

Curriculum Multi-Task Self-Supervision Improves Lightweight Architectures for Onboard Satellite Hyperspectral Image Segmentation

Hugo Carlesso^{1,2}, Josiane Mothe^{1,3}, Radu Tudor Ionescu⁴

¹Univ. Toulouse, IRIT, France ²CNRS, France ³CLLE, CNRS, France ⁴University of Bucharest, Romania

hugo.carlesso@cnrs.fr, josiane.mothe@irit.fr, raducu.ionescu@gmail.com

Abstract—Hyperspectral imaging (HSI) captures detailed spectral signatures across hundreds of contiguous bands per pixel, being indispensable for remote sensing applications such as land-cover classification, change detection, and environmental monitoring. Due to the high dimensionality of HSI data and the slow rate of data transfer in satellite-based systems, compact and efficient models are required to support onboard processing and minimize the transmission of redundant or low-value data. To this end, we introduce a novel curriculum multi-task self-supervised learning (CMTSSL) framework designed for lightweight architectures for HSI analysis. CMTSSL integrates masked image modeling with decoupled spatial and spectral jigsaw puzzle solving, guided by a curriculum learning strategy that progressively increases data difficulty during self-supervision. This enables the encoder to jointly capture fine-grained spectral continuity, spatial structure, and global semantic features. Unlike prior dual-task SSL methods, CMTSSL simultaneously addresses spatial and spectral reasoning within a unified and computationally efficient design, being particularly suitable for training lightweight models for onboard satellite deployment. We validate our approach on four public benchmark datasets, demonstrating consistent gains in downstream segmentation tasks, using architectures that are over 16,000× lighter than some state-of-the-art models. These results highlight the potential of CMTSSL in generalizable representation learning with lightweight architectures for real-world HSI applications. Our code is publicly available at: <https://github.com/hugocarlesso/CMTSSL>.

I. INTRODUCTION

Hyperspectral imaging (HSI) provides detailed spectral signatures across hundreds of contiguous bands per pixel, creating informative *data cubes* that enable fine-grained material recognition and environmental monitoring. This makes HSI an invaluable tool in satellite-based Earth observation, where robust scene understanding is essential for tasks such as land-cover classification, change detection, and resource monitoring [1]–[5]. The operational constraints of space-borne platforms pose unique challenges for HSI data processing. Satellites are required to perform onboard inference under strict computational and energy budgets, often relying on edge devices with limited memory and processing capabilities. Additionally, downlink bandwidth is limited [6], making it crucial to transmit only relevant or high-quality data, e.g. excluding scenes with excessive cloud cover. These constraints demand lightweight and efficient models, capable of operating on edge.

*This project has received financial support from the CNRS through the MITI interdisciplinary programs. This work was also supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS - UEFISCDI, project number PN-IV-P1-PCE-2023-0354, within PNCDI IV.




Fig. 1: Average accuracy vs. number of floating point operations (FLOPs) for various HSI architectures, including both lightweight models (1D Justo [6], 2D Justo [6], CLOLN [7], CUNet++ Reduced [6]) and foundation models (HyperSIGMA-B [8]). Circle areas indicate parameter counts. Performance is reported on the Pavia University dataset [9]. Our CMTSSL boosts performance of lightweight models, without affecting model size and FLOPs.

An orthogonal limitation is that acquiring high-quality labels for supervised training on hyperspectral images remains prohibitively expensive. Due to the scarcity of annotated large-scale HSI datasets, researchers in the community embraced self-supervised learning (SSL) as a means to harness the rich spatial-spectral structure of unlabeled data [10]. SSL frameworks employ label-free pretext tasks to learn meaningful representations, which can further be fine-tuned on downstream tasks with limited supervision [11]. For HSI, such approaches have demonstrated superior performance over traditional supervised baselines in classification, segmentation, and anomaly detection [8]. Among SSL frameworks, contrastive learning and masked image modeling (MIM) [12] have been widely adopted. Yet, these approaches suffer from key limitations. Contrastive methods [13], [14] may fail to capture fine-grained details, MIM pipelines [15] can struggle with high-level semantic separability, and discriminative frameworks [16], [17] often offer limited generalization. Furthermore, SSL methods are not optimized for lightweight architectures designed for onboard satellite processing.

To overcome these challenges, we introduce CMTSSL, a curriculum multi-task self-supervised learning framework for lightweight architectures explicitly designed for onboard hyperspectral image segmentation. CMTSSL integrates MIM with decoupled spatial and spectral jigsaw puzzle solving (JPS) tasks [18], within a curriculum learning paradigm

[19]. We find that 3D gradient magnitudes of HSI data cubes are highly correlated with the capability of solving self-supervised tasks on the respective cubes. Consequently, we design a curriculum learning procedure, organizing the samples into progressively difficult batches, from easy (low gradient magnitude) to hard (high gradient magnitude). By integrating multiple tasks and progressively increasing the difficulty of samples during self-supervision, our framework encourages the joint encoder to gradually learn robust complementary representations spanning spectral continuity, spatial arrangement, and global semantic structure. Our novel learning framework boosts the performance of lightweight architectures without increasing model size or the number of floating point operations (FLOPs), as shown in Figure 1. While the common approach is to train lightweight models via supervised learning [6], [7], we introduce CMTSSL as an encoder-agnostic pretraining stage, to be carried out before the conventional supervised training stage.

We carry out extensive experiments on four standard HSI datasets: Pavia University [9], Pavia Center [9], WHU-HI Hanchuan [20], and HYPISO [21]. The reported results confirm that introducing the CMTSSL stage is beneficial. Moreover, we obtain a new state-of-the-art average accuracy (93.5%) on the large-scale HYPISO benchmark. Hence, our contribution lays the groundwork for robust edge-deployable solutions for satellite-based HSI analysis.

In summary, our contribution is threefold:

- We propose a novel curriculum multi-task learning framework for HSI data, which encourages lightweight encoders to acquire robust representations by learning to solve decoupled spatial-spectral jigsaw puzzles and reconstruct masked data cubes.
- We adapt the jigsaw puzzle solving framework to hyperspectral images, decoupling the task across spatial and spectral dimensions, while using a joint encoder.
- We introduce a novel curriculum learning strategy for HSI data based on 3D gradient magnitudes, which progressively increases sample difficulty to guide the multi-task self-supervised learning.

II. RELATED WORK

A. Self-Supervised Learning

Self-supervised models learn data representations without human-annotated labels. Instead, pretext tasks serve to generate supervisory signals directly from the data, such as learning the arrow of time [22], solving jigsaw puzzles [23], [24], reconstructing masked information [12], [25], [26], or contrasting augmented views [27]–[29]. SSL enables effective training with limited labeled data and supports domain-specific pretraining, thereby mitigating domain gaps [30]. SSL is particularly promising in remote sensing [31], where large volumes of unlabeled Earth observation data are available, yet the cost to manually annotate this data is high. Recent advances demonstrate strong potential for both multispectral and hyperspectral imagery. For instance, TerraMind [32] is a large-scale multimodal foundation model for Earth

observation, while HyperSIGMA [8] and SpectralEarth [33] leverage large HSI datasets (450K and 539K images, respectively) to train foundation models with high transferability. SSL methods in hyperspectral imaging often employ contrastive learning with spectral augmentations [13], [14], MIM for spatial-spectral reconstruction [15], and auxiliary discriminative tasks enforcing spatial-spectral reasoning [17].

SSL exhibits strong potential in remote sensing, but its effectiveness heavily depends on the choice of a single pretext task. A given objective may capture some aspects of the data, such as invariance to augmentation or local spatial structures, while neglecting others, leading to suboptimal representations. To overcome this limitation, some recent frameworks [34]–[36] integrated multiple SSL tasks to obtain more robust representations. For example, CMID [34] fuses contrastive learning with masked image modeling through self-distillation, yielding representations that are both semantically rich and spatially grounded. However, solving multiple tasks can lead to an arduous learning process, requiring a careful balancing between learning objectives [37]. Unlike existing multi-task SSL frameworks for HSI, we introduce a curriculum learning strategy to gradually increase task complexity by adding more difficult images, guiding the learning process towards the desired objective.

B. Curriculum Learning

Curriculum learning [19] is a training strategy inspired by human learning: models first learn from simple examples before tackling more difficult ones. Replacing random data ordering with an easy-to-hard progression can speed up convergence and enhance generalization. A typical curriculum learning framework consists of two main elements: a difficulty criterion, which ranks training samples from easy to hard, and a scheduler, which determines when increasingly difficult examples are introduced during training. Curriculum learning can be applied at various levels, including the data [19], [38], [39], the task [40], the model capacity [41]–[43], and even the optimization objective [44].

Curriculum learning has also been applied in SSL frameworks based on masked information modeling [45], [46]. For instance, CL-MAE [45] integrates a learnable module that selects image patches for masking based on the difficulty to reconstruct the masked patches. However, the module and the masked autoencoder need to be trained in alternating steps. EH-MAM [46] is an audio SSL framework that employs a teacher model to decide which acoustic frames to mask. Unlike such frameworks [45], [46], our strategy exploits the correlation between SSL tasks and image gradient magnitudes to create a data-level curriculum, requiring no extra model to generate the curriculum. To the best of our knowledge, a data-level curriculum based on 3D gradients for SSL on HSI data has not been explored in previous literature. Our difficulty criterion stems from our observation that gradient magnitudes are well correlated to the SSL tasks integrated into the proposed framework.

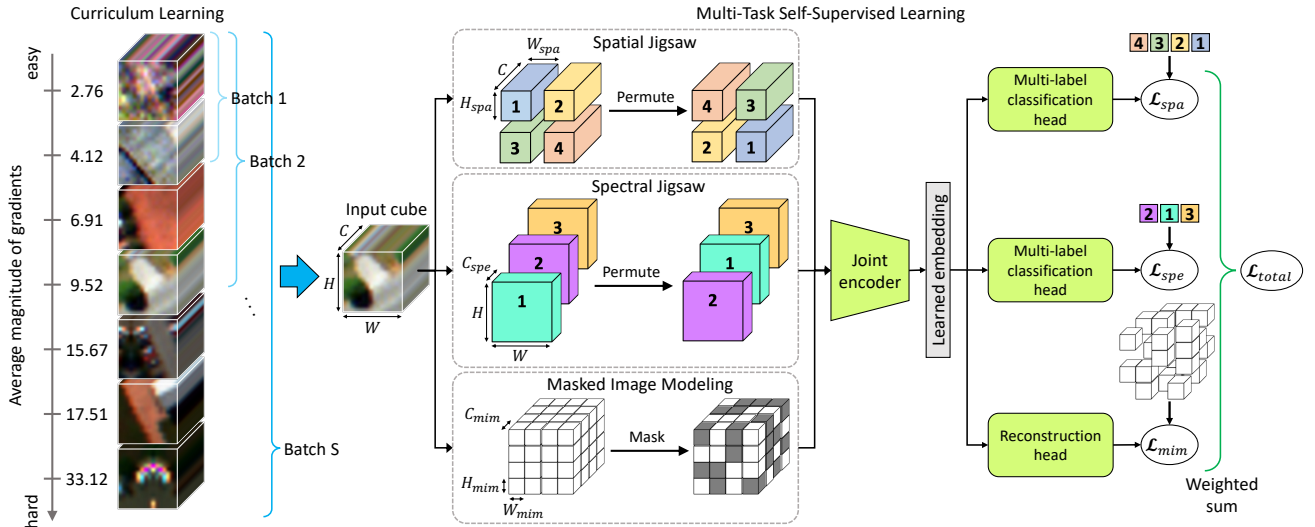


Fig. 2: Curriculum multi-task self-supervised learning (CMTSSL) pipeline. Input HSI cubes are sorted according to their 3D gradient magnitudes and divided into S curriculum batches. To generate the pretext tasks, each input cube undergoes three parallel transformations: random spatial permutation, spectral permutation, and patch masking. The transformed cubes are encoded via a shared encoder and routed to task-specific heads, namely spatial permutation prediction, spectral permutation prediction, and masked image modeling (MIM). The model is jointly optimized via a weighted loss. Best viewed in color.

C. Lightweight HSI Architectures

In remote sensing, especially for onboard processing, lightweight architectures that preserve high accuracy are of growing interest [6], [7], [47]. Li et al. [7] introduced the Channel-Layer-Oriented Lightweight Spectral-Spatial Network (CLOLN), which explicitly models both spectral and spatial interactions, while minimizing parameter and computational overhead. Justo et al. [6] proposed an entire lineup of lightweight 1D or 2D CNNs for onboard HSI segmentation. For example, 1D Justo-LiuNet [47] is a compact 1D CNN focused on spectral feature extraction, which was deployed on a satellite. Despite having only about 4,500 parameters, it outperforms more complex models, such as fast vision transformers (FastViT) [48]. Existing lightweight models suitable for deployment in scenarios with severe resource constraints, such as in-orbit processing, are all based on supervised learning. To the best of our knowledge, we are the first to show that multi-task SSL is beneficial for such lightweight models.

III. METHOD

We propose a multi-task self-supervised learning (MTSSL) framework that integrates decoupled spatial and spectral jigsaw puzzle solving tasks with masked image modeling (see Figure 2). We hypothesize that integrating multiple SSL tasks can enable the shared encoder to learn complementary cues about local continuity, structural arrangement, and global semantics within a single paradigm, essentially leading to a more robust representation (embedding). However, the integration of multiple objectives can make learning difficult and can even result in worse performance for each task than in single-task SSL [37]. To mitigate this issue, we propose a curriculum learning schedule constructed around the *saliency* of 3D input images, which is estimated via 3D image gradient

magnitudes. This approach is based on our preliminary findings suggesting that the difficulty of reconstructing masked patches or rearranging visible patches increases along with the saliency level of the input samples. Unlike heuristic loss balancing or auxiliary difficulty estimators, our scheme is data-driven, architecture-agnostic, and requires no extra model, mitigating negative transfer, while preserving the benefits of multiple SSL objectives. We formally present our curriculum multi-task self-supervised learning (CMTSSL) framework in Algorithm 1. We continue with a detailed description of each component, while referring to the corresponding steps of the algorithm along the way.

A. Mathematical Preliminaries and Notations

In the remainder of this work, we use bold letters to denote vectors, matrices and tensors, round letters to denote sets, lowercase letters to denote neural blocks, and uppercase letters to denote scalar values. We use \odot to represent the convolution operation, and \mathbf{X}^{*P} to represent the element-wise raise of matrix \mathbf{X} to power P . We denote a hyperspectral image (data cube) as $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$, where H is the height, W is the width and C is the number of channels (bands).

B. Curriculum Learning via Gradient Magnitudes

1) *Intuition and empirical motivation*: To quantify the visual difficulty (complexity) of input images, we introduce a criterion based on image gradients. Our approach leverages the intuition that images with higher gradient magnitudes are more salient, containing richer spatial-spectral structures, e.g. sharp edges, textures, or abrupt spectral transitions, and are therefore more challenging to model than smooth (homogeneous) images. To validate our intuition, we conduct preliminary experiments on the large-scale HYPSON dataset [21] and measure the correlation between image gradients and loss values for the masked image modeling and jigsaw puzzle solving tasks. At the same time, we consider

Algorithm 1 CMTSSL Framework

Require: $\mathcal{D} = \{\mathbf{I}_i\}_{i=1}^N$ - set of N training images; $\mathbf{K}_x, \mathbf{K}_y, \mathbf{K}_z$ - filters to compute gradients; S - number of curriculum batches; K - initial number of epochs; F - growth factor for the number of epochs; E - number of epochs for current batch; f_θ - joint encoder with weights θ ; h_ϕ^{spa} - spatial jigsaw multi-label classification head with weights ϕ ; h_β^{spe} - spectral jigsaw multi-label classification head with weights β ; h_γ^{mim} - masked patch reconstruction head with weights γ ; M - number of patches to be masked; $\alpha_{\text{spa}}, \alpha_{\text{spe}}, \alpha_{\text{mim}}$ - weights used for the individual losses; η - learning rate.

Ensure: f^θ - encoder pretrained via CMTSSL.

```

1:  $\mathcal{G} \leftarrow \emptyset$ 
2: for  $i = 1$  to  $N$  do
3:    $\mathbf{G}_x \leftarrow \mathbf{I}_i \odot \mathbf{K}_x, \mathbf{G}_y \leftarrow \mathbf{I}_i \odot \mathbf{K}_y, \mathbf{G}_z \leftarrow \mathbf{I}_i \odot \mathbf{K}_z$ 
4:    $G_i \leftarrow \text{average}(\sqrt{\mathbf{G}_x^{*2} + \mathbf{G}_y^{*2} + \mathbf{G}_z^{*2}})$ 
5:    $\mathcal{G} \leftarrow \mathcal{G} \cup \{G_i\}$ 
6: end for
7:  $\mathcal{D} \leftarrow \text{sort}(\mathcal{D}, \text{criterion} = \mathcal{G}, \text{ord} = \text{'asc'})$ 
8: Curriculum multi-task self-supervised learning
9:  $E \leftarrow K$ 
10: for batch  $k = 1$  to  $S$  do
11:   # Select data for batch  $k$ 
12:    $R \leftarrow \lfloor N \cdot \frac{k}{S} \rfloor$ 
13:    $\mathcal{D}_k \leftarrow \{\mathbf{I}_i \in \mathcal{D} \mid i \leq R\}$ 
14:   for epoch  $e = 1$  to  $E$  do
15:     for  $\mathbf{I} \in \mathcal{D}_k$  do
16:       # Set up each pretext task
17:        $(\mathbf{I}^{\text{spa}}, \mathbf{y}^{\text{spa}}) \leftarrow \text{spatialJigsaw}(\mathbf{I})$ 
18:        $(\mathbf{I}^{\text{spe}}, \mathbf{y}^{\text{spe}}) \leftarrow \text{spectralJigsaw}(\mathbf{I})$ 
19:        $(\mathbf{I}^{\text{vis}}, \mathbf{I}^{\text{mask}}) \leftarrow \text{masking}(\mathbf{I}, M)$ 
20:       # Compute predictions
21:        $\hat{\mathbf{y}}^{\text{spa}} \leftarrow h_\phi^{\text{spa}}(f_\theta(\mathbf{I}^{\text{spa}}))$ 
22:        $\hat{\mathbf{y}}^{\text{spe}} \leftarrow h_\beta^{\text{spe}}(f_\theta(\mathbf{I}^{\text{spe}}))$ 
23:        $\hat{\mathbf{I}}^{\text{mask}} \leftarrow h_\gamma^{\text{mim}}(f_\theta(\mathbf{I}^{\text{vis}}))$ 
24:       # Compute losses
25:        $\mathcal{L}_{\text{spa}} = \text{BCE}(\mathbf{y}^{\text{spa}}, \hat{\mathbf{y}}^{\text{spa}})$ 
26:        $\mathcal{L}_{\text{spe}} = \text{BCE}(\mathbf{y}^{\text{spe}}, \hat{\mathbf{y}}^{\text{spe}})$ 
27:        $\mathcal{L}_{\text{mim}} = \text{MAE}(\mathbf{I}^{\text{mask}}, \hat{\mathbf{I}}^{\text{mask}})$ 
28:        $\mathcal{L}_{\text{total}} = \alpha_{\text{spa}} \cdot \mathcal{L}_{\text{spa}} + \alpha_{\text{spe}} \cdot \mathcal{L}_{\text{spe}} + \alpha_{\text{mim}} \cdot \mathcal{L}_{\text{mim}}$ 
29:       # Update weights via backpropagation
30:        $\phi \leftarrow \phi - \eta \cdot \nabla_\phi \mathcal{L}_{\text{spa}}$ 
31:        $\beta \leftarrow \beta - \eta \cdot \nabla_\beta \mathcal{L}_{\text{spe}}$ 
32:        $\gamma \leftarrow \gamma - \eta \cdot \nabla_\gamma \mathcal{L}_{\text{mim}}$ 
33:        $\theta \leftarrow \theta - \eta \cdot \nabla_\theta \mathcal{L}_{\text{total}}$ 
34:     end for
35:   end for
36:    $E \leftarrow E \cdot F$ 
37: end for

```

various ways to aggregate image gradients and obtain a representative value at the image level, namely via computing the average, standard deviation, or maximum values. The correlations reported in Table I confirm our intuition, namely that gradient magnitudes are well correlated to SSL task objectives, regardless of the aggregation operator. The results also indicate that the average operator is more correlated with both task losses than the other operators. We thus employ the average gradient magnitude in our framework.

2) *Gradient computation:* For each image \mathbf{I} , we compute local gradients along both spatial and spectral dimensions. Spatial gradients are estimated using Scharr operators [49] in

TABLE I: Pearson correlation coefficients between individual task losses (MIM and JPS) and difficulty criteria computed on the HYPSSO [21] validation set. Gradient magnitudes are positively correlated with SSL task losses, regardless of the gradient aggregation operator.

| Gradient aggregation | MIM | JPS |
|----------------------|--------------|--------------|
| Average | 0.847 | 0.340 |
| Maximum | 0.742 | 0.241 |
| Standard Deviation | 0.521 | 0.227 |

the horizontal (x) and vertical (y) directions, while spectral differences between adjacent channels provide a proxy for inter-band variability:

$$\mathbf{G}_x = \mathbf{I} \odot \mathbf{K}_x, \quad \mathbf{G}_y = \mathbf{I} \odot \mathbf{K}_y, \quad \mathbf{G}_z = \mathbf{I} \odot \mathbf{K}_z, \quad (1)$$

where \mathbf{K}_x and \mathbf{K}_y are Scharr kernels, and \mathbf{K}_z denotes a kernel that performs the difference on the spectral axis (z), i.e. $[1, -1]$. The average gradient magnitude for image \mathbf{I} is given by:

$$G = \text{average}\left(\sqrt{\mathbf{G}_x^{*2} + \mathbf{G}_y^{*2} + \mathbf{G}_z^{*2}}\right). \quad (2)$$

The average gradient magnitudes are computed in steps 1-6 of Algorithm 1.

3) *Data organization and curriculum schedule:* The average gradient magnitude is used as a proxy for image difficulty in our curriculum learning framework. Images with low average gradient magnitude values (smooth, homogeneous content) are considered *easy*, while those with high gradient magnitude (salient content, complex patterns) are considered *hard*. We sort the samples in ascending order according to their average gradient magnitudes (step 7 of Algorithm 1). This ordering allows the model to first capture global spectral-spatial regularities, before tackling more challenging images containing fine-grained or high-frequency structures.

We next divide the sorted samples into S curriculum batches¹. Let \mathcal{D} denote the sorted dataset of N training images computed in step 7. Each batch \mathcal{D}_k contains all images $\mathbf{I}_i \in \mathcal{D}$, such that $i \leq R$ (step 11), where $R = \lfloor N \cdot \frac{k}{S} \rfloor$ (step 10). This selection procedure ensures that, in early stages, the training batches contain the easiest images, while harder images are progressively introduced until the entire dataset is covered. Notice that $\mathcal{D}_k \subset \mathcal{D}_{k+1}, \forall k < S$, which is required to avoid catastrophic forgetting [50].

We partition the total training budget (expressed in epochs) across the S curriculum batches, according to a schedule parameterized by the number of epochs $K \in \mathbb{N}^+$ for the first batch (step 8) and the growth factor $F > 0$ (step 30). Thus, the number of epochs on batch \mathcal{D}_k is given by $E = K \cdot F^{k-1}$, where k is the curriculum batch index.

C. Multi-Task Learning

We propose an SSL framework that comprises three tasks: spatial jigsaw puzzle solving, spectral jigsaw puzzle solving, and MIM. Jigsaw puzzle solving (JPS) [23] is a self-supervised task that produces semantically relevant representations. Inspired by Wang et al. [18], who decoupled the JPS task across spatial and temporal dimensions to

¹Not to be confused with mini-batches. A curriculum batch is much larger, containing multiple training mini-batches.

detect abnormal events in video, we adapt the decoupled JPS version to capture both spatial and spectral structures in hyperspectral imagery. While JPS encourages features that capture contextual relationships between image parts, MIM promotes features that capture fine-grained local details. Since these SSL tasks focus on different aspects, it makes sense to try to combine them to obtain more robust representations. We next present how each task is constructed and how the whole architecture is integrated.

1) *Spatial jigsaw puzzle solving*: Each input image is split into non-overlapping 3D patches across the spatial dimensions, such that the size of a 3D patch is $H_{\text{spa}} \times W_{\text{spa}} \times C$, where $H_{\text{spa}} < H$ and $W_{\text{spa}} < W$, respectively. The patches are permuted and reassembled into an output image denoted as \mathbf{I}^{spa} , while the permutation indices are used to compute the target labels for the pretext classification task. Let $N_{\text{spa}} = \frac{H}{H_{\text{spa}}} \cdot \frac{W}{W_{\text{spa}}}$ denote the number of 3D patches. To avoid the combinatorial explosion associated with classifying permutations into different classes, we frame the spatial JPS task as a multi-label classification problem, i.e. the position of each 3D patch needs to be independently predicted. This essentially reduces the number of output probabilities from $N_{\text{spa}}!$ to N_{spa}^2 . The multi-label targets are stored in \mathbf{y}^{spa} . The procedure to generate the spatial JPS task is executed in step 14 of Algorithm 1.

2) *Spectral jigsaw puzzle solving*: For the spectral JPS task, we permute contiguous spectral blocks, where the size of each block is $H \times W \times C_{\text{spe}}$, with $C_{\text{spe}} < C$. The number of spectral blocks N_{spe} is given by $N_{\text{spe}} = \frac{C}{C_{\text{spe}}}$. We adopt the same multi-label classification formulation as before, reducing the factorial number of classes to N_{spe}^2 . The procedure to generate the spectral JPS task executed in step 15 returns the shuffled input \mathbf{I}^{spe} and the target labels \mathbf{y}^{spe} .

3) *Masked image modeling*: For the MIM task, each hyperspectral image \mathbf{I} is divided into non-overlapping 3D patches, such that the size of a patch is $H_{\text{mim}} \times W_{\text{mim}} \times C_{\text{mim}}$, where $H_{\text{mim}} < H$, $W_{\text{mim}} < W$ and $C_{\text{mim}} < C$, respectively. A randomly chosen subset of M patches is masked, resulting in a visible version of the input, denoted as \mathbf{I}^{vis} , and a masked version, denoted as \mathbf{I}^{mask} (step 16). The visible patches \mathbf{I}^{vis} are given as input to the model, which is supposed to reconstruct the masked patches \mathbf{I}^{mask} .

4) *Multi-task learning architecture*: Our architecture comprises a shared encoder and three task-specific heads (see Figure 2). The shared encoder f_{θ} learns a joint representation space for all SSL tasks. One head, denoted as h_{γ}^{mim} , is dedicated to the MIM task, while the other two, h_{β}^{spa} and h_{ϕ}^{spe} , correspond to the spatial and spectral JPS tasks.

Each task generation procedure takes the training image \mathbf{I} as input and returns the task-specific input and target (steps 14-16). The resulting inputs are then passed through the shared encoder and routed to their respective task-specific heads. Each head produces a task-specific output (steps 17-19), which is compared with the corresponding target via a task-specific loss function (steps 20-22). The spatial and spectral JPS losses, denoted as \mathcal{L}_{spa} and \mathcal{L}_{spe} , predict the corresponding permutations via binary cross-entropy (BCE).

The reconstruction loss \mathcal{L}_{mim} is computed as the mean absolute error (MAE) between predicted and ground-truth pixel values in the masked regions. The three losses are subsequently combined via a weighted sum (step 23), modulated by the hyperparameters $\alpha_{\text{spa}}, \alpha_{\text{spe}}, \alpha_{\text{mim}}$, respectively. The final training objective is defined as follows:

$$\mathcal{L}_{\text{total}} = \alpha_{\text{spa}} \cdot \mathcal{L}_{\text{spa}} + \alpha_{\text{spe}} \cdot \mathcal{L}_{\text{spe}} + \alpha_{\text{mim}} \cdot \mathcal{L}_{\text{mim}}. \quad (3)$$

We update the parameters ϕ, β, γ of the task-specific heads via the corresponding loss functions (steps 24-26). The shared encoder f_{θ} is trained by back-propagating gradients from all three heads via $\mathcal{L}_{\text{total}}$ (step 27).

Upon completing the pretraining phase via CMTSSL, the pretrained encoder f_{θ} can be fine-tuned on downstream tasks by attaching an appropriate task-specific head.

IV. EXPERIMENTS

We conduct segmentation (a.k.a. pixel-wise classification) experiments to compare single-task and multi-task SSL, and to assess the benefits of curriculum learning in the context of multi-task SSL. We also compare the lightweight models trained via CMTSSL with state-of-the-art models, including both lightweight architectures and computationally-intensive foundation models. Finally, we perform ablation studies on various hyperparameters to complement the main results.

A. Datasets

The *HYPISO-1 “Sea-Land-Cloud”* dataset [21] comprises 179 hyperspectral images captured by the HYPISO-1 satellite across diverse global locations. Among these, 38 images include pixel-level annotations, providing approximately 25 million labeled spectral signatures for the sea, land, and cloud classes. The remaining 141 unlabeled images support unsupervised and self-supervised learning. Each image has a spatial resolution of 956×684 pixels with 120 spectral bands (400-800nm). The official split consists of 30 training, 3 validation, and 5 test images.

The *Pavia University (PU)* and *Pavia Center (PC)* datasets contain hyperspectral images acquired by the ROSIS sensor during an aerial survey over Pavia, northern Italy. Both provide a spatial resolution of 1.3m per pixel and include ground-truth annotations for nine land-cover types and a background class. PU consists of a single scene of 610×340 pixels with 42,776 labeled pixels (20% of total pixels) across 103 spectral bands (430-860nm). PC comprises a scene of 1096×715 pixels with 148,152 labeled pixels (19%) across 102 bands in the same range.

The *WHU-Hi-HanChuan (HC)* dataset is part of the WHU-Hi benchmark suite released by the RSDIA group at Wuhan University. It was captured using an UAV-mounted hyperspectral camera at 250m altitude. The scene contains 1217×303 pixels at 0.109m spatial resolution, with 274 spectral bands spanning 400-1000nm. It includes annotations for 16 distinct classes.

B. Experimental Setup

1) *Data preprocessing*: Following standard preprocessing in previous work [6], we divide the original high-resolution images from all datasets into subimages of 16×16 pixels,

TABLE II: Number of subimages per split for all datasets.

| Dataset | Pretraining | Training | Validation | Testing |
|---------|-------------|----------|------------|-----------|
| PU | 1,653 | 180 | 90 | 24,873 |
| PC | 6,364 | 180 | 90 | 90,146 |
| HC | 1,681 | 800 | 400 | 171,652 |
| HYPISO | 441,180 | 77,400 | 7,740 | 3,269,520 |

while preserving the number of channels (spectral bands). To ensure consistent input distributions and stabilize training, we normalize the data by computing the mean and standard deviation for each spectral band. For the three single-scene datasets (PU, PC and HC), we partition each scene into non-overlapping regions. These regions are distributed into training, validation, and test sets, to prevent data leakage during evaluation. For pretraining, we sample additional 16×16 subimages from the training regions, using an overlap (stride) of 8 pixels. In Table II, we report the resulting number of subimages (later called images) for each dataset.

2) *Chosen architectures:* We evaluate our CMTSSL framework on three lightweight architectures: 2D Justo-UNet-Simple (2D Justo) [6], 2D CUNet++ Reduced [6], and CLOLN [7]. The first two architectures are the best and second-best 2D CNN models on the HYPISO dataset [21]. All chosen architectures are suitable for onboard processing, containing between 4K and 11K learnable parameters.

3) *State-of-the-art models:* Aside from the chosen architectures, we compare with additional models, namely the HyperSIGMA-B foundation model [8], as well as the lightweight 1D Justo-LiuNet [6] and FastViT-S12 [48].

4) *Hyperparameter tuning:* All models are optimized with AdamW using a learning rate of $5 \cdot 10^{-4}$ and a mini-batch size of 16. Each model is pretrained for 120-250 epochs, depending on the dataset. The masking ratio (0.6) and patch dimensions are kept consistent across all experiments. The loss weights in the CMTSSL framework are carefully tuned to balance the optimization of individual objectives. More specifically, they are scaled to ensure that the gradient magnitudes of each loss component remain comparable, thereby preventing any single objective from disproportionately influencing the training process. Curriculum hyperparameters (S , K and F) are set in a dataset-dependent manner, considering the following ranges: $S \in \{3, 4, 5\}$, $K \in \{10, 11, \dots, 40\}$, $F \in \{1, 1.1, \dots, 2\}$. They are chosen such that the total number of batches processed during pretraining remains consistent between the curriculum-based and non-curriculum setups, ensuring a fair comparison. We tune hyperparameters using grid-search on validation data.

5) *Evaluation metrics:* We evaluate models using three commonly employed metrics in hyperspectral classification: Overall Accuracy (OA), which reflects the proportion of correctly classified pixels; Average Accuracy (AA), the mean per-class accuracy; and the Kappa coefficient, which adjusts OA by accounting for the accuracy expected by chance. These metrics provide a comprehensive assessment of model accuracy. Results are averaged over 3 runs to ensure statistical reliability.

6) *Computational environment:* Experiments are carried out on high-performance GPUs (NVIDIA A40 and Quadro

TABLE III: State-of-the-art comparison on the Pavia University, Pavia Center, WHU-Hancuan and HYPISO datasets. CMTSSL is applied on 2D Justo, CUNet++ Reduced, and CLOLN, respectively. The best score on each dataset is highlighted in **bold green**. * – results are taken from [6].

| | Model | AA | Parameters | FLOPs |
|------------------|----------------------------------|----------------------------------|------------|-------------|
| Pavia University | HyperSIGMA-B [8] | 62.2 ± 2.0 | 176,700K | 66,800,336K |
| | 1D Justo-LiuNet [6] | 65.6 ± 2.1 | 4.6K | 51K |
| | 2D Justo [6] | 73.5 ± 3.3 | 9K | 1,494K |
| | +CMTSSL (ours) | 75.8 ± 0.3 | 9K | 1,494K |
| | CUNet++ Reduced [6] | 59.8 ± 1.8 | 11K | 1,949K |
| | +CMTSSL (ours) | 68.1 ± 1.5 | 11K | 1,949K |
| | CLOLN [7] | 78.2 ± 5.6 | 3.5K | 926K |
| +CMTSSL (ours) | 79.4 ± 2.5 | 3.5K | 926K | |
| Pavia Center | HyperSIGMA-B [8] | 69.3 ± 5.5 | 176,700K | 66,800,336K |
| | 1D Justo-LiuNet [6] | 76.8 ± 0.3 | 4.6K | 51K |
| | 2D Justo [6] | 79.1 ± 6.8 | 9K | 1,494K |
| | +CMTSSL (ours) | 83.0 ± 8.0 | 9K | 1,494K |
| | CUNet++ Reduced [6] | 74.1 ± 4.7 | 11K | 1,949K |
| | +CMTSSL (ours) | 77.3 ± 4.2 | 11K | 1,949K |
| | CLOLN [7] | 82.7 ± 2.6 | 3.5K | 926K |
| +CMTSSL (ours) | 85.7 ± 4.2 | 3.5K | 926K | |
| WHU-HI Hancuan | HyperSIGMA-B [8] | 64.3 ± 2.5 | 176,700K | 66,800,336K |
| | 1D Justo-LiuNet [6] | 55.0 ± 1.0 | 9K | 206K |
| | 2D Justo [6] | 61.1 ± 1.9 | 18K | 3,827K |
| | +CMTSSL (ours) | 61.7 ± 1.5 | 18K | 3,827K |
| | CUNet++ Reduced [6] | 53.5 ± 1.3 | 23K | 5,045K |
| | +CMTSSL (ours) | 53.6 ± 0.2 | 23K | 5,046K |
| | CLOLN [7] | 60.9 ± 0.4 | 6K | 1,614K |
| +CMTSSL (ours) | 63.5 ± 2.4 | 6K | 1,614K | |
| HYPISO | FastViT-S12* [48] | 90.0 | 10,243K | - |
| | 1D Justo-LiuNet* [6] | 93.0 | 5K | 59K |
| | 2D Justo [6] | 92.9 ± 0.3 | 9K | 1,642K |
| | +CMTSSL (ours) | 93.5 ± 0.2 | 9K | 1,642K |
| | CUNet++ Reduced [6] | 92.8 ± 0.6 | 12K | 2,132K |
| | +CMTSSL (ours) | 92.9 ± 0.8 | 12K | 2,132K |
| | CLOLN [7] | 91.9 ± 0.1 | 4K | 967K |
| +CMTSSL (ours) | 92.4 ± 0.3 | 4K | 967K | |

RTX 6000) provided by a shared scientific computing facility, enabling consistent evaluation across multiple datasets.

C. Results

1) *Comparison with state of the art:* In Table III, we compare CMTSSL with various state-of-the-art models in terms of average accuracy, parameter count and FLOPs. For a fair comparison, all models are trained and tested on the same data splits. CMTSSL boosts the performance of all three models that is applied on, namely 2D Justo, CUNet++ Reduced, and CLOLN. Notably, the performance gains come without increasing parameter count or FLOPs. Lightweight models often outperform larger models, indicating that models optimized for onboard processing do not have to compromise on accuracy. On the HYPISO dataset, CMTSSL boosts the average accuracy of 2D Justo to 93.5% (+0.6%), setting a new state-of-the-art score that surpasses the previous best (93%), reported by 1D Justo-LiuNet [6]. Large standard deviations on the three one-scene datasets (PU, PC and HC) are likely due to the small amount of training samples per category (see Table II).

2) *Comparison of training procedures:* To assess the benefits of self-supervised pretraining, we compare models trained from scratch (with full supervision) with models that undergo fine-tuning, after pretraining with one the following SSL configurations: MIM only, decoupled JPS only, MTSSL, and CMTSSL (ours). In Table IV, we present the comparative

TABLE IV: Results of lightweight 2D Justo, CUNet++ and CLOLN encoders trained from scratch or pretrained via MIM, JPS, MTSSL, and CMTSSL, on the Pavia University, Pavia Center, WHI-Hanchuan and HYPISO datasets. The best and second-best scores of each model on each dataset are highlighted in **bold green** and underline blue, respectively.

| Data set | Training strategy | 2D Justo [6] | | | CUNet++ Reduced [6] | | | CLOLN [7] | | |
|------------------|-------------------|-------------------|-------------------|-------------------|---------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
| | | AA | OA | Kappa | AA | OA | Kappa | AA | OA | Kappa |
| Pavia University | From scratch | 73.5 ± 3.3 | 59.3 ± 4.8 | 51.5 ± 4.8 | 59.8 ± 1.8 | 53.0 ± 4.4 | 43.2 ± 4.7 | <u>78.2 ± 4.3</u> | <u>70.3 ± 5.6</u> | <u>64.1 ± 7.0</u> |
| | MIM | 75.5 ± 1.8 | 64.4 ± 1.6 | 56.7 ± 2.0 | 62.4 ± 1.0 | 52.0 ± 1.7 | 42.5 ± 1.7 | 74.7 ± 2.6 | 63.5 ± 2.7 | 55.8 ± 3.4 |
| | JPS | 73.8 ± 1.1 | 61.3 ± 3.0 | 53.9 ± 3.0 | 66.5 ± 3.3 | 57.0 ± 3.0 | 48.3 ± 3.3 | 77.0 ± 4.0 | 67.8 ± 6.0 | 61.4 ± 7.0 |
| | MTSSL | 72.7 ± 1.9 | 59.7 ± 2.8 | 51.6 ± 3.2 | 68.5 ± 1.2 | 55.4 ± 3.6 | 46.7 ± 3.4 | 75.3 ± 3.1 | 68.8 ± 2.7 | 61.7 ± 3.4 |
| | CMTSSL (ours) | 75.8 ± 0.3 | 64.8 ± 0.9 | 57.5 ± 0.8 | 68.1 ± 1.5 | 57.2 ± 1.3 | 48.8 ± 1.4 | 79.4 ± 2.5 | 70.6 ± 5.8 | 64.9 ± 7.0 |
| Pavia Center | From scratch | 79.1 ± 6.8 | 91.6 ± 3.8 | 87.5 ± 5.8 | 74.1 ± 4.7 | 86.1 ± 2.3 | 79.4 ± 3.4 | 82.7 ± 2.6 | 93.0 ± 2.6 | 89.5 ± 4.0 |
| | MIM | 83.7 ± 3.4 | 93.4 ± 2.2 | 90.1 ± 3.2 | 77.4 ± 3.4 | 87.8 ± 2.1 | 82.0 ± 3.0 | 83.4 ± 0.7 | 93.5 ± 1.0 | 90.3 ± 1.4 |
| | JPS | 81.8 ± 2.4 | 92.2 ± 1.0 | 88.4 ± 1.4 | 75.0 ± 4.4 | 87.5 ± 0.9 | 81.5 ± 1.4 | 84.0 ± 0.2 | 93.5 ± 0.7 | 90.3 ± 0.9 |
| | MTSSL | 82.9 ± 1.9 | 92.1 ± 0.7 | 88.3 ± 1.0 | 76.5 ± 2.2 | 87.9 ± 0.8 | 82.1 ± 1.1 | 82.2 ± 1.3 | 92.9 ± 1.4 | 89.4 ± 2.2 |
| | CMTSSL (ours) | 83.0 ± 8.0 | 92.4 ± 3.9 | 88.6 ± 5.7 | 77.3 ± 4.2 | 88.0 ± 2.4 | 82.2 ± 3.6 | 85.7 ± 4.2 | 94.7 ± 1.5 | 92.1 ± 2.2 |
| WHU-HI Hanchuan | From scratch | 61.1 ± 1.9 | 73.2 ± 1.6 | 69.5 ± 1.8 | 53.5 ± 1.3 | 65.0 ± 2.0 | 60.3 ± 2.1 | 60.9 ± 0.4 | 71.9 ± 3.7 | 68.0 ± 1.3 |
| | MIM | 58.6 ± 0.7 | 72.5 ± 1.4 | 68.6 ± 1.5 | 52.1 ± 1.6 | 64.3 ± 1.6 | 59.6 ± 1.7 | 61.5 ± 1.4 | 71.5 ± 0.5 | 67.7 ± 0.5 |
| | JPS | 59.2 ± 1.8 | 69.8 ± 3.0 | 65.8 ± 3.3 | 51.9 ± 2.4 | 60.5 ± 5.4 | 55.3 ± 5.7 | 61.1 ± 5.0 | 70.9 ± 4.0 | 66.9 ± 4.4 |
| | MTSSL | 62.1 ± 1.7 | 74.5 ± 0.6 | 70.9 ± 0.6 | 52.5 ± 1.6 | 66.1 ± 0.6 | 61.7 ± 0.7 | 62.6 ± 0.8 | 71.9 ± 1.4 | 68.1 ± 1.6 |
| | CMTSSL (ours) | 61.7 ± 1.5 | 73.8 ± 1.1 | 70.1 ± 1.1 | 53.6 ± 0.2 | 65.1 ± 0.9 | 60.1 ± 1.1 | 63.5 ± 2.4 | 72.3 ± 1.9 | 68.4 ± 2.2 |
| HYPISO | From scratch | 92.9 ± 0.3 | 95.4 ± 0.2 | 92.4 ± 0.4 | 92.8 ± 0.6 | 94.7 ± 0.1 | 91.2 ± 0.2 | 91.9 ± 0.1 | 94.7 ± 0.2 | 91.2 ± 0.3 |
| | MIM | 92.6 ± 0.6 | 94.8 ± 0.0 | 91.4 ± 0.1 | 92.7 ± 0.4 | 94.4 ± 0.5 | 90.7 ± 0.9 | 92.3 ± 0.9 | 94.7 ± 0.6 | 91.3 ± 0.9 |
| | JPS | 93.0 ± 0.2 | 95.3 ± 0.2 | 92.2 ± 0.3 | 92.9 ± 0.3 | 95.3 ± 0.3 | 92.2 ± 0.4 | 91.7 ± 0.8 | 95.1 ± 0.5 | 91.8 ± 0.8 |
| | MTSSL (ours) | 93.3 ± 0.7 | 95.5 ± 0.5 | 92.5 ± 0.8 | 93.1 ± 0.5 | 94.9 ± 1.0 | 91.6 ± 1.6 | 91.6 ± 1.3 | 94.8 ± 0.4 | 91.4 ± 0.8 |
| | CMTSSL (ours) | 93.5 ± 0.2 | 95.5 ± 0.1 | 92.6 ± 0.3 | 92.9 ± 0.8 | 94.9 ± 1.0 | 91.6 ± 1.5 | 92.4 ± 0.3 | 95.2 ± 0.4 | 92.0 ± 0.7 |

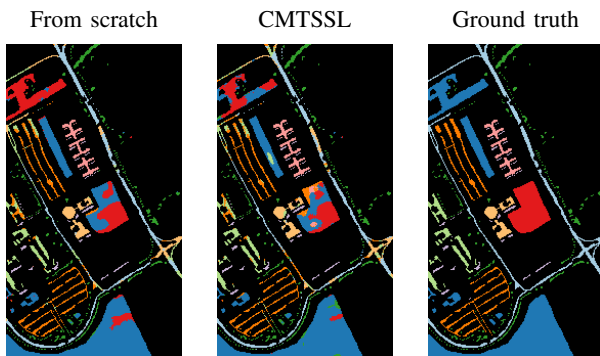


Fig. 3: Visual comparison on the Pavia University dataset for the 2D Justo [6] architecture, using training from scratch vs. CMTSSL. Each color represents a different class, with black being the background. Best viewed in color.

evaluation of training procedures. Note that MIM, JPS and MTSSL can be seen as ablated versions of our CMTSSL. The ablated versions are not always able to surpass the training from scratch. In most cases, CMTSSL outperforms the single-task pretraining frameworks across all evaluation metrics. Moreover, CMTSSL is the only self-supervised pretraining framework that is consistently better than training from scratch. The examples shown in Figure 3 further support our observation. Overall, the results confirm our hypothesis, namely that multi-task self-supervised learning is beneficial, but only with the help of curriculum learning.

3) *More ablation studies:* To assess the influence of individual hyperparameters, we conduct a sensitivity analysis by varying CMTSSL-specific parameters (α_{spa} , α_{spe} , α_{mim} , K , F , S), one at a time. The results shown in Figure 4 indicate that CMTSSL is a robust framework, consistently outperforming the baseline model (trained from scratch) across various hyperparameter configurations, demonstrating its robustness and stability.

V. CONCLUSION

In this paper, we introduced CMTSSL, a curriculum multi-task self-supervised learning framework tailored for HSI analysis with lightweight architectures. By jointly integrating MIM and decoupled JPS via curriculum learning based on 3D gradient magnitudes, CMTSSL enables efficient and robust representation learning, without increasing model complexity or requiring additional labeled data. Extensive experiments on four benchmark datasets demonstrate that CMTSSL consistently improves downstream segmentation performance across various encoder architectures. Our ablation results confirm the utility of the integrated components. By enabling lightweight models to acquire semantically rich and generalizable features through a unified and efficient pre-training pipeline, CMTSSL provides a promising foundation for next-generation remote sensing systems requiring fast, accurate, and compact hyperspectral processing.

REFERENCES

- [1] S. Li, W. Song, L. Fang *et al.*, “Deep learning for hyperspectral image classification: An overview,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 57, no. 9, pp. 6690–6709, 2019.
- [2] J. Shahmoradi, E. Talebi, P. Roghanchi *et al.*, “A comprehensive review of applications of drone technology in the mining industry,” *Drones*, vol. 4, no. 3, p. 34, 2020.
- [3] A. Wendel and J. Underwood, “Self-supervised weed detection in vegetable crops using ground based hyperspectral imaging,” in *ICRA*, 2016, pp. 5128–5135.
- [4] J. Yang, C. Wu, B. Du *et al.*, “Enhanced multiscale feature fusion network for HSI classification,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 59, no. 12, pp. 10328–10347, 2021.
- [5] W. Zhang, Z. Li, and T. Zhen, “Hybrid-KANet: a hyperspectral remote sensing crop classification method based on the Kolmogorov-Arnold network,” *Geo-spatial Information Science*, pp. 1–23, 2025.
- [6] J. A. Justo, A. Ghiță, D. Kováč *et al.*, “Semantic segmentation in satellite hyperspectral imagery by deep learning,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 273–293, 2025.
- [7] C. Li, B. Rasti, X. Tang *et al.*, “Channel-layer-oriented lightweight spectral-spatial network for hyperspectral image classification,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 62, pp. 1–14, 2024.




Fig. 4: Parameter sensitivity analysis on the Pavia Center dataset using the CLOLN architecture. Base values are $\alpha_{\text{spa}} = 1$, $\alpha_{\text{spe}} = 1$, $\alpha_{\text{min}} = 4$, $K = 32$, $F = 1.5$, and $S = 3$. Each plot varies a single hyperparameter, while the others remain fixed. The baseline corresponds to CLOLN without CMTSSL. Best viewed in color.

- [8] D. Wang, M. Hu, Y. Jin *et al.*, “HyperSIGMA: Hyperspectral Intelligence Comprehension Foundation Model,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 47, no. 8, pp. 6427–6444, 2025.
- [9] P. Gamba. Hyperspectral remote sensing scenes - grupo de inteligencia computacional (GIC). [Online]. Available: https://www.ehu.es/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes
- [10] Y. Wang, C. M. Albrecht, N. A. A. Braham *et al.*, “Self-supervised learning in remote sensing: A review,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 10, no. 4, pp. 213–247, 2022.
- [11] L. Ericsson, H. Gouk, and T. M. Hospedales, “How well do self-supervised models transfer?” in *CVPR*, 2021, pp. 5414–5423.
- [12] V. Hondru, F. A. Croitoru, S. Minaee *et al.*, “Masked image modeling: A survey,” *International Journal of Computer Vision*, vol. 133, pp. 7154–7200, 2025.
- [13] M. Zhu, J. Fan, Q. Yang *et al.*, “SC-EADNet: A Self-Supervised Contrastive Efficient Asymmetric Dilated Network for Hyperspectral Image Classification,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 60, pp. 1–17, 2022.
- [14] X. Huang, M. Dong, J. Li *et al.*, “A 3-D-Swin Transformer-based hierarchical contrastive learning method for hyperspectral image classification,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 60, pp. 1–15, 2022.
- [15] L. Scheibenreif, M. Mommert, and D. Borth, “Masked vision transformers for hyperspectral image classification,” in *CVPR*, 2023, pp. 2166–2176.
- [16] Y. Shu, A. Van den Hengel, and L. Liu, “Learning common rationale to improve self-supervised representation for fine-grained visual recognition problems,” in *CVPR*, 2023, pp. 11 392–11 401.
- [17] X. Yang, W. Cao, Y. Lu *et al.*, “Self-supervised learning with prediction of image scale and spectral order for hyperspectral image classification,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 60, pp. 1–15, 2022.
- [18] G. Wang, Y. Wang, J. Qin, D. Zhang *et al.*, “Video anomaly detection by solving decoupled spatio-temporal jigsaw puzzles,” in *ECCV*, 2022, pp. 494–511.
- [19] Y. Bengio, J. Louradour, R. Collobert *et al.*, “Curriculum learning,” in *ICML*, 2009, pp. 41–48.
- [20] X. Hu, Y. Zhong, C. Luo *et al.*, “WHU-Hi: UAV-borne hyperspectral with high spatial resolution (H2) benchmark datasets for hyperspectral image classification,” *arXiv:2012.13920*, 2020.
- [21] J. A. Justo, J. Garrett, D. D. Langer *et al.*, “An open hyperspectral dataset with sea-land-cloud ground-truth from the hypo-1 satellite,” in *WHISPERS*, 2023, pp. 1–5.
- [22] D. Wei, J. Lim, A. Zisserman *et al.*, “Learning and using the arrow of time,” in *CVPR*, 2018, pp. 8052–8060.
- [23] M. Noroozi and P. Favaro, “Unsupervised learning of visual representations by solving jigsaw puzzles,” in *ECCV*, 2016, pp. 69–84.
- [24] P. Chen, S. Liu, and J. Jia, “Jigsaw clustering for unsupervised visual representation learning,” in *CVPR*, 2021, pp. 11 526–11 535.
- [25] K. He, X. Chen, S. Xie *et al.*, “Masked autoencoders are scalable vision learners,” in *CVPR*, 2022, pp. 16 000–16 009.
- [26] Z. Xie, Z. Zhang, Y. Cao, Y. Lin *et al.*, “SimMIM: A Simple Framework for Masked Image Modeling,” in *CVPR*, 2022, pp. 9653–9663.
- [27] T. Chen, S. Kornblith, M. Norouzi *et al.*, “A simple framework for contrastive learning of visual representations,” in *ICML*, 2020, pp. 1597–1607.
- [28] J.-B. Grill, F. Strub, F. Alché *et al.*, “Bootstrap your own latent a new approach to self-supervised learning,” in *NeurIPS*, 2020, pp. 21 271–21 284.
- [29] K. He, H. Fan, Y. Wu *et al.*, “Momentum Contrast for Unsupervised Visual Representation Learning,” in *CVPR*, 2020, pp. 9726–9735.
- [30] L. Song, Z. Feng, S. Yang *et al.*, “Self-supervised assisted semi-supervised residual network for hyperspectral image classification,” *Remote Sensing*, vol. 14, no. 13, p. 2997, 2022.
- [31] D. L. Ayuba, J.-Y. Guillemaut, B. Marti-Cardona *et al.*, “SpecBPP: A Self-Supervised Learning Approach for Hyperspectral Representation and Soil Organic Carbon Estimation,” *arXiv:2507.19781*, 2025.
- [32] J. Jakubik, F. Yang, B. Blumenstiel *et al.*, “TerraMind: Large-Scale Generative Multimodality for Earth Observation,” in *ICCV*, 2025, pp. 7383–7394.
- [33] N. A. A. Braham, C. M. Albrecht, J. Mairal *et al.*, “SpectralEarth: Training hyperspectral foundation models at scale,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 16 780–16 797, 2025.
- [34] D. Muhtar, X. Zhang, P. Xiao *et al.*, “CMID: A unified self-supervised learning framework for remote sensing image understanding,” *IEEE Trans. on Geoscience and Remote Sensing*, vol. 61, pp. 1–17, 2023.
- [35] G. Zhang, P. Tan, X. Fang *et al.*, “A unified self-supervised learning framework for hyperspectral image classification,” *IEEE Access*, vol. 13, pp. 49 874–49 890, 2025.
- [36] W. Chen, Y. Zhang, Z. Xiao *et al.*, “Spectral-spatial self-supervised learning for few-shot hyperspectral image classification,” *arXiv:2505.12482*, 2025.
- [37] B. Liu, X. Liu, X. Jin *et al.*, “Conflict-averse gradient descent for multi-task learning,” in *NeurIPS*, vol. 34, 2021, pp. 18 878–18 890.
- [38] X. Chen and A. Gupta, “Webly supervised learning of convolutional networks,” in *ICCV*, 2015, pp. 1431–1439.
- [39] P. Soviany, R. T. Ionescu, P. Rota *et al.*, “Curriculum self-paced learning for cross-domain object detection,” *Computer Vision and Image Understanding*, vol. 204, p. 103166, 2021.
- [40] A. Caubrière, N. A. Tomashenko, A. Laurent *et al.*, “Curriculum-based transfer learning for an effective end-to-end spoken language understanding and domain portability,” in *INTERSPEECH*, 2019, pp. 1198–1202.
- [41] T. Karras, T. Aila, S. Laine *et al.*, “Progressive growing of GANs for improved quality, stability, and variation,” in *ICLR*, 2018.
- [42] F. A. Croitoru, N. C. Ristea, R. T. Ionescu *et al.*, “Learning rate curriculum,” *International Journal of Computer Vision*, vol. 133, no. 1, pp. 291–314, 2025.
- [43] S. Sinha, A. Garg, and H. Larochelle, “Curriculum by smoothing,” in *NeurIPS*, 2020, pp. 21 653–21 664.
- [44] P. Soviany, R. T. Ionescu, P. Rota *et al.*, “Curriculum learning: A survey,” *International Journal of Computer Vision*, vol. 130, no. 6, pp. 1526–1565, 2022.
- [45] N. Madan, N.-C. Ristea, K. Nasrollahi *et al.*, “CL-MAE: Curriculum-Learned Masked Autoencoders,” in *WACV*, 2024, pp. 2492–2502.
- [46] A. Seth, R. Selvakumar, S. Sakshi *et al.*, “EH-MAM: Easy-to-Hard Masked Acoustic Modeling for Self-Supervised Speech Representation Learning,” in *EMNLP*, 2024, pp. 6386–6400.
- [47] J. A. Justo, D. D. Langer, S. Berg *et al.*, “Hyperspectral Image Segmentation for Optimal Satellite Operations: In-Orbit Deployment of 1D-CNN,” *Remote Sensing*, vol. 17, no. 4, p. 642, 2025.
- [48] P. K. A. Vasu, J. Gabriel, J. Zhu *et al.*, “FastViT: A Fast Hybrid Vision Transformer Using Structural Reparameterization,” in *ICCV*, 2023, pp. 5785–5795.
- [49] B. Jähne, H. Scharr, S. Körkel *et al.*, “Principles of filter design,” *Handbook of Computer Vision and Applications*, vol. 2, pp. 125–151, 1999.
- [50] J. Kirkpatrick, R. Pascanu, N. Rabinowitz *et al.*, “Overcoming catastrophic forgetting in neural networks,” *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.