

# CromSS: Cross-Modal Pretraining With Noisy Labels for Remote Sensing Image Segmentation

Chenyang Liu<sup>1</sup>, Student Member, IEEE, Conrad M. Albrecht<sup>2</sup>, Member, IEEE,  
Yi Wang<sup>1</sup>, Student Member, IEEE, and Xiao Xiang Zhu<sup>3</sup>, Fellow, IEEE

**Abstract**—We explore the potential of large-scale noisily labeled data to enhance feature learning by pretraining semantic segmentation models within a multimodal framework for geospatial applications. We propose a novel cross-modal sample selection (CromSS) method, a weakly supervised pretraining strategy designed to improve feature representations through cross-modal consistency and noise mitigation techniques. Unlike conventional pretraining approaches, CromSS exploits massive amounts of noisy and easy-to-come-by labels for improved feature learning beneficial to semantic segmentation tasks. We investigate middle and late fusion strategies to optimize the multimodal pretraining architecture design. We also introduce a cross-modal sample selection module to mitigate the adverse effects of label noise, which employs a cross-modal entangling strategy to refine the estimated confidence masks within each modality to guide the sampling process. Additionally, we introduce a spatial-temporal label smoothing technique to counteract overconfidence for enhanced robustness against noisy labels. To validate our approach, we assembled the multimodal dataset, *NoLDO-S12*, which consists of a large-scale noisy label subset from Google’s Dynamic World (DW) dataset for pretraining and two downstream subsets with high-quality labels from Google DW and OpenStreetMap (OSM) for transfer learning. Experimental results on two downstream tasks and the publicly available *DFC2020* dataset demonstrate that when effectively utilized, the low-cost noisy labels can significantly enhance feature learning for segmentation tasks.

Received 24 October 2024; revised 14 January 2025 and 10 March 2025; accepted 11 March 2025. Date of publication 18 March 2025; date of current version 3 April 2025. This work was supported by the Helmholtz Association’s Initiative and Networking Fund on the HAICORE@FZJ Partition. The work of Chenyang Liu, Conrad M. Albrecht, and Yi Wang was supported by the Helmholtz Association through the Framework of HelmholtzAI under Grant ZT-I-PF-5-01-Local Unit Munich Unit @Aeronautics, Space and Transport (MASTR). The work of Chenyang Liu and Xiao Xiang Zhu was supported by Munich Center for Machine Learning (MCML). The work of Conrad M. Albrecht was supported by the European Union’s Horizon Europe Research and Innovation Program under Grant 101082130 (EvoLand). The work of Xiao Xiang Zhu was supported by the German Federal Ministry of Education and Research (BMBF) in the Framework of the International Future AI Laboratory “AI4EO—Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond” under Grant 01DD20001. (Corresponding author: Chenyang Liu.)

Chenyang Liu is with the Chair of Data Science in Earth Observation, Technical University of Munich (TUM), 80333 Munich, Germany, also with the Remote Sensing Technology Institute, German Aerospace Center (DLR), 82234 Weßling, Germany, and also with Munich Center for Machine Learning (MCML), 80538 Munich, Germany (e-mail: chenyang.liu@tum.de).

Conrad M. Albrecht is with the Remote Sensing Technology Institute, DLR, 82234 Weßling, Germany.

Yi Wang is with the Chair of Data Science in Earth Observation, TUM, 80333 Munich, Germany.

Xiao Xiang Zhu is with the Chair of Data Science in Earth Observation, TUM, 80333 Munich, Germany, and also with MCML, 80538 Munich, Germany.

Digital Object Identifier 10.1109/TGRS.2025.3552642

The data, codes, and pretrained weights are freely available at <https://github.com/zhu-xlab/CromSS>.

**Index Terms**—Geospatial artificial intelligence, multimodal deep learning, noisy labels, pretraining, sample selection, semantic segmentation, SSL4EO-S12 dataset.

## I. INTRODUCTION

**I**N SPITE of the effectiveness of deep learning in big data processing [1], obtaining enough accurate labeled data for deep learning model training is challenging due to the high cost of the annotation process [2]. Self-supervised learning (SSL) is promising to address this issue by means of learning hidden representations from unlabeled data [3], [4]. Popular SSL algorithms include generative masked autoencoders (MAEs) [5] and contrastive learning methods such as DINO [6] and MoCo [7]. MAE is inspired by image reconstruction, as most works utilizing vision transformers (ViTs) [8]. When applying MAE to remote sensing (RS), RingMo introduced a patch-incomplete strategy to reduce the loss of dense and small objects for masked image modeling (MIM) pretraining [9]. SatMAE employed extra temporal embeddings and spectral positional encodings to explore temporal information in multispectral images [10], which is improved with multiscale pretraining by satMAE++ [11]. Scale-MAE improved MIM toward better scale awareness [12]. On the other hand, contrastive learning methods, which push the semantically similar input patches to be close in representation space, can make a difference for both convolutional backbones like ResNets [13] and ViTs. To this end, the data augmentation strategies used to generate positive and negative samples are key. In the RS domain, such strategies are designed with geolocations [14], temporal or seasonal information [15], multisize cropping [16], and so on [17]. However, these methods primarily target image-level information, e.g., for classification tasks, leading to suboptimal performance when applied to tasks requiring pixel-level information such as semantic segmentation [18]. This discrepancy motivates us to explore alternative strategies to enhance the efficacy of pretrained models for dense mapping tasks like segmentation.

Before the adoption of SSL, pretrained models were dominantly trained on large-scale annotated datasets in computer vision, such as ImageNet [19]. Recently, the segment anything model (SAM) series [20], [21] has achieved remarkable success in image segmentation by training on

millions of labeled images and videos, demonstrating robust generalization capabilities for diverse downstream tasks like underwater image segmentation [22]. Similarly, tailored to RS image understanding, the SatlasPretrain dataset was developed to support large-scale pretraining on RS data. However, such vast amounts of accurately annotated labels are rare, driving researchers to explore sources for “weak” labels, which are noisy but much lower in cost and easy to obtain on a large scale. Recent studies have demonstrated the resilience of deep learning models to a certain degree of label noise [23], [24], showing that the deep learning models trained on large-scale noisy labels can exhibit strong feature learning capabilities and transfer learning performance across various applications. For instance, noisy social-media labels have been successfully used for image classification [25] and video analysis [26]. In RS, crowd-sourced maps such as OpenStreetMap (OSM) have been explored to pretrain models for tasks like building and road extraction [27]. Notably, Maggiori et al. [28] observed that the shallow layers of a network pretrained with noisy labels remain largely unchanged after fine-tuning. This implies that the resilience to label noise varies across different parts of a deep learning model, with shallow layers particularly adept at learning robust features even when trained on noisy labels. In our previous work [29], we analyzed the segmentation model’s (U-Net) robustness in a per-module fashion to further qualify this phenomenon. We compared two U-Nets trained with exact ground-truth (GT) and noisy labels. We observe the following: 1) the encoder features visually share similar spatial characteristics, in which case the closer the convolutional layer gets to the U-Net’s output, the more the features become contaminated by label noise and 2) similar weight statistics govern the encoder layers, while those in the decoder follow diverging weight statistics toward the semantic segmentation outputs. These observations highlight that encoder features are less biased by label noise, yet they benefit from the semantics provided by pixel-level noisy label masks. We attribute this phenomenon to encoders’ preference to learn basic spatial features more from the input data than from the labels. From this perspective, transferring such weakly supervised pretrained encoder weights to semantic segmentation downstream applications could be highly beneficial.

On the other hand, we have observed a surging tendency in multimodal learning, exemplified by vision-language learning such as image captioning [30] and visual question answering (VQA) [31]. In the RS domain, a plethora of satellite data modalities is becoming accessible: from multispectral, hyperspectral, to synthetic aperture radar (SAR), to light detection and ranging (LiDAR) data, and many more [32], [33], [34], [35], [36]. Different modalities can provide complementary information to each other, such as the radiation features of optical images and the structural features of SAR images. This cross-modal enhancement can potentially boost the mutual learning between modalities, even in SSL for RS data. For example, CROMA extends MAE from a single modality to multiple modalities alongside cross-modal contrastive learning [37]. SkySense, based on MAE, too, features a modular design accommodating diverse tasks and modalities [38]. DeCUR defines a contrastive loss to explicitly model features as shared or orthogonal among modalities [39]. DOFA, leveraging the concept of neural plasticity in brain

science, integrates various data modalities into a single framework by adjusting to different wavelengths [40]. These researches demonstrate that multimodal learning is beneficial in feature learning and, thus, is promising to make a difference in the noisy label pretraining framework.

Multimodal learning mainly includes two kinds of strategies: data fusion and co-learning. Data fusion is to add or concatenate data/features from different modalities in the input, feature, or decision levels [41]. This strategy usually requires all the modalities in both the training and inference stages. As a more flexible solution, co-learning leverages losses to integrate knowledge from different modalities by enforcing consistency between their outputs to enhance model training [42]. In this case, each modality has specific models and can be employed separately after training. Therefore, we adopt the co-learning framework for multimodal noisy label pretraining to maintain this flexibility in transfer learning settings. This structure is analogous to multimodal sample selection methods of learning from noisy labels (LNLs), such as coteaching [43] and decoupling [44], where two identical models are trained separately and select samples for each other to avoid overfitting to label noise. The basic idea behind this is that the samples with noisy labels have higher inconsistency than correctly annotated pixels during the training [45]. We expect some similar behaviors in the co-learning noisy label pretraining framework. The diverse features from different modalities can lead to different responses to label noise, making it possible to integrate some sample selection strategies to remove part of the reverse effects of label noise.

In this work, we introduce a novel cross-modal sample selection method, referred to as CromSS, to explore the potential of using large-scale noisy labels and multimodal learning to enhance the feature learning capabilities of RS image semantic segmentation models. We nest CromSS into a co-learning framework to maintain the flexibility of pretrained models when transferred to downstream tasks. To this end, we employ two U-Nets [46] (using ResNet-50 as backbones [13]) to separately extract features and make predictions for each modality, namely optical and SAR. Except for applying segmentation losses to each model, we push the predictions to be close with additional consistency losses between their softmax outputs. We also test an implicit consistency constraint of utilizing shared weights for the two decoders. Following the naming of data fusion, we name the pure co-learning architecture *late fusion* since the two models are only forced to mimic each other at the very end of the outputs. We call the shared decoder architecture *middle fusion* as the implicit constraint is additionally put on the feature level. Notice that fusion happens in a hidden way in our work instead of data addition or concatenation. We include the two architectures without explicit preferences, given that no evidence shows which one is superior. Our experiments provide some references to readers regarding the performance of the two architectures. Furthermore, we introduce a sample selection strategy within the multimodal framework encouraged by LNL multimodal-based sample selection methods to alleviate the adverse effects of label noise. We weight the loss functions with the sample selection masks obtained in three steps: confidence mask generation, cross-modal confidence enhancement, and

thresholding. We first estimate the two kinds of confidence, label-based and entity-based, for segmentation and consistency losses. Then, the confidence masks are enhanced by fortifying the shared information across modalities. Two thresholding strategies are applied to each enhanced confidence mask to derive final selection masks. Here, we introduce a spatial-temporal label smoothing technique tailored for the used SSL4EO-S12 dataset with four timestamps of each given geospatial area (patch) to replace the uniform distribution-based label smoothing to ensure robustness in pretraining.

In our experiments, we utilize Sentinel-1 (S1) of SAR and Sentinel-2 (S2) of multispectral data from the SSL4EO-S12 dataset [18] as two modalities. We propose a new dataset extended from SSL4EO-S12, called *NoLDO-S12*, which pairs Sentinel-1 and 2 data with noisy labels sourced from the Google Dynamic World (DW) project [47] for pretraining, alongside two land cover/land use segmentation downstream tasks with GT labels sourced from DW and OSM, to evaluate the proposed method. We also involve another publicly available dataset DFC2020 [48] to test the transferability of the model. Experimental results indicate the effectiveness of CromSS with its task-specific pretraining for RS image segmentation.

In summary, our main contributions are as follows.

- 1) The development of a novel, task-specific noisy label pretraining method called CromSS, which is proposed to explore the potential of noisy labels in feature learning for RS image segmentation, see Section III. CromSS is spotlighted with two characteristics. One is the multimodal framework nesting CromSS to boost feature learning with noisy labels. The other is the sample selection strategy based on the enhanced confidence masks by fortifying the shared information across modalities to enhance the performance by excluding partial label noise.
- 2) The curation of a multimodal dataset named NoLDO-S12 (Section II), which is composed of a pretraining subset with pixel-wise noisy labels and two segmentation downstream tasks for land cover/land use classification with Sentinel-1/2 images as inputs, see Section IV. The noisy label pretraining subset is the first global dataset for this purpose. The two downstream tasks complement existing land cover/land use classification datasets based on Sentinel data.
- 3) The experimental evaluation of the proposed CromSS compared with other state-of-the-art pretraining strategies on three different RS image segmentation downstream tasks, see Section IV. Furthermore, we analyzed the learned features by different pretraining methods to shed light on the future design of pretraining and feature learning methodologies.

## II. NoLDO-S12 DATASET

NoLDO-S12 contains two splits: SSL4EO-S12@NoL with noisy labels for pretraining based on the SSL4EO-S12 dataset [18], and two downstream datasets, SSL4EO-S12@DW and SSL4EO-S12@OSM, for transfer learning with exact labels. We detail them in Sections II-A and II-B, respectively.

### A. Pretraining Task

The pretraining dataset SSL4EO-S12@NoL extends the large-scale, multimodal, and multitemporal self-supervision dataset SSL4EO-S12 [18]. SSL4EO-S12 contains 251kx4 globally sampled Sentinel-1 (S1, radar) and Sentinel-2 (S2, optical) image pairs of  $264 \times 264$  pixels from four seasons. In our case, we employ S1 ground range detected (GRD) SAR images of two bands, and S2 level-1C (L1C) multispectral data of 13 bands. All bands are resampled to a spatial resolution of 10 m. We paired these data with the nine-class noisy labels sourced from the DW project on Google Earth Engine. DW generates noisy land cover labels with semisupervised deep learning based on the Sentinel-2 L1C product. As per [47], the labels bear an overall accuracy (OA) of approximately 70%. We picked the DW project for the following reasons: the nine-class schema aligns with other global land cover land use maps such as ESA's *Land Use and Coverage Area frame Survey* (LUCAS) [49]. DW offers near real-time mapping providing multitemporal label masks. As a result, SSL4EO-S12@NoL contains about 41% (103793 of the 251079 locations) of the SSL4EO-S12 dataset with noisy label masks. To keep the dataset's multitemporal characteristics, we only retain the (S1, S2, noisy label) triples from the locations where all four timestamps of (S1, S2) pairs in SSL4EO-S12 have corresponding DW labels, as illustrated in Fig. 1. This dataset well reflects real-world use cases where noisy labels remain more difficult to obtain compared to bare (S1, S2) image pairs.

### B. Downstream Tasks

We construct two downstream datasets, SSL4EO-S12@DW and SSL4EO-S12@OSM. Both are selected on the DW project's manually annotated training and validation datasets. However, they are paired with different label sources from DW and OSM. For DW, the human labeling process allows some ambiguous areas left unmarked (white spots in DW masks in Fig. 2). We exclusively utilize the expert labeled training subset of 4194 tiles with given dimensions of  $510 \times 510$  pixels (training data)<sup>1</sup> to construct our downstream tasks. The DW hold-out validation set of 409 tiles with given dimensions of  $512 \times 512$  is used to construct our test data.<sup>2</sup> The validation set was geographically sampled to ensure that its geographical distribution aligns with that of the training set. We spatially-temporally aligned the S1 and S2 data for the training and test tiles with Google Earth Engine. Consequently, 3574 training tiles and 340 test tiles were successfully paired. These serve as the first downstream task in SSL4EO-S12@DW with a total of 656 758 064 training pixels and 60 398 506 test pixels. The class distributions can be found in Fig. 2. SSL4EO-S12@DW test set has a slightly different distribution from the training set. We consider it an acceptable difference inherently caused by geospatial sampling of the splits.

To test the generability of pretrained models, we construct another downstream dataset with fine-grained labels derived

<sup>1</sup>The DW training labels were downloaded from <https://doi.pangaea.de/10.1594/PANGAEA.933475> (Expert directory).

<sup>2</sup>The DW hold-out validation set is built from the annotations where at least two of the three domain experts agree on the labels, *Expert Consensus*, downloaded from <https://zenodo.org/records/4766508>

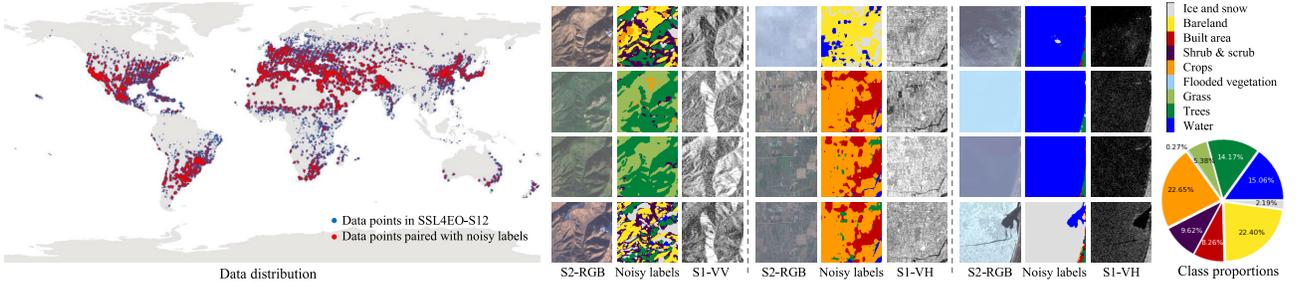


Fig. 1. Illustration of the pretraining set SSL4EO-S12@NoL in NoLDO-S12. (From left to right) Global distribution of samples (left), four-season samples (top-down) at three geolocations (middle), and statistics of the classes of the noisy labels (right).

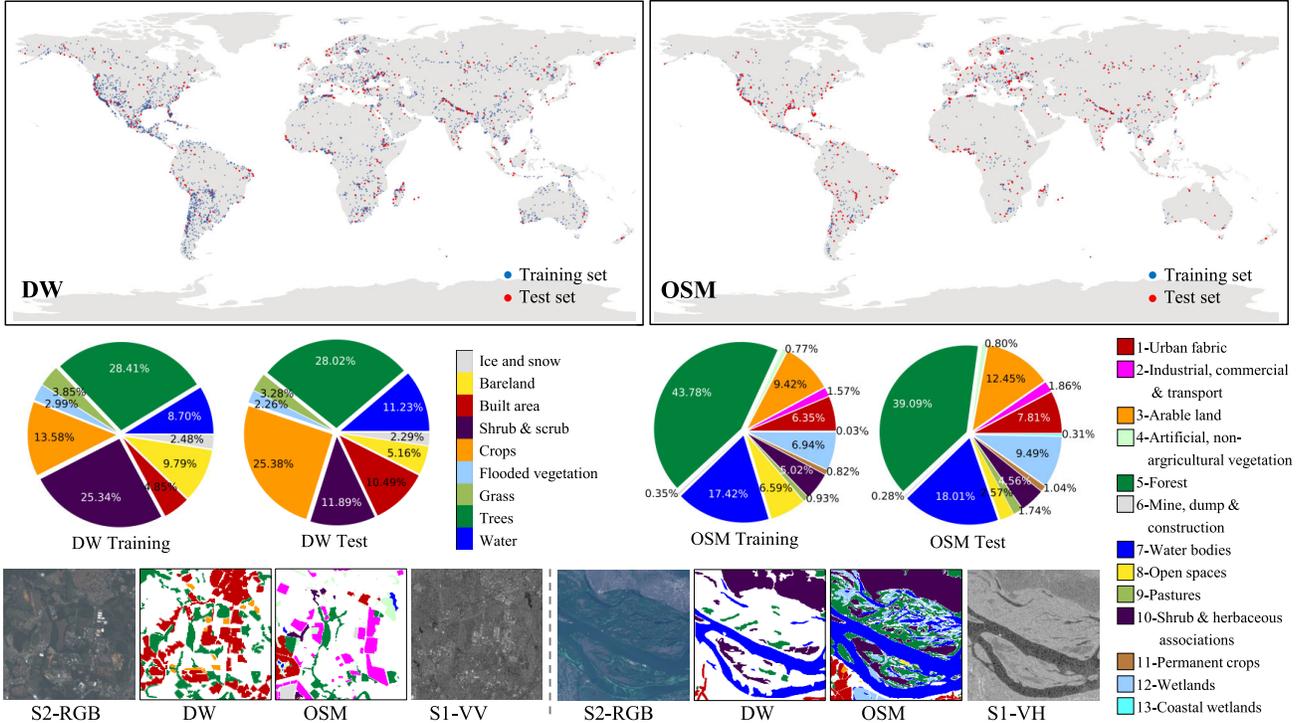


Fig. 2. Illustration of the two downstream tasks in NoLDO-S12 with different label sources (SSL4EO-S12@DW and SSL4EO-S12@OSM). Top (left and right) Global data distributions (DW and OSM). Middle (left and right) Class distributions of training and test sets along with corresponding legends (DW and OSM). (Bottom) Examples from two locations. The legend for DW labels is the same as that in Fig. 1.

TABLE I  
MANY-TO-MANY RELATION OF LAND COVER LABELS FROM THE OSM TO THE DW DATASETS

OSM label ID(s)	DW class names	OSM label ID	DW class names	OSM label ID	DW class names
1, 2	Built Area	3	Crops	4	Trees, Grass
5	Trees	6	Bare Ground	7	Water
8	Bare Ground, Snow & Ice	9	Grass	10	Grass, Shrub & Scrub
11	Trees, Crops	12	Flooded vegetation	13	Water, Flooded Vegetation

from OSM.<sup>3</sup> We follow the work of Schultz et al. [50], adopting a 13-class schema similar to the classification level 2 of the CORINE Landcover classes [51], as listed in Fig. 2. We retrieved 2996 OSM label masks among the  $3914 = 3574 + 340$  DW tiles. The remaining subset of 918 tiles is left without OSM labels. To keep the quality of OSM labels, we conduct an automatic check with DW labels as reference.

<sup>3</sup>OSM labels were downloaded from the OSMLanduse layer at <https://osmlanduse.org/#12/8.7/49.4/0/> via the Web Map Service (WMS) at <https://maps.heigit.org/osmlanduse/service>

By definition, we construct a many-to-many mapping of class schema from OSM (integers) to DW in Table I. Treating DW labels as GT and leveraging the OSM-DW relation in Table I, for each of the  $N$  labeled pixels of a tile, we check whether an OSM label agrees with any of the corresponding DW labels to compute the OA  $\text{acc} = \text{TP}/N$ . TP denotes all true positive matches according to Table I. Filtering for  $\text{acc} > .75$ , a total of 1772 tiles remain in the SSL4EO-S12@OSM downstream subset, with 1372 for training and 400 for test. Due to filtering, we note that only 106 of the 340 SSL4EO-S12@DW test tiles remain ( $512 \times 512$  pixels per tile). To substitute for

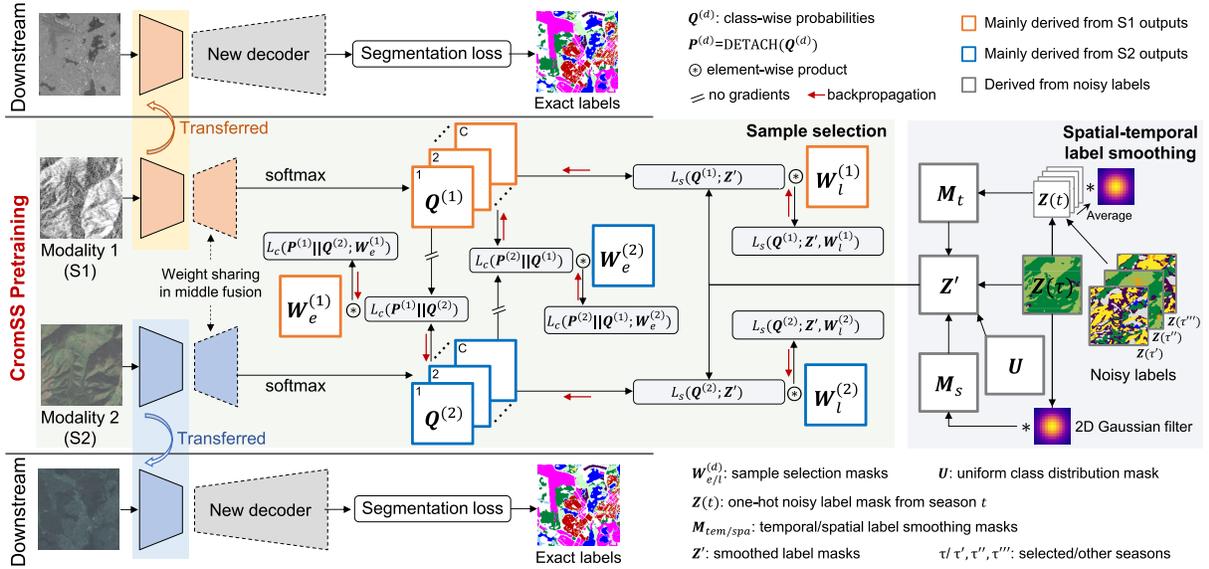


Fig. 3. Overview of the proposed CromSS method.

the reduction in tiles, we utilize additional 294 tiles from the SSL4EO-S12@DW training dataset ( $510 \times 510$  pixels per tile). For reasons of consistency, the 294 test tiles of size  $510 \times 510$  were padded by a one-pixel zero values boundary to increase them to a size of  $512 \times 512$  pixels. All our (convolutional) semantic segmentation models tested in Section IV (and many more from the literature) are able to consume  $512 \times 512$  pixel images as input. We implemented additional manual checks and corrections by human inspection on the test set to further improve the reliability of test results. The total numbers of training and test pixels are 165 993 707 and 44 535 192, respectively. Please refer to the upper right charts in Fig. 2 for additional data statistics of the SSL4EO-S12@OSM downstream dataset.

### III. METHODOLOGY

As illustrated in Fig. 3, the proposed CromSS is a multimodal weakly supervised pretraining method with noisy labels for RS image segmentation tasks. The pretrained encoders by CromSS can then be transferred to different segmentation downstream tasks in combination with various decoders. Overall, CromSS is nested into a co-learning framework using consistency losses to boost the feature learning across modalities in addition to segmentation losses for pretraining with noisy labels. As demonstrated in the middle-left part of Fig. 3, two models specific to each modality are employed to separately extract features and make predictions before connecting via consistency losses. We can optionally add an explicit cross-modal constraint by using shared weights for decoders. We name the two architectures *late fusion* and *middle fusion* regarding whether to include the explicit constraint. We present the details of architecture design in Section III-A. As highlighted in shadowed area (on the left in the middle), another spotlight of CromSS is the cross-modal sample selection strategy. We implement it by introducing weighting masks, i.e., selection masks ( $W_{l/e}$ ), into segmentation/consistency loss function calculation to attenuate

the impact of label noise. We employ a three-step strategy to generate these selection masks, which will be elaborated in Section III-B. Moreover, we include a spatial-temporal label smoothing technique to replace uniform distribution-based label smoothing to ensure robustness in pretraining. We briefly describe this small trick in Section III-C.

#### A. Multimodal Pretraining Architectures

We employ middle and late fusion architectures to explore complementary information across modalities to assist noisy label pretraining. Notice that our fusion strategies are implemented in an implicit way without adding or concatenating feature vectors of different modalities. As illustrated in Fig. 4, the major difference between multimodal and single-modal pretraining with noisy labels is the incorporation of additional consistency losses between the outputs from the two models. While late fusion retains individual decoders, middle fusion shares a common decoder for the two modalities. The loss function for multimodal pretraining with noisy labels can be written as follows:

$$L = L_s(\mathbf{Q}^{(1)}|\mathbf{Y}) + L_s(\mathbf{Q}^{(2)}|\mathbf{Y}) + L_c(\mathbf{P}^{(1)}\|\mathbf{Q}^{(2)}) + L_c(\mathbf{P}^{(2)}\|\mathbf{Q}^{(1)}). \quad (1)$$

Specifically, we utilize a combined segmentation loss summing the distribution-based CrossEntropy loss  $L_{ce}$  and the region-based Dice loss  $L_{dice}$  [52] to represent

$$L_s(\mathbf{Q}^{(d)}|\mathbf{Y}) = L_{ce}(\mathbf{Q}^{(d)}|\mathbf{Y}) + L_{dice}(\mathbf{Q}^{(d)}|\mathbf{Y}) \quad (2)$$

with

$$L_{ce}(\mathbf{Q}^{(d)}|\mathbf{Y}) = \left( - \sum_{i,c} (z)_{ic} \log(q_{ic}^{(d)}) \right) \quad (3)$$

and

$$L_{dice}(\mathbf{Q}^{(d)}|\mathbf{Y}) = \left( 1 - 2 \frac{\sum_{i,c} z_{ic} q_{ic}^{(d)}}{\sum_{i,c} z_{ic} + q_{ic}^{(d)}} \right) \quad (4)$$

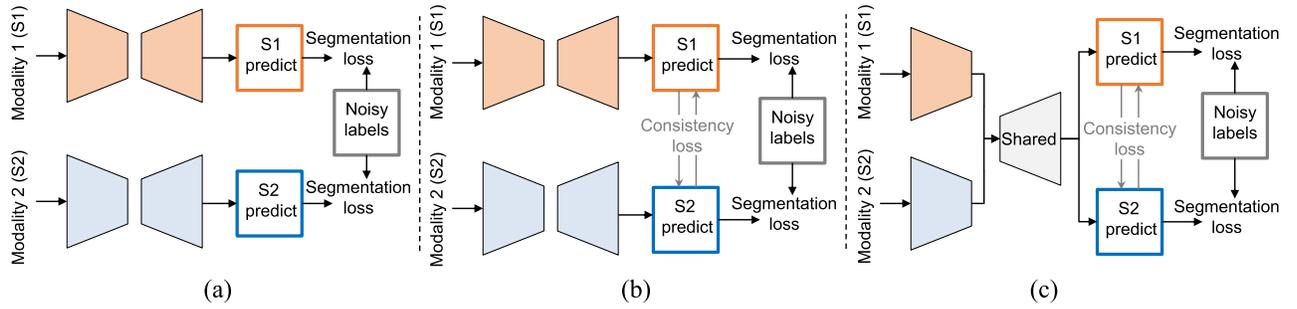


Fig. 4. Comparison of single-modal training and multimodal (middle/late) fusion strategies. (a) Single modality. (b) Multimodal: late fusion. (c) Multimodal: middle fusion.

where  $z_{ic} \in \{0, 1\}$  is the  $c$ th entry of the one-hot encoded noisy label vector  $\mathbf{z}_i$  for the  $i$ th noisy sample annotation  $y_i \in \mathbf{Y}$ . More specifically,  $z_{ic} = 1$  when  $c = y_i$ , and  $z_{ic} = 0$  otherwise. The class index  $c \in \mathbb{N}$  ranges from 1 to  $C$ , the total number of classes.  $N$  is the number of image pixels with the index  $i = 1, \dots, N$ , where  $N = W \times H$ ,  $W$  and  $H$  are the image width and height, respectively. The symbol  $q_{ic}^{(d)} \in \mathbf{Q}^{(d)}$  denotes the corresponding model prediction of modality  $d \in \{1, 2\}$  from the *softmax* layer. For the consistency loss  $L_c$ , we pick the Kullback–Leibler (KL) divergence

$$L_c(\mathbf{P}^{(d)} \parallel \mathbf{Q}^{(d')}) = - \sum_{i,c} p_{ic}^{(d)} \log \left( \frac{q_{ic}^{(d')}}{p_{ic}^{(d)}} \right) \quad (5)$$

where  $d \neq d' \in \{1, 2\}$  are two modalities, and  $p_{ic}^{(d)}$  indicates the detached version of  $q_{ic}^{(d)}$  without gradients in the backpropagation, see Fig. 5.

In CromSS, we utilize sample selection masks  $\mathbf{W}_{l/e}^{(d)}$  (details in Section III-B) to assign smaller weights to the samples possibly associated with noisy labels. In practice, we employ the weighted average strategy to aggregate loss values from each sample (see the right part in Fig. 5). As a consequence, (2) and (5) turn to

$$L_s(\mathbf{Q}^{(d)} | \mathbf{Z}', \mathbf{W}_l^{(d)}) = - \sum_{i,c} w_{l,i}^{(d)} z'_{ic} \log q_{ic}^{(d)} + \left( 1 - 2 \frac{\sum_{i,c} w_{l,i}^{(d)} z'_{ic} q_{ic}^{(d)}}{\sum_{i,c} w_{l,i}^{(d)} (z'_{ic} + q_{ic}^{(d)})} \right) \quad (6)$$

and

$$L_c(\mathbf{P}^{(d)} \parallel \mathbf{Q}^{(d')} | \mathbf{W}_e^{(d)}) = - \sum_{i,c} w_{e,i}^{(d)} p_{ic}^{(d)} \log \left[ q_{ic}^{(d')} / p_{ic}^{(d)} \right] \quad (7)$$

where  $w_{l/e,i}^{(d)} \in \mathbf{W}_{l/e}^{(d)}$  is the weight for sample  $i$ , and  $z'_{ic} \in \mathbf{Z}'$  is the smoothed label from the one-hot counterpart  $z_{ic}$  for sample  $i$  in class  $c$ .

### B. Sample Selection Mask Generation

In this section, we present the details of how to generate sample selection masks  $\mathbf{W}_{l/e}^{(d)}$  to weigh losses. As depicted by Fig. 5, the selection mask generation process is mainly composed of three steps: confidence mask generation, cross-modal confidence enhancement, and thresholding. The second

step, cross-modal confidence enhancement, is the same for both losses. In the first and third steps, we estimate the confidence values and apply different thresholding strategies for the two kinds of losses separately. The details of each step are given as follows.

In the first step, we generate the label-based and the entity-based confidence masks for the segmentation and the consistency losses, respectively. The segmentation loss is used to supervise the model to predict the pixel-wise classification results of an image in the given classes. It quantifies the difference between the predicted segmentation map and the GT. In our case, we do not have the real GT at hand. The labels we used for optimization contain noise or wrong annotations. We expect only the correctly assigned labels to contribute to the loss. That is, the confidence values should reflect the reliability of the labels. To this end, we utilize the softmax outputs at the given annotation class for each sample to generate the label-based confidence mask  $\mathbf{F}_l$  as follows:

$$f_{l,i}^{(d)} = p_{ic=y_i}^{(d)} \quad (8)$$

where  $f_{l,i} \in \mathbf{F}_l$  is the estimated label-based confidence score at the  $i$ th pixel, and  $y_i$  is the given noisy label class.

On the other hand, the consistency loss, in the absence of labels, encourages the models to extract similar features from both modalities for the same scene. We hope the models can learn more from each other's reliable features, namely, the less ambiguous predictions. To achieve this, we estimate entity-based confidence mask  $\mathbf{F}_e^{(d)}$  using the entropy  $H_i^{(d)} \in \mathbb{R}^+$  of the softmax output vectors as follows:

$$f_{e,i}^{(d)} = 1 - H_i^{(d)} / K = 1 + \frac{1}{K} \sum_c p_{ic}^{(d)} \log p_{ic}^{(d)} \quad (9)$$

where  $f_{e,i} \in \mathbf{F}_e$  is the estimated entity-based confidence score at the  $i$ th pixel, and  $K = \log C$  is the upper bound of the entropy  $H_i \in [0, K]$  when  $p_{ic} = 1/C$  for  $c = 1, \dots, C$ , i.e., equal distribution yields maximum entropy.

Then, we leverage cross-modal information to enhance the confidence masks of individual modalities before generating the final selection mask. Observations reveal that while a model from modality  $d = 1$  may exhibit high confidence in a label  $y$  or prediction  $\mathbf{p}$ , another model from modality  $d = 2$  might assign a low probability to  $y$  or express less confidence in  $\mathbf{p}$ . This discrepancy arises due to differences in modality characteristics and the inherent randomness in identical training processes. To address this, we reinforce consistent preferences between the two modalities and

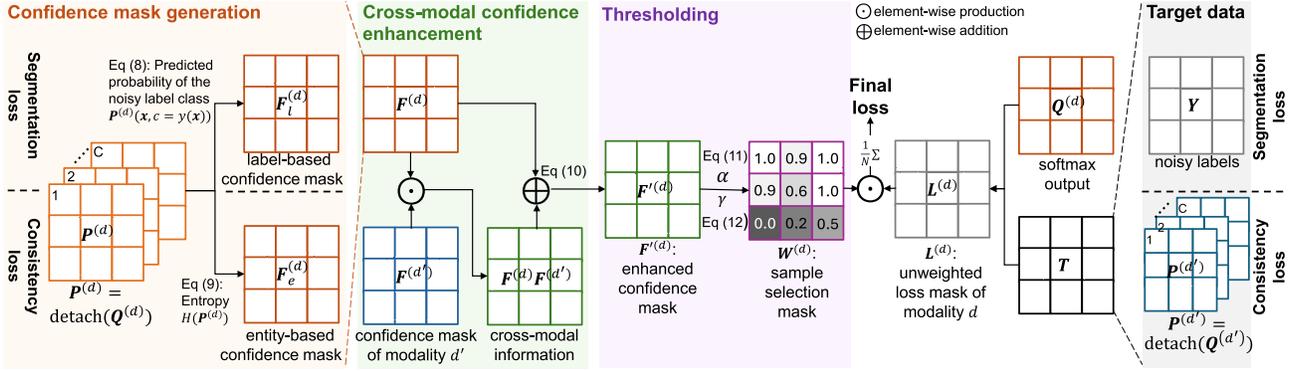


Fig. 5. Sample selection mask generation in three steps: confidence mask generation, cross-modal confidence enhancement, and thresholding, where  $d$  and  $d'$  represent two different modalities, and  $\alpha$  and  $\gamma$  are the thresholds used to derive the sample selection masks for the segmentation and consistency losses, respectively.

penalize divergent outcomes in the enhanced confidence masks by

$$F_{l/e}^{(d)} = \frac{1}{2} \left( F_{l/e}^{(d)} + F_{l/e}^{(1)} F_{l/e}^{(2)} \right) = \frac{1}{2} F_{l/e}^{(d)} \left( 1 + F_{l/e}^{(d')} \right) \quad (10)$$

where the factor  $F_{l/e}^{(1)} F_{l/e}^{(2)}$  serves to magnify the selection probabilities of the samples exhibiting high confidences for both modalities, while diminishing cases where both modalities agree on low confidence scores, and  $d' = \{1, 2\}/d$ . When  $F^{(1)}$  and  $F^{(2)}$  values vastly differ, or both exhibit values close to  $1/2$ , the selection probabilities tend to stay closer to  $3/8 \approx 1/2$ .<sup>4</sup>

Finally, we apply thresholding to the enhanced confidence masks to derive sample selection masks for loss weighing, with  $\alpha$  and  $\gamma$  denoting the two hyperparameters used for the segmentation and the consistency losses, respectively. Specifically, we define  $\alpha \in [0, 1]$  as the selection ratio to pick out the most reliable labels. We sort the label-based confidence values in the descending order, and assign the weights to the samples with the  $\alpha$  largest confidence values to 1. Notice that our goal is to pretrain the encoders for segmentation tasks. The encoders themselves can combat label noise to some extent [29]. As a result, we use soft weights ranging from 0 to 1 for those unselected ones to avoid the removal of valuable samples with hard weights of 0. The thresholding for  $F_l^{(d)}$  is formulated as follows:

$$w_{l,i}^{(d)}(\alpha) = \min \left[ 1, f_{l,i}^{(d)} / f_l^{(d)}(\alpha) \right] \quad (11)$$

where  $W_l$  is the label-based sample selection mask with  $w_{l,i} \in W_l$ , and the quantity  $f_l^{(d)}(\alpha)$  is the  $\lfloor \alpha N \rfloor$ th largest value<sup>5</sup> in  $F_l^{(d)}$ . Models' initial predictions are generally less accurate in the first couple of training epochs. Therefore, we set  $\alpha$  gradually decreasing from 1 to the final target selection ratio  $\alpha_0 > 0$  in the first  $n_s$  epochs.

<sup>4</sup>Define  $f^\pm = 1/2 \pm \epsilon$  with  $\epsilon \in [0, 1/2]$ , it follows for the quantity  $f^{\pm\pm} = (1/2)(f^\pm + f^\pm f^\pm) = 3/8 \pm \epsilon/2 + \mathcal{O}(\epsilon^2)$  such that  $f^{\pm\pm} \approx 3/8 \pm \epsilon/2$  with  $f^{\pm\pm} \approx 5/8$  or  $1/8$  when  $\epsilon = 1/2$  (the values vastly differ), and  $f^{\pm\pm} \approx 3/8$  when  $\epsilon = 0$  (both values are equal to  $1/2$ ), implying  $f^\pm \in [0, 1] \rightarrow f^{\pm\pm} \in [1/8, 5/8]$ .

<sup>5</sup>For practical purposes, it turned out that implementing the  $\alpha$ -dependent thresholding better performs when applied per class, that is, we utilize  $\lfloor \alpha N_c \rfloor$  with  $N = \sum_c N_c$ , where  $N_c$  represents the number of samples per class  $c$ . Correspondingly, each  $f_{l,i}^{(d)}$  gets divided by a class-specific  $w_l(\alpha)$  whenever the index  $i = 1, \dots, N$  denotes a pixel of class  $c = y_i$ .

The consistency loss uses the predicted class distributions from another modality as the target data, which can be regarded as a kind of soft label. In this sense, we keep the soft weight format for the entity-based selection masks. Like (11), we utilize a moderator  $\gamma$  to mitigate the negative impact of poor estimations in the early training stage. The entity-based selection mask  $W_e$  is actually a weighted sum between 1 and the entity-based confidence mask

$$w_{e,i}^{(d)}(\gamma) = (1 - \gamma) + \gamma f_{e,i}^{(d)} \quad (12)$$

where  $w_{e,i} \in W_e$ , and  $\gamma$  ramps up from 0 to 1 in the first  $n_s$  epochs.

### C. Spatial–Temporal Label Smoothing

As shown in Fig. 3 (right), we integrate a spatial–temporal label smoothing technique into CromSS to mitigate overconfidence in model predictions. The classic label smoothing technique replaces the one-hot label vector with the weighted averaging of it and the class uniform distribution [53]. In image segmentation, spatial correlations may get exploited for smoothing: weighted averaging in space is another option to improve label smoothing. Additionally, our pretraining dataset SSL4EO-S12@NoL contains four timestamps of a given geospatial area (patch) with S1 and S2 data from different seasons. Therefore, we suggest enhancing label smoothing with spatial–temporal information for image segmentation as follows:

$$Z'(\tau) = (1 - \beta - \mu)Z(\tau) + \beta U + \mu(M_s(\tau) + M_t)/2 \quad (13)$$

where  $Z(\tau)$  is the one-hot-encoded noisy label mask from a randomly selected season at timestamp  $\tau$  in the data augmentation stage. The hyperparameters  $\beta$  and  $\mu$  control the proportions of the uniform class distribution  $U = \{1/C\}^{H \times W \times C}$  and the spatial–temporal smoothing masks. The spatial smoothing mask  $M_s$  is generated with Gaussian kernel  $G$  convolved with  $Z(\tau)$

$$M_s(\tau) = Z(\tau) * G. \quad (14)$$

The temporal smoothing mask  $M_t$  is obtained with  $G$  applied to the temporally averaged label mask from four seasons  $\langle Z \rangle_t = (1/4) \sum_{t' \in \{\tau, \tau', \tau'', \tau'''\}} Z(t')$ , i.e.,

$$M_t = \langle Z \rangle_t * G. \quad (15)$$

#### IV. EXPERIMENTS

We tested the proposed CromSS methodology by pretraining ResNet-50 encoders nested in U-Nets and transferring them to three segmentation downstream tasks. In the following, we present the settings of our experiments in Section IV-A, ahead of the detailed results obtained for the downstream tasks compared to the state-of-the-art pretrained models in Section IV-B. We explore the roles of each part in CromSS in Section IV-C. Some insights for the pretraining stage are given in Section IV-D. Finally, we discuss lessons learned when generating features by different pretraining methods in Section IV-E.

##### A. Setups

In the pretraining stage, we train U-Nets with an Adam optimizer [54] and an initial learning rate of  $5 \cdot 10^{-3}$ . We employed the `ReduceLRonPlateau` scheduler to dynamically adjust the learning rate. It cuts in half the learning rate when the validation loss is not decreasing over 30 consecutive epochs. We randomly split off 1% of the entire training set as the validation set. The selection proportion  $\alpha$  for generating  $W_l^{(d)}$  exponentially ramps down from 1 to  $\alpha_0 = 0.5$  for the initial 80 epochs. The weighting factor  $\gamma$  for  $W_e^{(d)}$  is ramping up from 0 to 1 in parallel. The model input data (Sentinel-2 L1C) is normalized to  $[0, 1]$  after discarding the values outside  $\langle x \rangle \pm 2\sigma(x)$  using the mean  $\langle x \rangle$  and variance values  $\sigma^2(x)$  of the SSL4EO-S12 dataset. The data augmentation strategies we use before feeding the data into the U-Net are the following:

- 1) random cropping to  $256 \times 256$ ;
- 2) random flipping with 0.5 probability;
- 3) random rotating at probability 0.2; and
- 4) random season selection which feeds the tiles of a randomly selected season.

We implemented pretraining on SSL4EO-S12@NoL with 4 NVIDIA A100 GPUs running approximately 22 h for 200 epochs. We configured a batch size of 128 for each GPU working in data-parallel mode, i.e., backpropagated gradients are averaged over a total batch size of 512 data samples.

When transferred to the downstream segmentation tasks, we embed the pretrained ResNet-50 encoders into different segmentation frameworks. We fine-tuned with the Adam optimizer given a combined loss of `CrossEntropy` and `Dice` at a learning rate of  $5 \cdot 10^{-4}$  over 50 epochs. Aside from the two downstream tasks in NoLDO-S12, we also incorporate the DFC2020 dataset [48] as an additional downstream task, where we utilize the provided validation set comprising of 986 patches as the fine-tuning training set, and the original test set consisting of 5128 patches for test. The three datasets (SSL4EO-S12@{DW, OSM}, and DFC2020) represent various testing scenarios for pretrained models. SSL4EO-S12@DW involves a transfer learning task that mirrors the noisy label pretraining data SSL4EO-S12@NoL. However, SSL4EO-S12@OSM and DFC2020 differ in their class definitions from SSL4EO-S12@NoL. Moreover, DFC2020 is characterized by dense annotations, whereas the other two datasets contain unlabeled areas within each patch. The latter requires the use of binary masks to calculate the loss omitting unlabeled pixels.

For DFC2020, SSL4EO-S12@DW, and SSL4EO-S12@OSM, we utilize PSPNets [55], DeepLabV3+ [56], and FPN [57] as frameworks to nest ResNet-50 backbones, respectively.<sup>6</sup> The batch size for the three artificial neural networks we set to is 32, 56, and 32. For data augmentation, we employed the random flip with a rate of 0.5 and the random rotation with a rate of 0.2. Besides, random cropping was also included for the SSL4EO-S12@DW and SSL4EO-S12@OSM datasets to resize the training inputs from  $510 \times 510$  as  $496 \times 496$  and  $480 \times 480$ . We compare CromSS with seven state-of-the-art SSL pretrained models in Earth observation including DINO [6], MoCo [7], DeCUR [39], scaleMAE [12], satMAE [10], satMAE++ [11], and DOFA [40].<sup>7</sup> DINO, MoCo, and DeCUR provide pretrained ResNet-50 backbones, while the rest are ViT-large models. We fine-tuned UperNets [58] for ViT backbones with a learning rate of  $10^{-4}$ . DeCUR is a multimodal extension of Barlow Twins that explicitly distinguishes inter- and intramodal embeddings through multimodal redundancy reduction. ScaleMAE implements a scale-aware masking strategy for geospatial representation learning, yet only takes RGB bands as inputs. SatMAE and satMAE++ encode multispectral data (taking ten bands of S2 excluding B1, B9, and B10) as groups of bands with distinct spectral positional encodings. SatMAE++ additionally performs multiscale pretraining. DOFA adaptively integrates various data modalities into a single framework by encoding wavelength information at the input end. In addition, we also compare CromSS with a supervised pretrained ResNet-50 backbone by Satlas [59], which was derived on the SatlasPretrain dataset in a multitask fashion with separate heads—however, we took nine bands of S2 (excluding B1, B8a, B9, and B10) as inputs following [59]. We also took nine bands for DOFA since it was pretrained on the SatlasPretrain dataset. For experiments with S1 data, we compare CromSS pretrained weights with the MoCo, DeCUR, and DOFA weights.

We implemented noisy label pretraining with CromSS for two cases of S2 data: taking either all the 13 bands or the nine high-resolution bands following Satlas as inputs. We report the results of both cases for S2 while only reporting those pretrained along with all S2 bands for S1 in the downstream tasks. To differentiate between the two multimodal architectures, we append `midF` and `lateF` to CromSS to represent the middle and late fusion. As metrics, we picked OA, average accuracy (AA), mean intersection over union (mIoU), and mean F1-scores for each class (mF1). The results listed below stem from three repeated runs of each downstream setup.

##### B. Transfer Learning Results

We first present the transfer learning results on the DFC2020 dataset in Table II. The proposed CromSS demonstrates a notable enhancement in the efficacy of pretrained encoders for RS image segmentation, particularly for S2 multispectral data. On this dataset, middle fusion slightly outperforms late

<sup>6</sup>Models are constructed with the `segmentation-models-pytorch` package available at [https://github.com/qubvel/segmentation\\_models\\_pytorch](https://github.com/qubvel/segmentation_models_pytorch)

<sup>7</sup>The pretrained weights with DINO and MoCo were provided by [18] on the SSL4EO-S12 dataset.

TABLE II

TRANSFER LEARNING RESULTS (%) ON THE DFC2020 DATASET FROM PSPNET WITH RESNET-50 BACKBONES AND UPERNET WITH ViT-LARGE BACKBONES. 3B/9B/10B/13B IN BRACKETS INDICATE THE NUMBER OF S2 BANDS AS INPUTS FOR EACH MODEL. FROZEN AND FINE-TUNED REPRESENT NO UPDATES OF ENCODERS AND OPTIMIZING ENCODERS ALONG WITH DECODERS IN THE TRANSFER LEARNING SETTING. THE BEST RESULTS ARE HIGHLIGHTED IN BOLD. THE ANNOTATIONS ARE THE SAME IN THE FOLLOWING TABLES

Inputs	Encoder		Frozen				Fine-tuned			
	Backbone	Weights (#S2 bands)	OA	AA	mIoU	mF1	OA	AA	mIoU	mF1
S1	ViT-Large	DOFA	49.15±1.89	39.59±1.40	27.51±0.72	36.74±1.13	57.41±1.44	43.74±1.37	31.12±1.43	41.35±1.80
		random	53.89±0.56	40.65±0.38	28.96±0.16	38.74±0.22	52.03±0.56	42.02±0.34	28.02±0.23	37.45±0.29
	ResNet50	MoCo	60.74±0.36	<b>47.48±0.49</b>	34.17±0.27	44.40±0.36	60.30±0.45	45.00±0.65	31.78±0.42	41.78±0.42
		DeCUR	60.22±0.18	46.75±0.58	34.02±0.18	44.24±0.35	57.79±0.48	44.53±0.55	31.23±0.40	41.20±0.39
		CromSS-midF	60.75±0.30	45.66±0.66	34.12±0.28	<b>44.47±0.31</b>	60.07±0.50	<b>46.44±0.63</b>	33.48±0.43	<b>44.05±0.48</b>
		<b>CromSS-lateF</b>	<b>61.65±0.34</b>	44.80±0.09	<b>34.20±0.10</b>	44.14±0.15	<b>60.69±0.19</b>	45.75±0.36	<b>33.63±0.12</b>	43.86±0.10
S2	ViT-Large	scaleMAE (3B)	63.49±0.26	53.31±0.29	38.51±0.20	49.90±0.12	68.42±0.46	55.13±0.25	40.87±0.31	52.18±0.27
		DOFA (9B)	63.49±0.58	48.79±0.20	37.19±0.14	48.87±0.44	64.44±0.50	49.50±0.11	37.65±0.31	48.91±0.59
		satMAE (10B)	67.02±0.41	51.91±0.04	40.13±0.16	51.52±0.31	66.47±0.21	50.51±0.18	39.02±0.16	50.48±0.66
		satMAE++ (10B)	68.41±0.06	53.15±0.77	40.96±0.27	52.93±0.32	67.21±0.17	51.44±0.04	39.90±0.26	51.32±0.66
		random (9B)	61.16±1.32	46.76±0.34	34.00±0.53	44.76±0.61	63.63±0.53	50.28±0.33	36.57±0.31	47.58±0.21
	ResNet50	random (13B)	56.11±0.55	44.63±0.22	31.26±0.14	41.89±0.19	58.83±0.87	45.89±0.36	33.63±0.24	44.14±0.29
		Satlas (9B)	65.91±0.35	51.03±0.41	38.74±0.33	50.27±0.46	64.76±0.35	52.02±0.61	37.85±0.13	49.23±0.21
		DINO (13B)	64.65±0.12	48.53±0.30	37.65±0.13	48.77±0.15	63.75±0.70	49.80±1.06	37.04±0.55	48.95±0.73
		MoCo (13B)	63.19±0.70	50.76±0.34	37.57±0.72	49.83±0.60	60.98±0.64	47.07±0.36	34.75±0.74	46.64±0.40
		DeCUR (13B)	63.02±0.70	47.99±0.57	36.82±0.37	47.90±0.44	64.69±1.03	50.90±0.97	38.02±0.94	50.10±1.10
		CromSS-midF (9B)	66.00±0.15	55.81±0.07	40.31±0.40	52.59±0.17	69.00±0.26	<b>57.07±0.07</b>	42.17±0.13	54.02±0.24
		CromSS-lateF (9B)	68.95±0.67	<b>56.01±0.41</b>	41.69±0.78	53.61±0.94	68.82±0.17	55.75±0.10	41.51±0.28	53.34±0.20
		<b>CromSS-midF (13B)</b>	<b>69.56±0.32</b>	54.61±0.25	<b>42.52±0.11</b>	<b>54.12±0.16</b>	<b>70.41±0.35</b>	55.11±0.30	<b>43.21±0.16</b>	<b>54.91±0.23</b>
		CromSS-lateF (13B)	67.96±0.49	52.62±0.64	41.38±0.34	53.25±0.49	68.49±0.49	53.19±0.59	41.91±0.47	53.59±0.39

TABLE III

TRANSFER LEARNING RESULTS (%) ON THE SSL4EO-S12@DW DATASET FROM DEEPLABV3+ WITH RESNET-50 AND UPERNET WITH ViT-LARGE

Inputs	Encoder		Frozen				Fine-tuned			
	Backbone	Weights (#S2 bands)	OA	AA	mIoU	mF1	OA	AA	mIoU	mF1
S1	ViT-Large	DOFA	62.65±0.20	51.53±0.26	36.92±0.34	49.97±0.45	66.91±0.24	57.75±0.83	41.31±0.19	54.66±0.10
		random	56.84±0.53	44.67±0.55	29.85±0.45	41.10±0.58	65.77±0.28	54.98±0.41	39.80±0.36	53.04±0.50
	ResNet50	MoCo	68.47±0.20	62.24±0.22	45.19±0.47	58.95±0.48	68.69±0.22	60.71±0.69	<b>44.76±0.37</b>	<b>58.41±0.42</b>
		<b>DeCUR</b>	<b>69.38±0.16</b>	<b>62.49±0.38</b>	<b>45.53±0.08</b>	<b>59.36±0.21</b>	<b>69.10±0.18</b>	<b>61.67±1.62</b>	44.72±0.46	58.37±0.45
		CromSS-midF (13B)	66.25±0.16	54.29±0.23	39.63±0.12	52.17±0.26	68.04±0.54	59.75±1.19	43.71±0.42	57.21±0.58
		<b>CromSS-lateF (13B)</b>	<b>66.03±0.14</b>	53.05±0.12	39.63±0.14	52.22±0.18	69.00±0.44	60.47±1.01	44.00±0.82	57.29±0.94
S2	ViT-Large	scaleMAE (3B)	75.89±1.23	72.60±1.67	57.37±1.33	68.95±1.40	78.80±1.50	73.57±2.12	59.89±2.07	72.67±1.78
		DOFA (9B)	82.18±0.27	72.07±0.09	58.45±0.40	71.21±0.39	82.25±0.16	73.24±0.89	60.11±1.13	72.53±0.93
		satMAE (10B)	77.54±0.10	71.34±0.10	58.37±0.11	71.52±0.05	77.65±0.11	71.84±0.36	57.84±0.26	70.58±0.19
		satMAE++ (10B)	76.64±0.18	70.44±0.18	57.04±0.17	69.71±0.05	78.18±0.10	72.83±0.24	58.87±0.23	71.54±0.22
		random (9B)	66.79±0.32	62.68±0.47	48.22±0.32	61.22±0.37	74.64±0.43	70.91±0.75	56.51±0.68	69.56±0.70
	ResNet50	random (13B)	64.80±0.47	60.30±0.20	46.21±0.33	59.11±0.30	73.36±0.21	68.58±0.49	54.72±0.25	67.76±0.26
		Satlas (9B)	75.15±0.10	70.26±0.15	56.39±0.12	69.29±0.24	77.92±0.67	72.59±0.86	59.17±1.02	71.87±1.00
		DINO (13B)	77.64±0.16	72.55±0.53	59.36±0.46	71.98±0.50	77.81±0.23	71.36±0.58	58.05±0.26	70.38±0.39
		MoCo (13B)	77.44±0.62	73.05±0.40	59.52±0.90	72.14±0.83	77.53±0.08	72.76±0.61	58.82±0.56	71.57±0.48
		DeCUR (13B)	76.37±0.18	71.28±0.51	57.45±0.15	70.15±0.28	77.20±0.22	72.05±0.30	58.54±0.39	71.32±0.33
		CromSS-midF (9B)	78.85±0.18	<b>73.77±0.28</b>	<b>61.39±0.09</b>	<b>73.79±0.10</b>	79.19±0.34	73.89±0.56	61.12±0.43	73.64±0.37
		<b>CromSS-lateF (9B)</b>	<b>78.98±0.21</b>	73.39±0.60	60.93±0.60	73.37±0.52	<b>79.93±0.14</b>	<b>74.03±0.45</b>	<b>62.05±0.24</b>	<b>74.18±0.20</b>
		CromSS-midF (13B)	78.60±0.35	73.04±0.23	60.56±0.28	73.01±0.24	78.64±0.19	73.16±0.14	60.79±0.12	73.20±0.13
		CromSS-lateF (13B)	78.48±0.16	72.24±0.18	60.50±0.13	72.86±0.12	79.33±0.24	73.46±0.32	61.10±0.66	73.33±0.61

fusion. CromSS exhibits advantages in most metrics compared to other state-of-the-art pretraining methods, indicating the potential of using noisy labels for semantic segmentation in the pretraining setups. However, the improvement of CromSS for S1 radar data turned out to be less significant. We attribute this discrepancy to the different characteristics of two modalities involved in the pretraining task, that is, land cover classification in our case. S1, which we consider a *weak* modality for land cover semantic segmentation, is more susceptible to label noise, thereby requiring stronger cross-modal strategies to combat it with the assistance of S2—research for future investigation.

We list the transfer learning results acquired from the SSL4EO-S12@DW downstream dataset in Table III. Notably, CromSS utilizing noisy labels for pretraining performs better than the other considered SSL methods for S2 data, showcasing a significant margin. Also, we found that “more”

does not necessarily imply “better.” For example, the models taking fewer bands as inputs tend to perform better, such as CromSS (9B) versus CromSS (13B). This is probably because the low-resolution bands introduce more ambiguity when classifying the nine basic classes defined for this dataset. Regarding S1, the encoders pretrained with noisy labels showcase similar shortcomings, where DeCUR achieves the best results for S1, benefitting from the decoupled learning of two modalities in the pretraining stage. Nevertheless, after fine-tuning, the models pretrained with CromSS can obtain comparable results to DeCUR, demonstrating the effectiveness and potential of CromSS for S1. Furthermore, we present some segmentation maps generated with different pretrained encoders in Fig. 6 for visual comparison. Whether fixed or fine-tuned, encoders pretrained with CromSS can capture more intricate details than those obtained through other pretraining methods listed on the upper lines—particularly when the

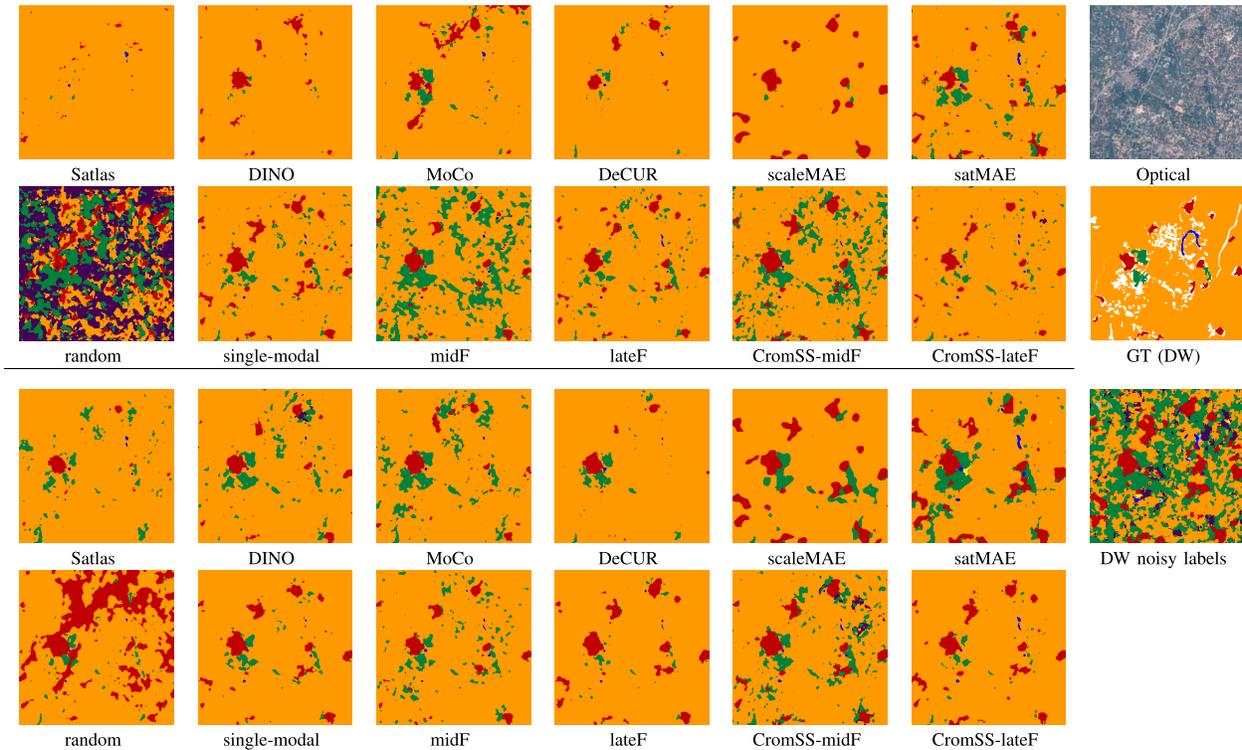


Fig. 6. Segmentation maps obtained on the SSL4EO-S12@DW downstream dataset with S2 as inputs from UperNet with ViT-large pretrained by scaleMAE and satMAE, and the rest from PSPNets with ResNet-50 pretrained by different methods (upper: with frozen encoders, lower: with fine-tuned encoders).

encoder weights are fixed. However, we also observe that segmentation maps generated by the models pretrained on noisy labels tend to exhibit more pepper-and-salt noise. This is likely a side effect of the model’s ability to capture finer details, which may also amplify minor inconsistencies or artifacts. While this noise can be seen as a drawback, it underscores the model’s sensitivity to subtle features, which could be advantageous in applications where detail preservation is critical.

Finally, we present the results generated on the SSL4EO-S12@OSM downstream dataset in Table IV. Consistent with the findings on the other two datasets, CromSS demonstrates superior performance compared to randomly initialized and SSL-pretrained models for S2 data. This dataset is challenging since it contains more fine-grained classes with fewer training samples than the other two downstream tasks. In this case, all the pretraining methods exhibit a significant advantage over random initialization, particularly our proposed CromSS. Besides, the models taking more bands as inputs perform better on this dataset since more complicated scenes require more information to parse. Moreover, when encoders are fine-tuned alongside decoders, the proposed methods exhibit less severe overfitting when dealing with very limited training samples. These observations align with the insights drawn from segmentation maps illustrated in Fig. 7, where we also find that patch-based ViT models easily lead to oversmoothed segmentation maps, especially using a larger patch size such as scaleMAE with a patch size of 16 compared to satMAE with a patch size of 8. In contrast, the effectiveness of noisy label pretraining for S1 data requires further improvements compared to that for S2.

### C. Ablation Experiments

In this section, we explore the roles of the components of CromSS, including multimodal pretraining architectures, cross-modal sample selection, and spatial-temporal label smoothing, through a series of ablation experiments.

We first test the effectiveness of the first two components. Table V lists the transfer learning results of single-modal noisy label pretraining (S1/S2) and multimodal pretraining without cross-modal sample selection (midF and lateF denote middle and late fusion), in comparison to random initialization and CromSS (multimodal pretraining with cross-modal sample selection), along with the corresponding segmentation maps illustrated in Figs. 6 and 7. Notably, single-modal pretraining with noisy labels brings a significant performance gain compared to random initialization for both S1 and S2 data on three datasets. This margin is further amplified after introducing multimodal learning into the noisy label pretraining framework (midF and lateF). These results suggest that multimodal training can inherently aid models in alleviating the adverse impacts of label noise. However, the margin is smaller for S1 than that for S2. As mentioned above, S1 is a “weak” modality in the pretraining setup, where the learning is more easily affected by label noise, thus benefiting less from noisy labels. This further impacts the effectiveness of the sample selection module for S1 data. On the contrary, we observe some further improvements for S2 data after integrating cross-modal sample selection into the multimodal pretraining framework. However, CromSS fails to surpass midF and lateF on the SSL4EO-S12@DW dataset. We attribute this shortcoming to the following: although the

TABLE IV  
TRANSFER LEARNING RESULTS (%) ON THE SSL4EO-S12@OSM DATASET FROM FPN WITH RESNET-50 AND UPERNET WITH ViT-LARGE

Inputs	Encoder		Frozen				Fine-tuned			
	Backbone	Weights (#S2 bands)	OA	AA	mIoU	mF1	OA	AA	mIoU	mF1
S1	ViT-Large	DOFA	72.18±0.45	49.22±0.31	28.94±0.44	38.48±0.06	75.30±0.68	52.63±0.40	33.71±0.43	44.60±0.32
		random	63.94±1.01	<b>57.75±5.09</b>	18.21±1.14	24.76±1.66	66.44±3.76	48.70±3.79	24.01±2.01	32.58±2.35
	ResNet50	MoCo	78.37±0.54	51.16±2.00	36.37±0.69	46.24±0.71	71.94±3.21	49.08±3.31	29.55±1.68	38.99±1.67
		<b>DeCUR</b>	<b>80.10±0.28</b>	51.13±0.67	<b>38.16±0.57</b>	<b>48.25±0.59</b>	<b>79.83±0.53</b>	<b>58.91±0.69</b>	<b>36.55±0.29</b>	<b>46.92±0.23</b>
		CromSS-midF	76.69±0.58	48.11±0.78	33.61±0.57	43.83±0.66	77.72±1.18	51.46±2.12	36.15±0.97	46.65±0.85
		CromSS-lateF	76.57±0.54	47.89±1.18	32.92±0.80	42.80±0.89	77.89±0.78	48.53±1.36	35.13±0.73	45.04±0.84
S2	ViT-Large	scaleMAE (3B)	82.52±0.12	66.32±1.45	44.96±0.55	55.76±0.81	80.80±0.80	62.20±1.72	43.18±0.88	54.38±1.01
		DOFA (9B)	82.18±0.27	60.82±0.36	43.41±0.09	53.86±0.23	82.25±0.16	61.08±0.19	45.72±0.92	57.09±0.98
		satMAE (10B)	84.59±0.35	<b>67.32±1.98</b>	48.07±1.55	59.32±2.19	84.26±0.06	69.03±0.27	46.20±0.47	56.74±0.55
		satMAE++ (10B)	83.65±0.13	65.48±0.36	45.60±0.07	56.71±0.10	84.24±0.27	69.51±0.79	46.19±1.04	56.92±1.25
	ResNet50	random (9B)	70.89±0.40	60.04±3.57	28.35±0.53	36.71±0.57	67.38±7.02	50.80±3.17	27.98±3.90	36.39±4.29
		random (13B)	71.26±0.71	60.68±3.61	28.42±0.55	36.72±0.63	68.77±4.29	54.23±4.69	27.64±2.82	36.03±3.37
		Satlas (9B)	81.89±0.36	60.25±0.49	43.21±0.30	54.34±0.31	83.16±1.01	62.70±1.39	44.51±0.47	55.60±0.43
		DINO (13B)	83.46±0.30	60.29±1.77	42.77±0.73	53.68±0.93	78.53±2.30	57.87±2.85	37.54±2.58	47.60±2.63
		MoCo (13B)	83.87±0.43	61.82±1.28	46.79±0.75	56.93±0.80	80.67±1.25	62.75±2.84	41.99±1.93	52.74±2.07
		DeCUR (13B)	84.58±0.27	60.07±1.17	44.96±0.71	56.11±0.81	83.64±0.39	61.47±0.15	45.21±0.32	56.36±0.49
		CromSS-midF (9B)	83.98±0.14	62.27±2.20	46.58±0.49	57.31±0.56	84.57±0.63	64.01±3.74	47.69±0.97	58.76±1.04
		CromSS-lateF (9B)	83.64±0.30	60.55±0.25	45.72±0.36	56.61±0.42	85.25±0.12	64.45±0.27	49.38±0.16	60.62±0.36
		<b>CromSS-midF (13B)</b>	<b>84.84±0.35</b>	65.99±1.06	<b>49.06±0.60</b>	<b>60.28±0.81</b>	<b>86.04±0.30</b>	<b>70.05±2.71</b>	<b>50.11±0.64</b>	<b>62.59±0.64</b>
		CromSS-lateF (13B)	84.75±0.43	66.63±2.48	47.38±0.79	58.49±0.85	85.36±0.37	65.84±1.87	49.38±0.66	60.67±0.56

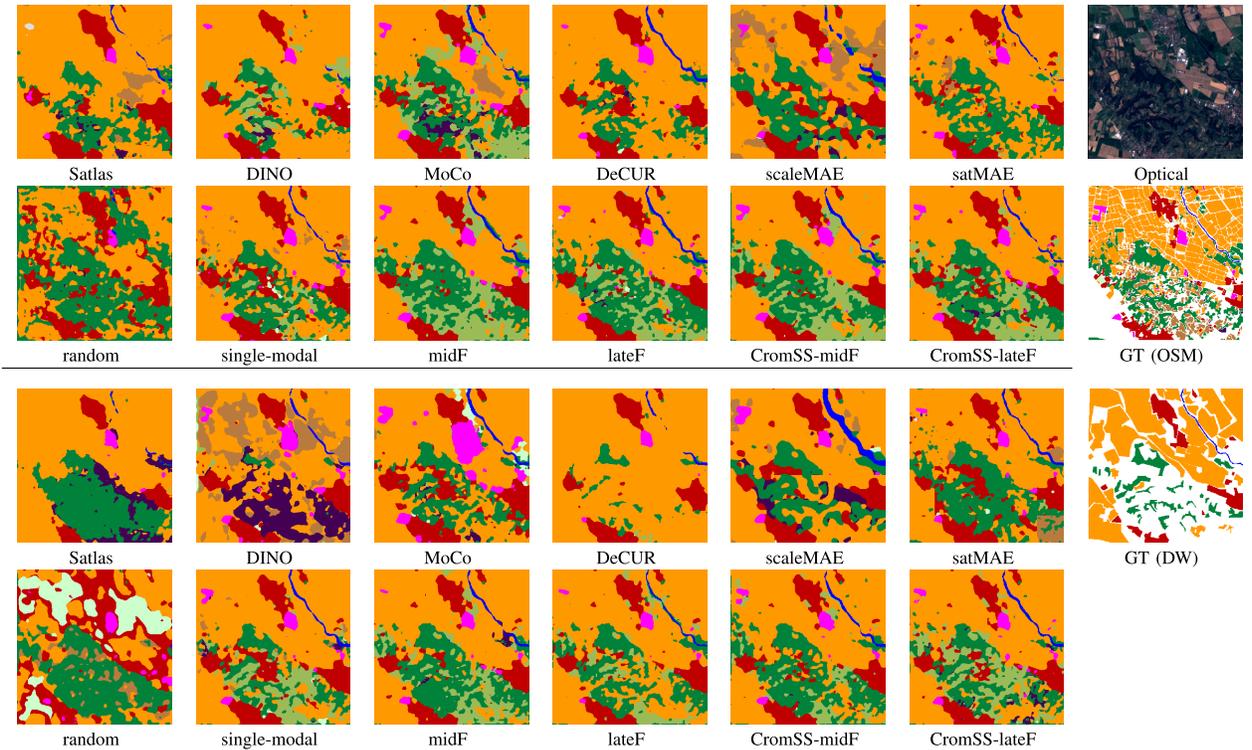


Fig. 7. Segmentation maps obtained on the SSL4EO-S12@OSM downstream dataset with S2 as inputs from UperNet with ViT-large pretrained by scaleMAE and satMAE, and the rest from FPN with ResNet-50 pretrained by different methods (upper: with frozen encoders, lower: with fine-tuned encoders).

sample selection masks are designed to suppress label noise, they may block relevant information, such as boundaries between areas of two classes or fine details where many different classes coexist in a small neighborhood of pixels. If sample selection masks block this information, the model performance degrades when training is continued with a limited number of exact labels under the same definitions (we will analyze this issue in detail in Section IV-D). Nevertheless, the performance of CromSS is still comparable to midF and lateF. Remarkably, CromSS exhibits a significant advantage on the other two datasets. Interestingly, the partial loss of specific details during the pretraining stage contributes to

enhanced generality when transferred to other datasets. The advantage is more significant in the middle fusion cases, implying that explicit constraint via decoder weight sharing can potentially boost sample selection. Moreover, when encoders are fine-tuned alongside decoders, CromSS shows less severe overfitting when dealing with limited training samples.

We then evaluate the effectiveness of the spatial-temporal label smoothing technique within the CromSS-midF framework. As shown in (13), in addition to the commonly used class uniform term (weighed by  $\beta$ ), we introduce a spatial-temporal prior term (weighed by  $\mu$ ), derived from four-

TABLE V

TRANSFER LEARNING RESULTS (%) ON THE THREE DOWNSTREAM DATASETS OBTAINED BY RESNET-50 WITH DIFFERENT NOISY LABEL PRETRAINING STRATEGIES: SINGLE-MODAL, MULTIMODAL (MIDF AND LATEF), AND MULTIMODAL WITH SAMPLE SELECTION (CROMSS)

Inputs	Weights	DFC2020						SSL4EO-S12@DW						SSL4EO-S12@OSM					
		Frozen			Fine-tuned			Frozen			Fine-tuned			Frozen			Fine-tuned		
		OA	AA	mIoU	OA	AA	mIoU	OA	AA	mIoU	OA	AA	mIoU	OA	AA	mIoU	OA	AA	mIoU
S1	random	53.89	40.65	28.96	52.03	42.02	28.02	56.84	44.67	29.85	65.77	54.98	39.80	63.94	57.75	18.21	66.44	48.70	24.01
	single-modal	61.15	46.14	34.17	59.91	45.68	33.46	65.80	56.15	39.98	68.09	59.88	43.26	77.12	49.04	33.96	78.30	49.85	36.2
	midF	60.00	44.16	33.29	59.99	45.69	33.13	67.04	54.73	40.30	67.89	57.68	42.39	77.21	49.53	34.10	78.44	51.94	36.88
	lateF	61.52	45.98	34.12	61.06	46.83	33.66	67.04	56.51	40.16	68.45	61.10	43.49	77.33	48.70	34.54	78.18	49.64	36.01
	CromSS-midF	60.75	45.66	34.12	60.07	46.44	33.48	66.25	54.29	39.63	68.04	59.75	43.71	76.69	48.11	33.61	77.72	51.46	36.15
	CromSS-lateF	61.65	44.80	34.20	60.69	45.75	33.63	66.03	53.05	39.63	69.00	60.47	44.00	76.57	47.89	32.92	77.89	48.53	35.13
S2	random	56.11	44.63	31.26	58.83	45.89	33.63	64.80	60.30	46.21	73.36	68.58	54.72	71.26	60.68	28.42	68.77	54.23	27.64
	single-modal	66.36	52.52	40.27	67.47	53.22	41.15	78.46	72.54	59.30	78.36	72.37	59.56	83.94	64.87	45.29	83.72	64.72	46.68
	midF	68.57	52.49	41.31	68.70	52.94	41.46	78.82	73.28	60.91	78.60	73.12	60.32	84.05	65.26	47.37	83.50	63.98	47.96
	lateF	68.88	53.58	41.30	68.42	53.24	41.40	78.69	72.93	60.29	78.71	73.80	61.39	84.73	64.42	47.62	84.43	66.50	48.89
	CromSS-midF	69.56	54.61	42.52	70.41	55.11	43.21	78.60	73.04	60.56	78.64	73.16	60.79	84.84	65.99	49.06	86.04	70.05	51.11
	CromSS-lateF	68.93	52.62	41.35	68.37	53.19	41.83	78.48	72.24	60.50	79.33	73.46	61.10	84.75	66.63	47.38	85.36	65.84	49.38

TABLE VI

TRANSFER LEARNING RESULTS (mIoU, %) OBTAINED USING RESNET-50 PRETRAINED BY CROMSS-MIDF UNDER DIFFERENT LABEL SMOOTHING SETTINGS, WHERE WE REPORT THE SCORES WITH “FROZEN/FINE-TUNED” ENCODERS ON THE THREE DOWNSTREAM DATASETS

Spatial-temporal ( $\mu$ )	Uniform ( $\beta$ )	DFC2020	SSL4EO-S12 @DW	SSL4EO-S12 @OSM
0	0	43.26/42.63	61.07/61.26	47.24/49.80
0	0.2	41.64/42.32	58.65/59.36	47.67/50.57
0.2	0	42.3/43.52	60.98/61.61	46.98/49.49
0.15	0.05	42.52/43.21	60.64/60.66	49.00/51.28

seasonal label masks at each location for label smoothing. Table VI presents the transfer learning results obtained with various combinations of these two terms. The results indicate that the combined approach ( $\mu = 0.15$ ,  $\beta = 0.05$ ) used in our method yields more stable performance across all datasets. While it may not outperform other settings on each dataset, it provides comparable results without significant degradation. In contrast, the plain uniform label smoothing strategy leads to a notable performance drop on nearly all datasets, and the pure spatial-temporal label smoothing strategy does not perform satisfactorily on the SSL4EO-S12@OSM dataset. This highlights the robustness and adaptability of our proposed spatial-temporal label smoothing technique across different data sources.

#### D. Model Performance in the Pretraining Stage

To further demonstrate the effectiveness of CromSS, we evaluated the accuracy of the pretrained U-Net models with noisy labels using the SSL4EO-S12@DW test set. As depicted in Fig. 8(a) and (b), multimodal training, either through middle or late fusion, consistently improves model performance for both S1 and S2 data compared to single-modal training. Remarkably, incorporating cross-modal sample selection enhances the ability of middle fusion to mitigate label noise for both S1 and S2 data. In contrast, the performance gain from late fusion is most prominent for S2 data. This observation again underlines the potential of sharing decoders as a robust constraint for training models with noisy labels.

To analyze the disparity between encoders and decoders, we quantify their similarity in feature statistics (the mean and variance values stored in the batch normalization layers)

across models pretrained in single-modal and multimodal setups for S2 data, as illustrated in Fig. 8(c) and (d). An intriguing observation we notice is that encoders trained in distinct setups tend to converge to similar representations as training progresses, while, at the same time, decoders exhibit more noticeable discrepancies. This observation becomes more prominent in the later stages of model training. A potential mechanism behind may read as follows: although encoders prioritize extracting general features, decoders focus on more specific ones. The discrepancy is more prominent in middle fusion settings than in late fusion. The KL divergence values of single-modal models and the late fusion counterparts approach zero, indicating similar behavior and possibly less need for adjustments.

Furthermore, we present two sets of segmentation maps generated from S2 optical images by different U-Net models trained with noisy labels in Fig. 9. In the first set of maps (top row), employing sample selection leads to a more accurate delineation of water bodies. Conversely, in the second set (bottom row), maps generated by CromSS appear somewhat oversmoothed, with spatial details missing. This observation aligns with the degraded performance of encoders pretrained with CromSS on the SSL4EO-S12@DW dataset compared to other multimodal pretrained schemes, as discussed in Section IV-C. The sample selection bias is one potential root cause of our observation of oversmoothed semantic segmentation maps that discard spatial details. As illustrated in Fig. 11, the small objects (see dark spots in the selection masks  $W_l$ ) or class boundaries may get treated as uncertain and thus will get assigned smaller weight values during the selection process. The absence of spatial details may lead to degraded accuracy for a downstream task (e.g., exact DW dataset labels) identical to the noisy labels in the pretraining phase. However, masking geospatially small objects can partly avoid model overfitting to the pretraining dataset for other downstream tasks (e.g., exact OSM dataset labels). Consequently, our sample selection strategy can enhance the effectiveness of noisy label pretrained encoders when transferred to other datasets. Besides, as observed from the last row of Fig. 11, informing the confidence mask of S2 ( $W_l$  from  $F_l$ ) with that of S1 can reduce uncertainty in the ice region within the final selection mask of S2 ( $W_l$  from  $F_l'$ ). This example indicates the effectiveness of the cross-modal confidence integration in calibrating the single-modal selection masks.

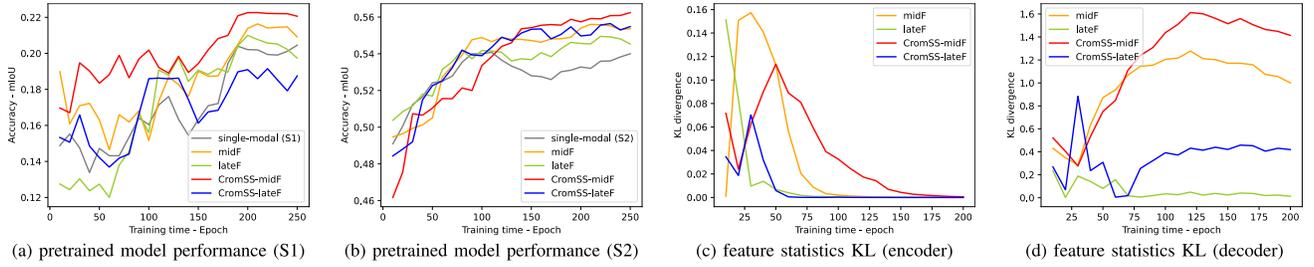


Fig. 8. Accuracy (mIoU) of the pretrained models versus training time for (a) S1 and (b) S2 data, and the similarity (KL divergence) of the feature means and standard deviations in the batch normalization layer between the single-modal and multimodal pretraining for S2 data versus training time for (c) encoders and (d) decoders. The curves were obtained with an interval of 10 epochs, and smoothed by average moving with a window size of 5.

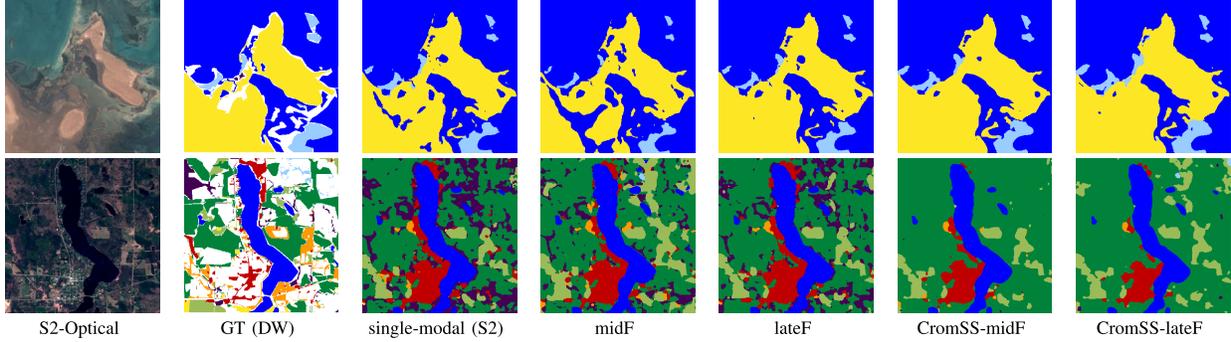


Fig. 9. Segmentation maps generated by the pretrained models (encoder-decoder for modality 2 in Fig. 5) with noisy labels taking S2 as inputs. The color coding is the same as that in Fig. 1 for the DW dataset.

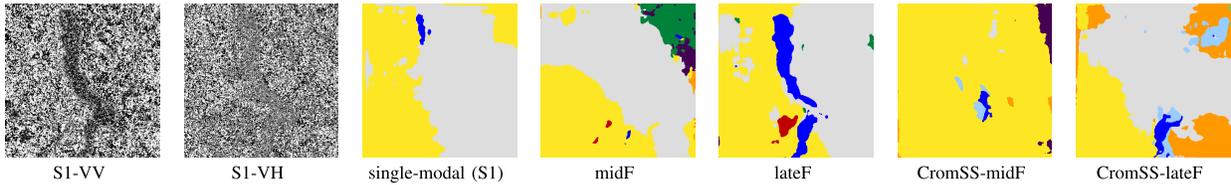


Fig. 10. Segmentation maps inferred by the pretrained model (encoder-decoder for modality 1 in Fig. 5) with noisy labels taking S1 as inputs. The color coding is the same as that in Fig. 1 for the DW dataset, and the scene illuminated by the radar sensor is the same as in Fig. 9 (bottom).

Finally, we display segmentation maps generated from S1 data by U-Net models in Fig. 10. On the contrary to Fig. 9, these segmentation maps are much less accurate. As discussed, S1 is a “weak” modality for land cover land use classification. We can hardly distinguish different types of landscapes from the VV and VH images by eye, given pronounced noise patterns. Correspondingly, the training on S1 data is less effective than training on S2 data, leading to suboptimal transfer learning results in the downstream tasks.

### E. Comparison in Pretrained Encoder Characteristics

To directly compare encoders pretrained with different methods, we present the KL divergence values of encoder (ResNet-50) and decoder weights before and after fine-tuning across various downstream tasks. We choose two classical image-wise SSL methods, DINO and MoCo, as representatives to compare with the noisy label pretraining ones to support the discussion. We use the discrete form of KL divergence and obtain the discrete distributions by binning data with  $n = 100$  even bins. As shown in Fig. 12, the differences in the KL divergence of various pretraining methods tend to be smaller for decoders when compared to the corresponding encoders. Specifically, encoders pretrained with noisy labels

exhibit less optimization change than those pretrained by DINO and MoCo, where MoCo exhibits the largest variation. This suggests that encoders pretrained with noisy labels focusing on pixel-wise information are closer to the optimal solution than those mainly utilizing image-level information. Furthermore, we observe that variation in the weights of decoders in the fine-tuning stage is more substantial when using the encoders pretrained on noisy labels. A performant encoder enables decoders to generate better solutions, albeit such an encoder potentially renders the decoder prone to overfitting. These adverse effects need to be taken care of in stages of fine-tuning.

Finally, we present a comparison in feature extraction patterns of the ResNet-50 encoders trained by image-wise SSL methods, DINO and MoCo, and our proposed pixel-wise pretraining with noisy labels. We extract the features from each ResNet-50 block/module and then apply principal component analysis (PCA) to get the accumulated variance ratios<sup>8</sup> for each feature set. We plot these accumulated variance ratios as curves per module in Fig. 13, where the upper and lower rows correspond to the results obtained

<sup>8</sup>Estimated using the function of `sklearn.decomposition.PCA` (<https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.PCA.html>).

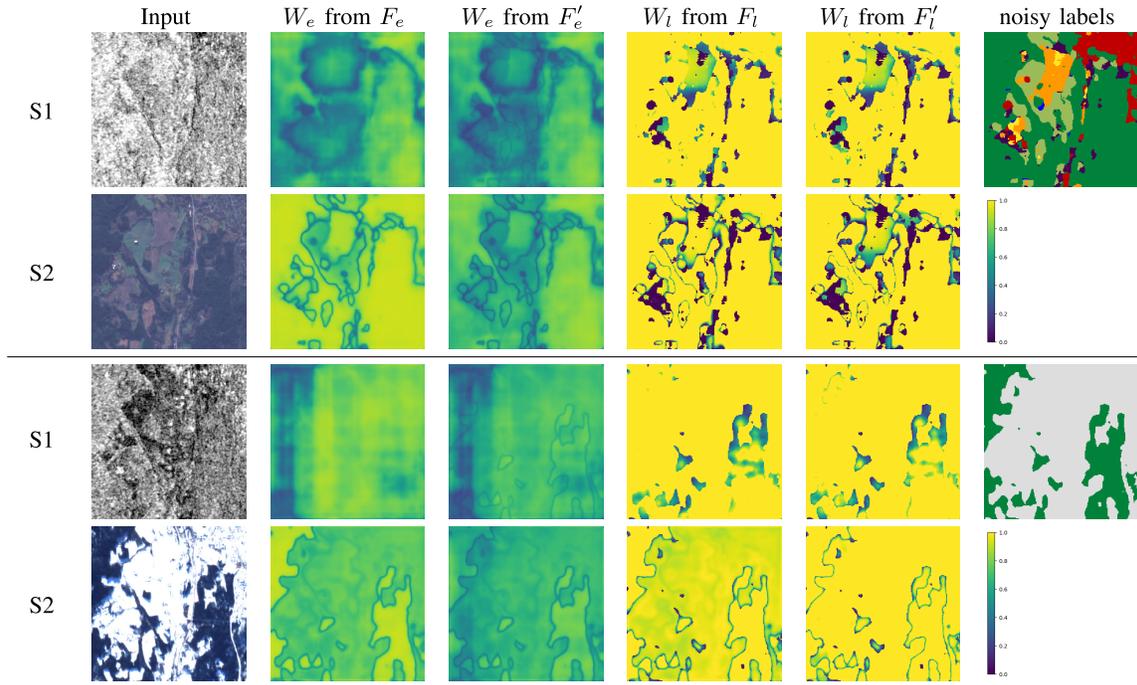


Fig. 11. Entity-based and label-based sample selection masks ( $W_e$  and  $W_l$ ) generated by CromSS-midF (13B) from the single-modal and cross-modal confidence masks ( $F$  and  $F'$ ) for a given scene at two different seasons. The colorbar for the sample selection maps is displayed on the right-hand side, while the legend for the noisy labels can be found in Fig. 2.

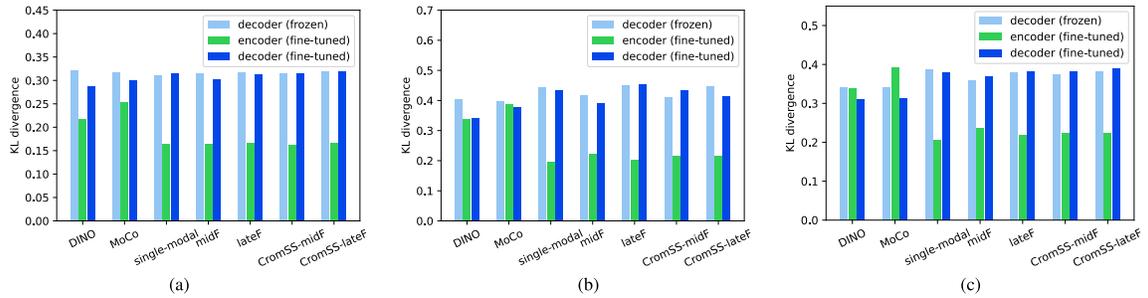


Fig. 12. KL divergence of the encoder and decoder weights before and after fine-tuning, where the brackets indicate the type of encoder weights. We display  $-1/\log(\text{KL})$  instead of the raw KL divergence values for improved visualization. (a) DFC2020. (b) DW. (c) OSM.

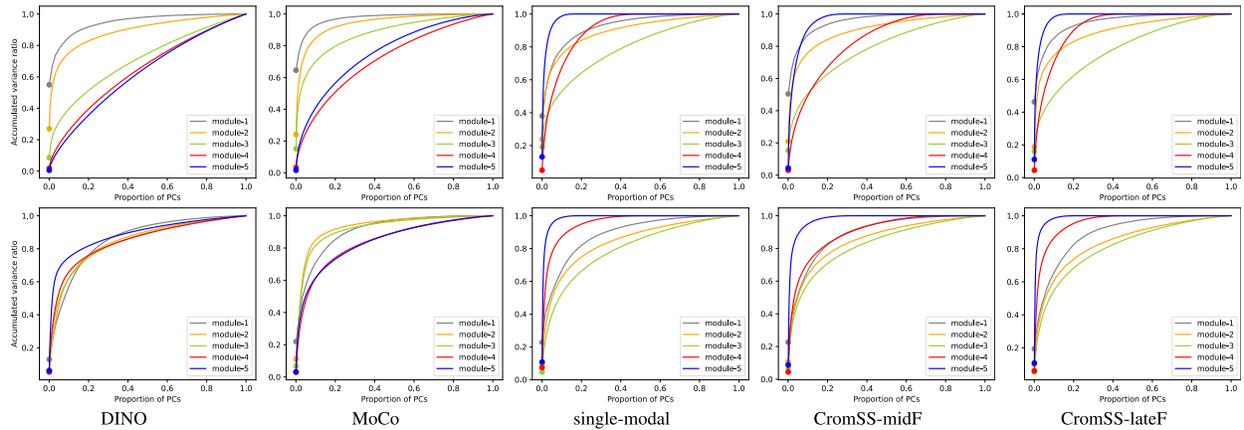


Fig. 13. Accumulated variance ratio plotted against the number of principal components (PCs, in proportion) of output features from various modules/blocks in the ResNet-50 encoders using S2 as inputs. Modules 1–5 represent the modules from the input to the bottleneck. (Upper) Image-level features generated from the EuroSAT dataset [60]. (Lower) Pixel-level features generated from the SSL4EO-S12@DW test set.

at the image and pixel levels for EuroSAT and SSL4EO-S12@DW, respectively. Empirically, a closer alignment of the accumulated variance ratio curves  $y = y(x)$  with the

diagonal  $y = x$  indicates a more equal importance of all features. A steep slope  $dy/dx$  for  $x \ll 1$  underlines the dominance of a small subset of the features, suggesting higher

feature redundancy. It is important to note that these variance ratios are relative within each feature set, and absolute values are not directly comparable across different sets. Therefore, accumulated variance ratios primarily aim to illustrate the differences in feature distributions across modules within each encoder rather than for direct comparison between models. As depicted in Fig. 13, the encoders pretrained by image-level SSL methods, DINO and MoCo, typically produce the image-level feature sets with the variance ratio curves of the bottleneck module (module 5) more closely aligning with the diagonal. This tendency gets weak on the pixel-level feature sets. In contrast, the encoders pretrained with pixel-wise noisy labels exhibit the least feature redundancy in the middle layers (module 3), with redundancy increasing in both the shallower layers (modules 1 and 2) and the deeper layers (modules 4 and 5), for either image-level or pixel-level features. This increase is particularly pronounced in the deeper layers (modules 4 and 5). These findings highlight the orthogonal nature of encoders pretrained with image-wise (discriminative-like DINO and MoCo) versus pixel-wise (generative-like noisy label pretraining) methods. While pixel-wise pretraining methods emphasize learning detailed, high-resolution features in the shallower layers, they appear to be less sensitive to global semantic information. Nevertheless, the image-wise pretraining is the opposite. This suggests that the choice of pretraining methods significantly influences the model's focus on specific aspects of the images.

## V. CONCLUSION AND PERSPECTIVES

This work introduces a pretraining strategy called CromSS, guided by noisy labels for large-scale RS image segmentation. CromSS exploits a cross-modal sample selection strategy based on the pixel-wise estimated class distributions over classes modeled by multiple sensors/modalities of a given geospatial scene. Two kinds of confidence masks are generated for sample selection purposes: label-based confidence for noisy label sampling with segmentation losses and entity-level confidence for cross-modal consistency modeling with consistency losses. On top of this, a confidence correction strategy is introduced to enhance the confidence masks from one modality through those from another. Furthermore, we test middle and late fusion for pretraining architecture design with noisy labels. To evaluate the effectiveness of the proposed CromSS, we collected a NoLDO-S12 dataset with S1 and S2 as two modalities, alongside the noisy labels sourced from DW for pretraining and GT labels retrieved from DW and OSM serving as downstream tasks. Transfer learning results on three differently defined datasets demonstrate the effectiveness of the CromSS-pretrained ResNet-50 encoders. Specifically, we find that multimodal training has a strong ability to combat label noise. Sample selection risks losing detailed information when the pretrained models are transferred to the downstream tasks identical to the pretraining one. At the same time, it can boost the generability of pretrained models to some extent when applied to a different setting.

Our work represents an initial exploration of using noisy labels for pretraining in segmentation tasks. Although containing noise, these labels are easier to acquire in greater quantities on a larger scale than exact labels. However, the relationship between label quantity and quality remains an open question. In future work, we will investigate this aspect to

gain deeper insights into optimizing this trade-off and refining the usage of noisy labels for feature learning. Additionally, we will explore the potential of CromSS for ViT pretraining, comparing it to MIM for RS image segmentation. Another promising research direction is to study the discrepancies in CromSS performance across modalities, particularly focusing on improving its performance on weaker modalities like S1, as observed in our experiments.

## REFERENCES

- [1] T. Zhao et al., "Artificial intelligence for geoscience: Progress, challenges and perspectives," *Innovation*, vol. 5, no. 5, Aug. 2024, Art. no. 100691. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666675824001292>
- [2] Y. Li, T. Shi, Y. Zhang, W. Chen, Z. Wang, and H. Li, "Learning deep semantic segmentation network under multiple weakly-supervised constraints for cross-domain remote sensing image semantic segmentation," *ISPRS J. Photogramm. Remote Sens.*, vol. 175, pp. 20–33, May 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271621000423>
- [3] Y. Wang, C. M. Albrecht, N. A. A. Braham, L. Mou, and X. X. Zhu, "Self-supervised learning in remote sensing: A review," *IEEE Geosci. Remote Sens. Mag.*, vol. 10, no. 4, pp. 213–247, Dec. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9875399>
- [4] X. X. Zhu et al., "On the foundations of Earth and climate foundation models," 2024, *arXiv:2405.04285*.
- [5] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16000–16009. [Online]. Available: [https://openaccess.thecvf.com/content/CVPR2022/html/He\\_Masked\\_Autoencoders\\_Are\\_Scalable\\_Vision\\_Learners\\_CVPR\\_2022\\_paper](https://openaccess.thecvf.com/content/CVPR2022/html/He_Masked_Autoencoders_Are_Scalable_Vision_Learners_CVPR_2022_paper)
- [6] M. Caron et al., "Emerging properties in self-supervised vision transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 9630–9640. [Online]. Available: [https://openaccess.thecvf.com/content/ICCV2021/html/Caron\\_Emerging\\_Properties\\_in\\_Self-Supervised\\_Vision\\_Transformers\\_ICCV\\_2021\\_paper.html](https://openaccess.thecvf.com/content/ICCV2021/html/Caron_Emerging_Properties_in_Self-Supervised_Vision_Transformers_ICCV_2021_paper.html)
- [7] X. Chen, H. Fan, R. Girshick, and K. He, "Improved baselines with momentum contrastive learning," 2020, *arXiv:2003.04297*.
- [8] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," 2020, *arXiv:2010.11929*.
- [9] X. Sun et al., "RingMo: A remote sensing foundation model with masked image modeling," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5612822. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9844015>
- [10] Y. Cong et al., "SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery," in *Proc. Adv. Neural Inf. Process. Syst.*, Jan. 2022, pp. 197–211. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2022/hash/01c561df365429f33cd7a7faa44c985-Abstract-Conference.html](https://proceedings.neurips.cc/paper_files/paper/2022/hash/01c561df365429f33cd7a7faa44c985-Abstract-Conference.html)
- [11] M. Noman, M. Naseer, H. Cholakkal, R. M. Anwar, S. Khan, and F. S. Khan, "Rethinking transformers pre-training for multi-spectral satellite imagery," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 27811–27819.
- [12] C. J. Reed et al., "Scale-MAE: A scale-aware masked autoencoder for multiscale geospatial representation learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 4088–4099.
- [13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778. [Online]. Available: [https://openaccess.thecvf.com/content\\_cvpr\\_2016/html/He\\_Deep\\_Residual\\_Learning\\_CVPR\\_2016\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html)
- [14] K. Ayush et al., "Geography-aware self-supervised learning," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 10181–10190. [Online]. Available: [https://openaccess.thecvf.com/content/ICCV2021/html/Ayush\\_Geography-Aware\\_Self-Supervised\\_Learning\\_ICCV\\_2021\\_paper.html](https://openaccess.thecvf.com/content/ICCV2021/html/Ayush_Geography-Aware_Self-Supervised_Learning_ICCV_2021_paper.html)
- [15] O. Mañas, A. Lacoste, X. Giro-i-Nieto, D. Vazquez, and P. Rodriguez, "Seasonal contrast: Unsupervised pre-training from uncurated remote sensing data," 2021, *arXiv:2103.16607*.
- [16] X. Wanyan, S. Seneviratne, S. Shen, and M. Kirley, "Extending global-local view alignment for self-supervised learning with remote sensing imagery," 2023, *arXiv:2303.06670*.

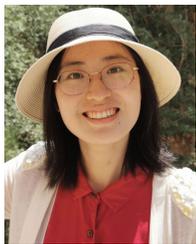
- [17] U. Mall, B. Hariharan, and K. Bala, "Change-aware sampling and contrastive learning for satellite images," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 5261–5270. [Online]. Available: [https://openaccess.thecvf.com/content/CVPR2023/html/Mall\\_Change-Aware\\_Sampling\\_and\\_Contrastive\\_Learning\\_for\\_Satellite\\_Images\\_CVPR\\_2023\\_paper.html](https://openaccess.thecvf.com/content/CVPR2023/html/Mall_Change-Aware_Sampling_and_Contrastive_Learning_for_Satellite_Images_CVPR_2023_paper.html)
- [18] Y. Wang, N. A. A. Braham, Z. Xiong, C. Liu, C. M. Albrecht, and X. X. Zhu, "SSL4EO-s12: A large-scale multimodal, multi-temporal dataset for self-supervised learning in Earth observation [software and data sets]," *IEEE Geosci. Remote Sens. Mag.*, vol. 11, no. 3, pp. 98–106, Sep. 2023. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10261879>
- [19] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Miami, FL, USA, Jun. 2009, pp. 248–255. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/5206848>
- [20] A. Kirillov et al., "Segment anything," 2023, *arXiv:2304.02643*.
- [21] N. Ravi et al., "SAM 2: Segment anything in images and videos," in *Proc. 13th Int. Conf. Learn. Represent.*, 2025. [Online]. Available: <https://openreview.net/forum?id=Ha6RTeWMd0>
- [22] H. Wang, K. Köser, and P. Ren, "Large foundation model empowered discriminative underwater image enhancement," *IEEE Trans. Geosci. Remote Sens.*, vol. 63, 2025, Art. no. 5609317. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10824846>
- [23] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning (still) requires rethinking generalization," *Commun. ACM*, vol. 64, no. 3, pp. 107–115, 2021, doi: [10.1145/3446776](https://doi.org/10.1145/3446776).
- [24] C. Liu, C. M. Albrecht, Y. Wang, Q. Li, and X. X. Zhu, "AIO2: Online correction of object labels for deep learning with incomplete annotation in remote sensing image segmentation," *IEEE Trans. Geosci. Remote Sens.*, vol. 62, 2024, Art. no. 5613917.
- [25] D. Mahajan et al., "Exploring the limits of weakly supervised pretraining," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Sep. 2018, pp. 181–196. [Online]. Available: [https://openaccess.thecvf.com/content\\_ECCV\\_2018/html/Dhruv\\_Mahajan\\_Exploring\\_the\\_Limits\\_ECCV\\_2018\\_paper.html](https://openaccess.thecvf.com/content_ECCV_2018/html/Dhruv_Mahajan_Exploring_the_Limits_ECCV_2018_paper.html)
- [26] D. Ghadiyaram, D. Tran, and D. Mahajan, "Large-scale weakly-supervised pre-training for video action recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 12038–12047. [Online]. Available: [https://openaccess.thecvf.com/content\\_CVPR\\_2019/html/Ghadiyaram\\_Large-Scale\\_Weakly-Supervised\\_Pre-Training\\_for\\_Video\\_Action\\_Recognition\\_CVPR\\_2019\\_paper.html](https://openaccess.thecvf.com/content_CVPR_2019/html/Ghadiyaram_Large-Scale_Weakly-Supervised_Pre-Training_for_Video_Action_Recognition_CVPR_2019_paper.html)
- [27] P. Kaiser, J. D. Wegner, A. Lucchi, M. Jaggi, T. Hofmann, and K. Schindler, "Learning aerial image segmentation from online maps," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 11, pp. 6054–6068, Nov. 2017.
- [28] E. Maggiori, Y. Tarabalka, G. Charpiat, and P. Alliez, "Convolutional neural networks for large-scale remote-sensing image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 2, pp. 645–657, Feb. 2017.
- [29] C. Liu, C. M. Albrecht, Y. Wang, and X. X. Zhu, "Task specific pretraining with noisy labels for remote sensing image segmentation," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Jul. 2024, pp. 7040–7044.
- [30] H. Li, H. Wang, Y. Zhang, L. Li, and P. Ren, "Underwater image captioning: Challenges, models, and datasets," *ISPRS J. Photogramm. Remote Sens.*, vol. 220, pp. 440–453, Feb. 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271624004726>
- [31] Z. Yuan, L. Mou, Q. Wang, and X. X. Zhu, "From easy to hard: Learning language-guided curriculum for visual question answering on remote sensing data," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5623111. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9771224>
- [32] M. Schmitt and X. X. Zhu, "Data fusion and remote sensing: An ever-growing relationship," *IEEE Geosci. Remote Sens. Mag.*, vol. 4, no. 4, pp. 6–23, Dec. 2016.
- [33] D. Hong, J. Hu, J. Yao, J. Chanussot, and X. X. Zhu, "Multimodal remote sensing benchmark datasets for land cover classification with a shared and specific feature learning model," *ISPRS J. Photogramm. Remote Sens.*, vol. 178, pp. 68–80, Aug. 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271621001362>
- [34] V. S. F. Garnot, L. Landrieu, and N. Chehata, "Multi-modal temporal attention models for crop mapping from satellite time series," *ISPRS J. Photogramm. Remote Sens.*, vol. 187, pp. 294–305, May 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271622000855>
- [35] Z. Cai et al., "Improving agricultural field parcel delineation with a dual branch spatiotemporal fusion network by integrating multimodal satellite data," *ISPRS J. Photogramm. Remote Sens.*, vol. 205, pp. 34–49, Nov. 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271623002630>
- [36] Y. Sun, L. Lei, Z. Li, and G. Kuang, "Similarity and dissimilarity relationships based graphs for multimodal change detection," *ISPRS J. Photogramm. Remote Sens.*, vol. 208, pp. 70–88, Feb. 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0924271624000029>
- [37] A. T. Fuller, K. Millard, and J. R. Green, "CROMA: Remote sensing representations with contrastive radar-optical masked autoencoders," in *Proc. Adv. Neural Inf. Process. Syst.*, Jan. 2023, pp. 5506–5538. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2023/hash/11822e84689e631615199d3b75cd0e4-Abstract-Conference.html](https://proceedings.neurips.cc/paper_files/paper/2023/hash/11822e84689e631615199d3b75cd0e4-Abstract-Conference.html)
- [38] X. Guo et al., "SkySense: A multi-modal remote sensing foundation model towards universal interpretation for Earth observation imagery," 2023, *arXiv:2312.10115*.
- [39] Y. Wang, C. M. Albrecht, N. A. A. Braham, C. Liu, Z. Xiong, and X. X. Zhu, "Decoupling common and unique representations for multimodal self-supervised learning," 2023, *arXiv:2309.05300*.
- [40] Z. Xiong et al., "Neural plasticity-inspired multimodal foundation model for Earth observation," 2024, *arXiv:2403.15356*.
- [41] Y. Chen and L. Bruzzone, "Self-supervised SAR-optical data fusion of Sentinel-1/2 images," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5406011. [Online]. Available: <https://ieeexplore.ieee.org/document/9614157>
- [42] Y. Xie, J. Tian, and X. X. Zhu, "A co-learning method to utilize optical images and photogrammetric point clouds for building extraction," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 116, Feb. 2023, Art. no. 103165. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1569843222003533>
- [43] Z. Huang, H. Zhang, A. Laine, E. Angelini, C. Hendon, and Y. Gan, "Co-seg: An image segmentation framework against label corruption," in *Proc. IEEE 18th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2021, pp. 550–553.
- [44] E. Malach and S. Shalev-Shwartz, "Decoupling 'when to update' from 'how to update,'" in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, Jun. 2017, pp. 960–970. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2017/hash/58d4d1e7b1e97b258c9ed0b37e02d087-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2017/hash/58d4d1e7b1e97b258c9ed0b37e02d087-Abstract.html)
- [45] B. Han et al., "Co-teaching: Robust training of deep neural networks with extremely noisy labels," in *Proc. NIPS*, 2018, pp. 8536–8546.
- [46] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. 18th Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, vol. 9351. Cham, Switzerland: Springer, 2015, pp. 234–241.
- [47] C. F. Brown et al., "Dynamic World, near real-time global 10 M land use land cover mapping," *Sci. Data*, vol. 9, no. 1, p. 251, Jun. 2022. [Online]. Available: <https://www.nature.com/articles/s41597-022-01307-4>
- [48] M. Schmitt, L. Hughes, P. Ghamisi, N. Yokoya, and R. Hänsch, "2020 IEEE GRSS data fusion contest," *IEEE Dataport*, 2019. [Online]. Available: <https://dx.doi.org/10.21227/rha7-m332>
- [49] A. Jones, O. Fernández-Ugalde, and S. Scarpa, *LUCAS 2015 Topsoil Survey: Presentation of Dataset and Results*. Luxembourg, Luxembourg: Publications Office of the European Union, 2020. [Online]. Available: <https://data.europa.eu/doi/10.2760/616084>
- [50] M. Schultz, J. Voss, M. Auer, S. Carter, and A. Zipf, "Open land cover from OpenStreetMap and remote sensing," *Int. J. Appl. Earth Observ. Geoinf.*, vol. 63, pp. 206–213, Dec. 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0303243417301605>
- [51] *CORINE Land Cover*. Accessed: Feb. 26, 2024. [Online]. Available: <https://land.copernicus.eu/en/products/corine-land-cover>
- [52] S. Jadon, "A survey of loss functions for semantic segmentation," in *Proc. IEEE Conf. Comput. Intell. Bioinf. Comput. Biol. (CIBCB)*, Oct. 2020, pp. 1–7. [Online]. Available: <https://ieeexplore.ieee.org/document/9277638>
- [53] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2818–2826. [Online]. Available: [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2016/html/Szegedy\\_Rethinking\\_the\\_Inception\\_CVPR\\_2016\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Szegedy_Rethinking_the_Inception_CVPR_2016_paper.html)
- [54] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.

- [55] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2881–2890.
- [56] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder–decoder with atrous separable convolution for semantic image segmentation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Jan. 2018, pp. 833–851.
- [57] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 936–944.
- [58] T. Xiao, Y. Liu, B. Zhou, Y. Jiang, and J. Sun, "Unified perceptual parsing for scene understanding," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, Jan. 2018, pp. 432–448.
- [59] F. Bastani, P. Wolters, R. Gupta, J. Ferdinando, and A. Kembhavi, "SatlasPretrain: A large-scale dataset for remote sensing image understanding," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 16726–16736.
- [60] P. Helber, B. Bischke, A. Dengel, and D. Borth, "EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 7, pp. 2217–2226, Jul. 2019.



**Yi Wang** (Student Member, IEEE) received the B.S. degree in remote sensing science and technology from Wuhan University, Wuhan, China, in 2018, and the M.Sc. degree in geomatics engineering from the University of Stuttgart, Stuttgart, Germany, in 2021. He is currently pursuing the Ph.D. degree with the Technical University of Munich (TUM), Munich, Germany.

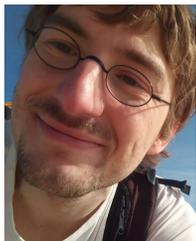
In 2020, he spent three months as a Research Intern with the Perception System Group, Sony Corporation, Stuttgart. From 2021 to 2024, he was a Research Associate with the Remote Sensing Technology Institute, German Aerospace Center (DLR), Weßling, Germany. His research interests include remote sensing, deep learning, computer vision, and self-supervised learning.



**Chenying Liu** (Student Member, IEEE) received the B.S. and M.S. degrees (Hons.) from Sun Yat-sen University, Guangzhou, China, in 2017 and 2020, respectively. She is currently pursuing the Ph.D. degree with the Technical University of Munich, Munich, Germany.

She was a Research Associate with Sun Yat-sen University from 2020 to 2021 and the German Aerospace Center (DLR), Weßling, Germany, from 2021 to 2024. Her research interests include remote sensing and Earth observation, weakly supervised learning, multimodal learning, computer vision, and signal processing.

supervised learning, multimodal learning, computer vision, and signal processing.



**Conrad M. Albrecht** (Member, IEEE) received his PhD in physics from Heidelberg University, Germany, in 2014. Since April 2021 he is a researcher at the Earth Observation Center of the German Aerospace Center (DLR), Oberpfaffenhofen, Germany. As PI of the HelmholtzAI Young Investigator Group (YIG) "Large-Scale Data Mining in Earth Observation" (DM4EO) he advanced weakly-supervised deep learning methodologies for remote sensing. In July 2023, he was appointed a visiting associate professor with the Institute of

Nasca at Yamagata University, Yamagata, Japan contributing to research in machine learning for the UNESCO World Heritage of the Nasca culture in Peru. Jointly with his collaborators, his work has been awarded the 2024 Cozzarelli Prize by the National Academy of Sciences.

For over 6 years Conrad was a research scientist in the Physical Sciences department at the IBM T.J. Watson Research Center in Yorktown, NY, USA. While at the Institute for Theoretical Physics, Heidelberg, Germany, he graduated in physics (International Max Planck Research School for Quantum Dynamics in Physics, Chemistry and Biology) with an extra certification in computer science (Cluster- & Detector Management team at CERN, Switzerland). He received a corresponding Ph.D. degree from Heidelberg University, Germany in 2014 working on distributed computing to study physics at low temperatures.

Conrad's research agenda interconnect physical models and numerical analysis, employing Big Data technologies and machine learning through open-science research, <https://conrad-m-albrecht.github.io>. Conrad co-organized workshops at the IEEE BigData conference, IGARSS conferences, the EGU General Assembly, and the AAAS annual meeting. Conrad is home in Europe and the United States. Some of his transatlantic initiatives include: In 2023, he was invited by the Alexander-von-Humboldt Foundation to present at the German-American Frontiers of Engineering Symposium. Beginning in 2024, and initiated by Helmholtz Imaging and Data Science, he is in close touch with the German Department at Princeton University for cross-atlantic undergraduate internships. In fall 2024, Conrad taught as Adjunct Professor of Earth and Environmental Engineering at Columbia University, NY, USA.



**Xiao Xiang Zhu** (Fellow, IEEE) received the Master (M.Sc.) degree, her doctor of engineering (Dr.-Ing.) degree and her "Habilitation" in the field of signal processing from Technical University of Munich (TUM), Munich, Germany, in 2008, 2011 and 2013, respectively.

She is the Chair Professor for Data Science in Earth Observation at Technical University of Munich (TUM) and was the founding Head of the Department "EO Data Science" at the Remote Sensing Technology Institute, German Aerospace Center (DLR). Since 2022, she serves as the PI and director of the national center of excellence ML4Earth (ml4earth.de). Since May 2020, she is the PI and director of the international future AI lab "AI4EO – Artificial Intelligence for Earth Observation: Reasoning, Uncertainties, Ethics and Beyond", Munich, Germany. Since October 2020, she also serves as a Director of the Munich Data Science Institute (MDSI), TUM. From 2019 to 2022, Zhu has been a co-coordinator of the Munich Data Science Research School ([www.muds.de](http://www.muds.de)) and the head of the Helmholtz Artificial Intelligence – Research Field "Aeronautics, Space and Transport". Prof. Zhu was a guest scientist or visiting professor at the Italian National Research Council (CNR-IREA), Naples, Italy, Fudan University, Shanghai, China, the University of Tokyo, Tokyo, Japan and University of California, Los Angeles, United States in 2009, 2014, 2015 and 2016, respectively. She is currently a visiting AI professor at ESA's Philab, Frascati, Italy. Her main research interests are remote sensing and Earth observation, signal processing, machine learning and data science, with their applications in tackling societal grand challenges, e.g. Global Urbanization, UN's SDGs and Climate Change.

Dr. Zhu has been a member of young academy (Junge Akademie/Junges Kolleg) at the Berlin-Brandenburg Academy of Sciences and Humanities and the German National Academy of Sciences Leopoldina and the Bavarian Academy of Sciences and Humanities. She is a Fellow of the Academia Europaea (the Academy of Europe). She serves in the scientific advisory board in several research organizations, among others the German Research Center for Geosciences (GFZ, 2020-2023) and Potsdam Institute for Climate Impact Research (PIK). She is an associate Editor of IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, *Pattern Recognition* and served as the area editor responsible for special issues of *IEEE Signal Processing Magazine* (2021-2023). She is a Fellow of AIAA, and ELLIS.