'depths': [2, 2, 18, 2],
12                    'num_heads': [8, 16, 32, 64]}),
13        patch_size=4,
14        in_channels=3,
15        window_size=7,
16        drop_rate=0.,
17        drop_path_rate=0.2,
18        out_indices=(3, ),
19        out_after_downsample=False,
20        use_abs_pos_embed=False,
21        interpolate_mode='bicubic',
22        with_cp=False,
23        frozen_stages=-1,
24        norm_eval=False,
25        pad_small_map=False,
26        norm_cfg=dict(type='LN'),
27        init_cfg=dict(type='Pretrained',
28                    checkpoint=checkpoint,
29                    prefix='backbone')),
30    neck=dict(type='GlobalAveragePooling'),
31    head=dict(
32        type='LinearClsHead',
33        num_classes=30,
34        in_channels=2816,
35        loss=dict(type='CrossEntropyLoss',
36                loss_weight=1.0),
37        topk=(1, 5)))

```

Figure 11. Example configuration for using SkySense pre-trained weights in single-modal scene classification task on HSRIOIs.

can be seen in Tab. 9 (ii).

Multi-Modal Scene Classification. We further employ the representative BigEarthNet-MM dataset to assess the performance of the Skysense in large-scale scene classification task, considering the integration of optical and SAR data. Detailed implementation details are provided in Tab. 9 (iii).

1) *BigEarthNet-MM (BEN-MM) [35]*. The BigEarthNet-MM dataset expands upon the aforementioned BigEarthNet-Sentinel2 dataset by incorporating corresponding Sentinel-1 SAR data, facilitating the evaluation of multi-modal (optical and SAR) multi-label scene classification task. The BigEarthNet-MM dataset supplements each Sentinel-2 image patch in the BigEarthNet-Sentinel2 dataset with a corresponding preprocessed Sentinel-1 image patch that shares a similar timestamp. Additionally, each Sentinel-1 image patch inherits the annotation information from its

corresponding Sentinel-2 image patch. The resulting Sentinel-1 image patches possess a GSD of 10 meters, providing dual-polarization information channels (VV and VH), and are based on interferometric wide-swath mode. Following [12, 42, 43], we adopt the same data splits as the BigEarthNet-Sentinel2 dataset.

12. Downstream usage with SkySense pre-trained weights

Our released SkySense pre-trained weights encompass five pivotal modules (as depicted in the upper section of Fig. 10):

- The Spatial Encoders g_{HR} , g_{Ms} , and g_{SAR} , which are responsible for extracting representations of corresponding modal images.
- The Multi-Modal Temporal Fusion Transformer module, designed to integrate multi-modal temporal representations.
- The Attentional Geo-Context Integration module, which introduces the region-specific prototype set to enhance the features.

The aforementioned key components of the pre-trained weights can be used alone or combined with the others flexibly to accommodate the requirements of various downstream tasks. We categorize the downstream tasks into three broad types based on the input data types, and we discuss how to utilize the corresponding pre-trained components to suit each of them respectively.

- 1) *Single-modal static downstream tasks.* As shown in the Fig. 10 (a), the Spatial Encoders for the corresponding modalities are employed to extract spatial representations, with their parameters being either freezable or tunable. Moreover, if the geographic coordinates corresponding to the images are available, the Attentional Geo-Context Integration module can be involved to attentively integrate pre-trained geo-context information. It is noteworthy that the use of the Attentional Geo-Context Integration module is optional for the user. The output features are fed into task-specific heads for further fine-tuning to obtain the desired outputs.
- 2) *Single-modal temporal downstream tasks.* Different from *single-modal static downstream tasks*, the parameter-learnable Multi-Modal Temporal Fusion Transformer module is applied after the Spatial Encoder to merge features from temporal sequences, as shown in the Fig. 10 (b).
- 3) *Multi-modal downstream tasks.* Fig. 10 (c) demonstrates examples of various multi-modal data combinations. The respective Spatial Encoders are applied to extract features from the corresponding modalities (static or temporal) data. After concatenation and reshaping, these features are fed into the parameter-learnable Multi-

Modal Temporal Fusion Transformer module to obtain a fused representation. The optional Attentional Geo-Context Integration module may be considered to further enhance the features before they are input into various task-specific decoders. This process allows for the effective integration and enhancement of multi-modal data, which is crucial for complex downstream tasks that require a comprehensive understanding of both spatial and temporal cues.

SkySense pre-trained weights exhibits compatibility with the MMCV framework⁹ as well as other prevalent codebase repositories (e.g., torchvision¹⁰ and TIMM¹¹), necessitating only rudimentary conversions. Herein, we provide a configuration example for single-modal scene classification based on MMPretrain codebase¹², as illustrate in Fig. 11. More comprehensive usage guidelines can be found within our project repository.

References

- [1] 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. 1, 2, 4
- [2] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020. 11
- [3] Keumgang Cha, Junghoon Seo, and Taekyung Lee. A billion-scale foundation model for remote sensing images. *arXiv preprint arXiv:2304.05215*, 2023. 1, 13
- [4] 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. 5, 14
- [5] 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. 14
- [6] 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, 14
- [7] 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. 13
- [8] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6172–6180, 2018. 14

⁹<https://github.com/open-mmlab/mmcv>

¹⁰<https://pytorch.org/vision/stable/index.html>

¹¹<https://github.com/huggingface/pytorch-image-models>

¹²<https://github.com/open-mmlab/mmpretrain/tree/mmcv-0.x>

- [9] 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, 3, 14
- [10] 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. 14
- [11] 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. 11
- [12] 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. 14, 16
- [13] 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. 3, 15
- [14] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 11
- [15] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 11
- [16] M. Joseph Hughes and Robert Kennedy. High-quality cloud masking of landsat 8 imagery using convolutional neural networks. *Remote Sensing*, 2019. 2
- [17] 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. 1, 2, 13
- [18] 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. 13
- [19] 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. 2, 11
- [20] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 15

- [21] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. **11**
- [22] Ilya Loshchilov and Frank Hutter. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017. **11**
- [23] Rose M Rustowicz, Robin Cheong, Lijing Wang, Stefano Ermon, Marshall Burke, and David Lobell. Semantic segmentation of crop type in africa: A novel dataset and analysis of deep learning methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 75–82, 2019. **3, 4**
- [24] 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. **14**
- [25] 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. **14**
- [26] 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. **1, 2**
- [27] 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. **1, 2**
- [28] 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. **3, 4**
- [29] 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. **1, 2, 3, 4**
- [30] 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. **13**
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. **14**
- [32] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. **6**
- [33] Jamie Sherrah. Fully convolutional networks for dense semantic labelling of high-resolution aerial imagery. *arXiv preprint arXiv:1606.02585*, 2016. **2, 11**
- [34] 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. **14**
- [35] 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. **16**
- [36] 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. **11, 13, 14**
- [37] 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. **5, 13**
- [38] Michail Tarasiou, Erik Chavez, and Stefanos Zafeiriou. Vits for sits: Vision transformers for satellite image time series. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10418–10428, 2023. **3, 4**
- [39] 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, 11, 12, 14**
- [40] Xin-Yi Tong, Gui-Song Xia, and Xiao Xiang Zhu. Enabling country-scale land cover mapping with meter-resolution satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 196:178–196, 2023. **2**
- [41] 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. **11, 13, 14**
- [42] Yi Wang, Conrad M Albrecht, Nassim Ait Ali Braham, Chenying Liu, Zhitong Xiong, and Xiao Xiang Zhu. Decur: decoupling common & unique representations for multimodal self-supervision. *arXiv preprint arXiv:2309.05300*, 2023. **16**
- [43] 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. **14, 16**

- [44] Zhirui Wang, Xuan Zeng, Zhiyuan Yan, Jian Kang, and Xian Sun. Air-polsar-seg: A large-scale data set for terrain segmentation in complex-scene polsar images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2022. [2](#)
- [45] 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. [14](#)
- [46] 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. [1](#), [3](#), [4](#), [11](#)
- [47] Long Wen, Yu Cheng, Yi Fang, and Xinyu Li. A comprehensive survey of oriented object detection in remote sensing images. *Expert Systems with Applications*, page 119960, 2023. [13](#)
- [48] 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. [1](#), [2](#), [6](#), [14](#)
- [49] 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. [13](#)
- [50] Pengchuan Zhang, Xiyang Dai, Jianwei Yang, Bin Xiao, Lu Yuan, Lei Zhang, and Jianfeng Gao. Multi-scale vision long-former: A new vision transformer for high-resolution image encoding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2998–3008, 2021. [11](#)