

Towards Privacy-preserved Pre-training of Remote Sensing Foundation Models with Federated Mutual-guidance Learning

Jieyi Tan^{1*} Chengwei Zhang^{2*} Bo Dang¹ Yansheng Li^{1†}

¹Wuhan University ²University of Cambridge

tanjieyi@whu.edu.cn cz431@cam.ac.uk yansheng.li@whu.edu.cn

Abstract

Traditional Remote Sensing Foundation Models (RSFMs) are pre-trained with a data-centralized paradigm, through self-supervision on large-scale curated remote sensing data. For each institution, however, pre-training RSFMs with limited data in a standalone manner may lead to suboptimal performance, while aggregating remote sensing data from multiple institutions for centralized pre-training raises privacy concerns. Seeking for collaboration is a promising solution to resolve this dilemma, where multiple institutions can collaboratively train RSFMs without sharing private data. In this paper, we propose a novel privacy-preserved pre-training framework (**FedSense**), which enables multiple institutions to collaboratively train RSFMs without sharing private data. However, it is a non-trivial task hindered by a vicious cycle, which results from model drift by remote sensing data heterogeneity and high communication overhead. To break this vicious cycle, we introduce federated mutual-guidance learning. Specifically, we propose a Server-to-Clients Guidance (SCG) mechanism to guide clients' updates towards global-flatness optimal solutions. Additionally, we propose a Clients-to-Server Guidance (CSG) mechanism to inject local knowledge into the server by low-bit communication. Extensive experiments on four downstream tasks demonstrate the effectiveness of our FedSense in both full-precision and communication-reduced scenarios, showcasing remarkable communication efficiency and performance gains.

1. Introduction

Remote Sensing Foundation Models (RSFMs) have recently attracted increasing attention for their impressive applicability and performance across various Earth Observation tasks by producing general-purpose visual features [47, 66]. Traditionally, existing RSFMs follow a

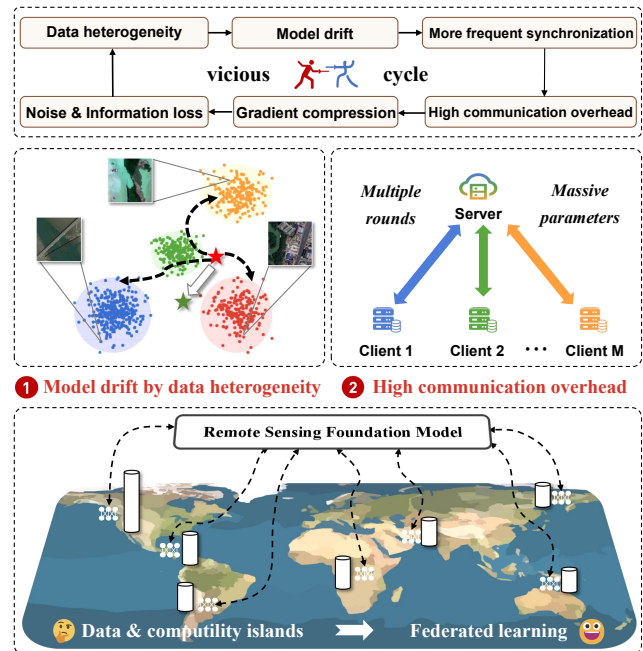


Figure 1. **Privacy-preserved pre-training of RSFMs with FL to bridge data islands.** The vicious cycle between data heterogeneity-induced model drift and communication bottlenecks reveals a critical trade-off between performance and efficiency.

data-centralized training paradigm. They are built through self-supervision on large-scale curated remote sensing data, gathered from diverse sources [15]. These data are collected via satellites, drones, or aerial platforms and are stored in centralized archives by different institutions. However, a single institution often lacks the data scale and diversity required to pre-train RSFMs [28]. Meanwhile, aggregating multi-institutional data for centralized pre-training poses challenges due to concerns over geo-information security, storage limitations, and competitive interests [4, 19, 43, 49]. These contradictory demands urgently demand paradigm-shifting frameworks that enable collaborative training without compromising privacy [32, 67].

A more practical and realistic solution is to collabora-

*Equal contribution.

†Corresponding author.

tively inject remote sensing knowledge learned from private data owned by institutions into foundation models in a distributed manner [10, 18, 25, 44, 45]. In recent years, federated learning (FL) has emerged as a promising privacy-preserving alternative, enabling collaborative model training without raw data exchange via periodic model aggregation [13, 17, 34]. Self-supervised learning (SSL) operates on the principle of latent structure exploitation, where models learn transferable representations by solving pretext tasks derived from intrinsic attributes in data without human annotations [63]. However, directly combining SSL and FL to pre-train RSFMs is non-trivial [14]. Existing studies on federated self-supervised learning (FSSL) mainly focus on natural images, showing potential for distributed visual representation learning. For example, MocoSFL [23] focuses on cross-device SSL by leveraging momentum contrast on mobile devices. FedU² designs disperses local data representations uniformly using spherical Gaussian sampling and optimizes global-local model consistency. FedMKD [24] introduces a resource-adaptive FSSL to handle architecture heterogeneity and class skew. A recent commentary [28] envisions collaborative pre-training of RSFMs via FSSL to bridge data and computility islands. However, few works have explored the **vicious cycle** challenge in FSSL, which manifests more severely in remote sensing [3, 20, 53, 54].

As illustrated in Fig. 1, this vicious cycle comprises two tightly coupled challenges: ❶ **model drift by data heterogeneity**. Remote sensing data is inherently heterogeneous due to diverse sensor types, resolutions, and geographic distribution. Such heterogeneity leads to significant variability in data distributions across different institutions, causing convergence inefficiencies and degraded model performance [4, 27]. ❷ **high communication overhead**. Foundation models are characterized by massive parameters (unlike neural networks with shallow architectures), leading to excessive communication costs and bandwidth demands [33]. The two challenges form a vicious cycle, where data heterogeneity necessitates more frequent model synchronization to mitigate model drift, thereby amplifying communication costs [57, 58]. Conversely, communication compression may reduce overhead, but it introduces noise and information loss, which further exacerbates client-side model drift. This bidirectional aggravation fundamentally undermines the efficiency and model consistency in distributed pre-training of RSFMs via FSSL.

In this paper, we introduce a new challenging yet meaningful task: **privacy-preserved pre-training of RSFMs**. We propose FedSense, a novel FSSL framework with federated mutual-guidance learning, to address the vicious cycle challenges. First, we introduce a Server-to-Clients Guidance (SCG) mechanism to guide client models towards a global consensus, mitigating model drift by data heterogeneity. Second, we propose a Clients-to-Server Guidance

(CSG) mechanism to distill client models’ knowledge into a server-side reference model, reducing communication overhead. FedSense thus provides an integrated solution to resolve the vicious cycle challenges. As a result, our FedSense achieves state-of-the-art performance compared to existing FSSL methods with higher performance and lower communication overhead.

To sum up, this paper takes the first step towards privacy-preserved pre-training of RSFMs as far as we know. The main contributions of this work are as follows:

- We propose FedSense, establishing a new paradigm for privacy-preserved pre-training of RSFMs. To the best of our knowledge, it is the first generic FL framework that supports mainstream pre-training methods, including contrastive learning and masked image modeling.
- The vicious cycle problem in FSSL is alleviated by federated mutual-guidance learning, which is composed of the Server-to-Clients Guidance (SCG) and Clients-to-Server Guidance (CSG) mechanisms. The integrated solution effectively mitigates model drift due to data heterogeneity and reduces communication overhead.
- We pre-train an RSFM with million-scale remote sensing data from ten participants and evaluate the performance on four downstream tasks. Experimental results on eight datasets demonstrate that our FedSense achieves superior performance and lower communication overhead compared to existing FSSL methods.

2. Related Work

2.1. Centralized Pre-training for RSFMs

Recent years have witnessed significant progress in data-centralized pre-training paradigms for RSFMs [26, 35, 37]. A series of studies focuses on constructing large-scale pre-training datasets and developing specialized algorithms tailored to remote sensing data characteristics. For instance, [51] introduces a rotation-variable window attention mechanism to handle large-scale and arbitrarily oriented geospatial objects, accompanied by MillionAID, a billion-parameter vision foundation model pre-trained on massive remote sensing imagery. To address dense object detection challenges in remote sensing, RingMo [46] optimizes a masked autoencoder framework by incorporating multi-scale hierarchical representations and task-specific decoding strategies. For multispectral data with rich spectral information, SpectralGPT [16] treats multispectral images as 3D tensors and proposes a multi-objective reconstruction loss to jointly capture spatial-spectral correlations and spectral sequence dependencies. SkySense [15] further extends multimodal capabilities through spatiotemporal disentanglement and temporal-aware embedding mechanisms, enabling joint contrastive learning across high-resolution optical imagery and time-series optical and Synthetic Aperture

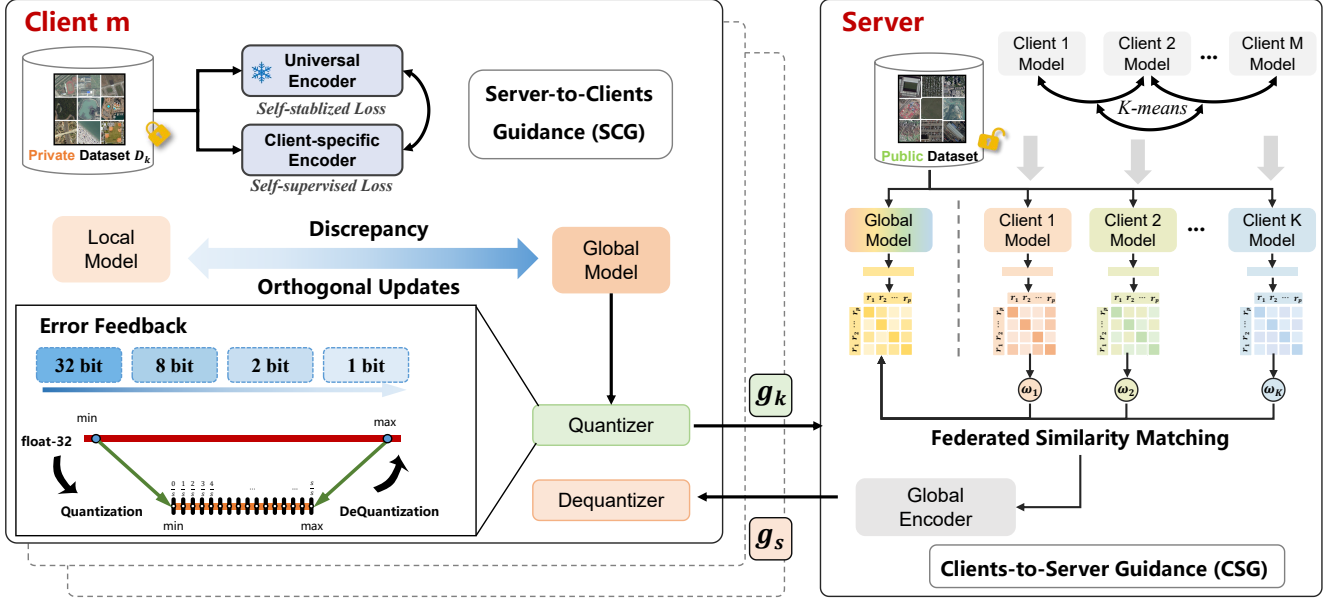


Figure 2. **Overview of our FedSense.** The framework includes two components: the Server-to-Clients Guidance (SCG) mechanism and Clients-to-Server Guidance (CSG) mechanism. SCG guides clients’ updates towards solutions with optimal global flatness, while CSG injects local knowledge into the server via low-bit communication.

Radar (SAR) data. This unified framework supports diverse tasks ranging from image-level classification to pixel-level segmentation and crop phenology monitoring. AnySat [1] is a versatile model designed to handle diverse data across resolutions, scales, and modalities.

Different from the above studies, our work is the first to explore privacy-preserved pre-training of RSFMs, and we aim to collaboratively pre-train RSFMs in a novel decentralized manner. This research is orthogonal to the existing centralized pre-training methods tailored to RSFMs, and can be complementary to scale up the performance of RSFMs in real-world applications.

2.2. Federated Self-Supervised Learning (FSSL)

FL [7, 22, 39] has emerged as a promising paradigm for collaborative model training across decentralized data sources while preserving data privacy. Moreover, FSSL [40, 61, 68] has demonstrated potential for privacy-preserving model training. However, existing FSSL methods predominantly address cross-device scenarios characterized by numerous resource-constrained clients (e.g., mobile devices) with homogeneous data distributions, focusing on computational efficiency and communication compression. For instance, L-DAWA [41] mitigates data heterogeneity through hierarchical angular divergence weighting, while FedU² [31] and FedMKD [24] enhance representation consistency via unified embedding alignment and adaptive knowledge distillation. Despite these advances, current methods remain inadequate for cross-institution pre-training of RSFMs.

Existing FSSL methods focus on optimizing federated model convergence and mainly support cross-device scenarios [21, 36, 50]. Additionally, these methods have limited applicability to support self-supervised learning frameworks (contrastive learning and masked image modeling) consistently. Thirdly, these methods ignore the vicious cycle in collaboratively pre-training of RSFMs, which requires handling the drift challenges and communication efficiency simultaneously. Our work sheds light on these limitations and proposes federated mutual-guidance learning to unleash the power of FSSL for RSFM pre-training.

3. Methodology

3.1. Problem Definition

An FL framework for privacy-preserving collaborative pre-training of RSFMs aims to learn general-purpose visual representations from distributed and unlabeled remote sensing images. We assume that this framework comprises M institutions (*i.e.*, clients) and a central server. Let the m -th **client** c_m possess a private unlabeled dataset \mathcal{D}_{c_m} , and let the **server** s maintain a public unlabeled dataset \mathcal{D}_s . For each client c_m with model parameters θ_m , its local self-supervised objective function \mathcal{F}_{c_m} can be formulated as:

$$\min_{\theta_m} \mathcal{F}_{c_m} = \mathbb{E}_{x \sim \mathcal{D}_{c_m}} \left[\psi(f_{\theta_m}(x^{(a)}), f_{\theta_m}(x^{(b)})) \right], \quad (1)$$

where $x^{(a)}$ and $x^{(b)}$ are view augmentations or masked modeling of the input image x , and $\psi(\cdot, \cdot)$ denotes the self-

supervised loss function. For contrastive learning, ψ measures feature consistency between augmented views. For masked modeling, ψ computes the reconstruction error between predicted and original pixel values in masked regions. In the t -th communication round, upon receiving the local parameters $\{\theta_m^{(t)}\}_{m=1}^M$ from clients, the server s with public data \mathcal{D}_s applies aggregation $\Theta^{(t)} = \sum_{m=1}^M p_m \theta_m^{(t)}$, where $p_m = \frac{|\mathcal{D}_{c_m}|}{\sum_{m=1}^M |\mathcal{D}_{c_m}|}$. The server's optimization objective is:

$$\min_{\Theta} \mathcal{F}_s = \mathbb{E}_{x_s \sim \mathcal{D}_s} \left[\phi(\Theta, \{\theta_m^{(t)}\}_{m=1}^M, \dots, x_s) \right], \quad (2)$$

where x_s is an image sampled from the public dataset, and ϕ denotes the loss function to further inject public data knowledge into the collaboratively pre-trained foundation model.

3.2. Overview of Our FedSense

As illustrated in Fig. 2, the core of FedSense is a bi-directional guidance mechanism, comprising two primary components: the SCG and CSG mechanisms. These components are detailed in the subsequent sections. The complete algorithmic procedure is presented in Algorithm 1. For conciseness, we omit the explicit conversion between gradients and parameters during the transmission phase.

3.3. Server-to-Clients Guidance

The core insight of SCG is to strike a balance between global knowledge preservation and local model optimization. The server guides clients to optimize their local models while constraining their deviation from the global model. To achieve this, we design a dual loss for each client, consisting of a self-supervised loss and a self-stabilized loss in Eq. (7), which are complementary. The former requires models to learn orthogonality to the discrepancy between the local and global models, while the latter regularizes the model with stabilized knowledge information in federated updates. Thus, the overall training loss for the m -th client is:

$$\mathcal{L}_m^{\text{total}} = \underbrace{\mathcal{L}_m^{\text{ssl}}(\theta_m; \Theta; \epsilon; \mathcal{R}_C(\theta_m))}_{\text{self-supervised term}} + \underbrace{\mathcal{L}_m^{\text{sst}}(\theta_m; \theta_{\text{uni}})}_{\text{self-stabilized term}}, \quad (3)$$

where $\mathcal{L}_m^{\text{ssl}}$ denotes the self-supervised loss. It corresponds to SSL objectives commonly used in the field, such as contrastive loss and masked reconstruction loss. Note that this is a drift-aware loss incorporated with our proposed optimization method $\mathcal{R}_C(\theta_m)$. We conclude it as the following min-max optimization problem:

$$\min_{\theta_m} \max_{|\epsilon|_2 < \rho} \mathcal{L}_m^{\text{disc}}(\theta_m + \epsilon, f_{\theta_m}(x^{(a)}), f_{\theta_m}(x^{(b)})) \quad (4)$$

The objective of the weight perturbation on the client-side is to find ϵ , which causes the maximum increase in the parameter discrepancy. Given local model θ_m and global model

Θ , the parameter discrepancy is calculated by:

$$\nabla_{\theta_m} \mathcal{L}_m^{\text{disc}}(\theta_m; \Theta) = \nabla_{\theta_m} (\beta(\theta_m - \Theta)), \quad (5)$$

where β is a hyperparameter of the discrepancy term. Then, the optimal perturbation ϵ is approximated by:

$$\tilde{\epsilon} = \operatorname{argmax}_{\epsilon} \mathcal{L}_m^{\text{disc}}(\theta_m + \epsilon; \Theta) \approx \lambda \frac{\nabla_{\theta_m} \mathcal{L}_m^{\text{disc}}(\theta_m; \Theta)}{\|\nabla_{\theta_m} \mathcal{L}_m^{\text{disc}}(\theta_m; \Theta)\|_2}, \quad (6)$$

where λ is a scaling factor. The optimal perturbation ϵ will be later used to update the local model θ_m .

Moreover, the self-stabilized loss $\mathcal{L}_m^{\text{sst}}$ is designed to leverage the knowledge from a universal encoder, $f_{\theta_{\text{uni}}}$, which is pre-trained on a large-scale dataset (e.g., ImageNet-22k) [12] and broadcast to all clients by the server at the beginning of the FL process. It guides the client-specific encoder to output representations that align with the universal encoder. Formally, this loss can be defined as:

$$\mathcal{L}_m^{\text{sst}} = \mathbb{E}_{x \sim \mathcal{D}_{c_m}} \left[-\frac{f_{\theta_m}(x)}{\|f_{\theta_m}(x)\|_2} \cdot \frac{f_{\theta_{\text{uni}}}(x)}{\|f_{\theta_{\text{uni}}}(x)\|_2} \right], \quad (7)$$

where $f_{\theta_{\text{uni}}}$ and f_{θ_m} are the feature extractors of the universal encoder and the client-specific encoder, respectively.

Lastly, the optimal perturbation ϵ is then used to update the local model θ_m :

$$\theta_m \leftarrow \theta_m - \gamma \nabla_{\theta_m} \mathcal{L}_m^{\text{total}}(\theta_m + \epsilon; \Theta), \quad (8)$$

where γ is the learning rate. In this way, the server guides the clients to optimize their local models while preserving the global knowledge with orthogonal property and stabilized information.

3.4. Clients-to-Server Guidance

As the model size and number of clients increase, the communication overhead becomes a more serious bottleneck. The clients need to transmit the updates to the server in each communication round and wait for the server to aggregate them, then send back the updated parameters.

Since the uplink communication is the most time-consuming part in this process, we propose a CSG mechanism to reduce the communication cost. Quantization is a widely used technique to reduce the communication cost. However, it inherently introduces quantization error and information loss, which may accumulate over time and degrade the model's performance. To guide the server to aggregate the quantized updates from clients, we propose a feedback error mechanism to compensate for the quantization error.

Assume the feedback error at the t -th round is e_m^t , it is added to compensate updates of each client:

$$\mathcal{G}_m^{t+1} = \Delta \theta_m^t + e_m^t, \quad (9)$$

Algorithm 1: Our FedSense

Input : T : communication round; M : number of clients; E : local epochs;

Output: Weight Θ^T of the RSFM at the T -th round.

```
1 Server-side:
2 Initialize global model  $\Theta^0$ 
3 for  $t = 1$  to  $T$  do // the  $t$ -th round
4   Server broadcast  $\Theta^{t-1}$  to selected clients
5   for each client  $c_m \in \{c_m\}_{m=1}^M$  in parallel do
6     Server-to-Clients Guidance (SCG)
7   end
8    $\{\Delta\theta_m^t\}_{m=1}^M = \text{DQF}\{\{\Delta\tilde{\theta}_m^t\}_{m=1}^M; b\}$ 
9    $\Theta^t \leftarrow \text{ServerAggregation}(\{\Delta\theta_m^t\}_{m=1}^M)$ 
10  Clients-to-Server Guidance (CSG)
11  Similarity alignment with public data Eq. (18)
12 end
13 Client-side:
14 Initialize the local model  $\theta_m^{t-1} \leftarrow \Theta^{t-1}$ 
15 for  $e = 1$  to  $E$  do // the  $e$ -th epoch
16   Server-to-Clients Guidance (SCG)
17   Optimization with knowledge preservation
18    $\theta_m \leftarrow \theta_m - \gamma \nabla_{\theta_m} \mathcal{L}_m^{\text{total}}(\theta_m + \epsilon; \Theta)$  Eq. (8)
19 end
20  $\Delta\tilde{\theta}_m^t = \text{QF}\{\{\theta_m^t - \theta_m^{t-1}\}; b\}$ 
21 Send the updates  $\Delta\tilde{\theta}_m^t$  back to the server.
22 Downstream tasks: Institutions utilize the
    pre-trained model as a backbone, fine-tuning it on
    labeled data for specific tasks.
```

where $\Delta\theta_m^t$ is the updates of client m at round t . The quantized updates are computed as follows:

$$\tilde{\mathcal{G}}_m^{t+1}[i] = \underbrace{\|\mathcal{G}_m^{t+1}\|}_{\text{L2 norm of raw updates}} \cdot \underbrace{\text{sgn}(\mathcal{G}_m^{t+1}[i])}_{\text{sign of element (1 bit)}} \cdot \underbrace{\xi(\mathcal{G}_m^{t+1}[i]; s)}_{\text{unbiased stochastic function}(b-1 \text{ bits})}, \quad (10)$$

where $\xi(\cdot)$ is a unbiased stochastic function mapping $|\mathcal{G}_m^{t+1}[i]|/|\mathcal{G}_m^{t+1}|$ to the quantization level s .

$$e_m^t = \mathcal{G}_m^t - \text{DQF}(\tilde{\mathcal{G}}_m^t; b), \quad (11)$$

where DQF is the dequantization function. The feedback error is updated by:

$$e_m^t = \alpha \cdot e_m^{t-1} + (1 - \alpha) \cdot e_m^t, \quad (12)$$

where α is the momentum factor, and b is the bit-width of quantization. The feedback error is accumulated over time to preserve the gradient information across communication rounds. Inspired by the dynamic optimization characteristics of neural network training [59], we implement periodic feedback error resetting to address the fast-evolving loss

landscape. We reset the feedback error to zero at the frequency of T_{reset} rounds to ensure the feedback error remains aligned with the current optimization state.

On the server side, we propose federated similarity distillation to provide public remote sensing data guidance. Our FedSense consists of three core steps: server-side aggregation, local model clustering, and cross-model knowledge distillation. In this way, clients can leverage the public data to enhance the global model's performance. The server aggregates client models $\{\theta_m\}_{m=1}^M$ using data volume weights. They are clustered into K groups via K-means to accelerate the multi-model forward pass cost and mitigate the bias of some models:

$$\{\theta^{(k)}\}_{k=1}^K = \text{K-means}(\{\theta_m\}_{m=1}^M). \quad (13)$$

For each public batch $\mathcal{B} = \{x_i\}_{i=1}^p$, the group of models produce feature matrices:

$$\mathbf{Z}^{(k)} = [z_1^{(k)}, \dots, z_p^{(k)}]^\top \in \mathbb{R}^{p \times d}, \quad (14)$$

$$\mathbf{Z}^g = [z_1^g, \dots, z_p^g]^\top \in \mathbb{R}^{p \times d}, \quad (15)$$

where $z_i^{(k)} = f_{\theta^{(k)}}(x_i)$ and $z_i^g = f_{\Theta}(x_i)$, and p and d are the batch size and feature dimension, respectively. The similarity matrices are computed as:

$$\mathbf{S}^{(k)} = \frac{\mathbf{Z}^{(k)}(\mathbf{Z}^{(k)})^\top}{\|\mathbf{Z}^{(k)}\|_F}, \quad \mathbf{S}^g = \frac{\mathbf{Z}^g(\mathbf{Z}^g)^\top}{\|\mathbf{Z}^g\|_F}. \quad (16)$$

The weighted consensus similarity combines local expertise:

$$\mathbf{S}_{\text{consensus}} = \sum_{k=1}^K \omega_k \mathbf{S}^{(k)}, \quad \omega_k = \frac{|\mathcal{D}_k|}{\sum_{i=1}^N |\mathcal{D}_i|}. \quad (17)$$

The global model is optimized by matching similarity distributions:

$$\mathcal{L}_{\text{distill}} = \frac{1}{p^2} \|\mathbf{S}^g - \mathbf{S}_{\text{consensus}}\|_F^2. \quad (18)$$

By distilling the similarity knowledge from multiple models, the global model can learn to capture the intrinsic structure of the public data, which is beneficial for enhancing the generalization ability of the global model.

4. Experiments

4.1. Federated Experimental Setups

Distributed Pre-training Datasets. A distributed dataset is constructed for federated pre-training of RSFMs, comprising ten clients with heterogeneous private remote sensing data and supplementary public datasets maintained by a server (Fig. 3). Notably, the dataset includes satellite images (e.g., WorldView-2 and JL-1) and aerial images (e.g.,

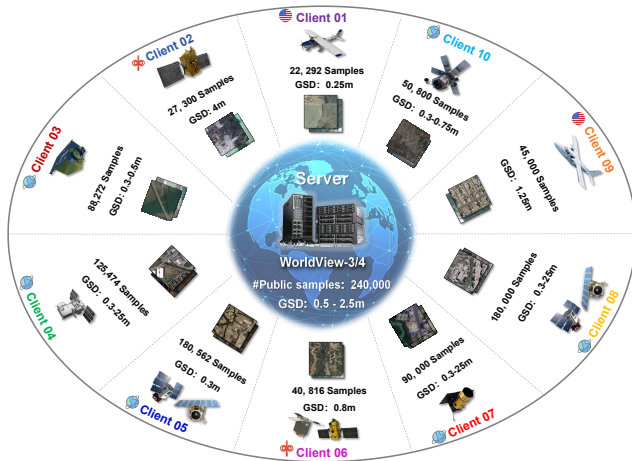


Figure 3. **Illustration of distributed pre-training datasets.** The whole datasets consist of ten clients with million-scale private remote sensing datasets and a public one maintained by a server.

NAIP and NOAA), while some institutions possess multi-sourced collections, featuring heterogeneous sensor configurations and spatial resolutions ranging from 0.25m to 25m [2, 11, 29, 30, 48]. The coverage spans diverse geographical regions, including global, the USA, and China. Such heterogeneity across data volume, geographic distribution, resolution variance, and platform diversity establishes a representative million-scale dataset that simulates real-world scenarios.

Downstream Tasks. We conduct experiments on four typical downstream tasks of remote sensing image analysis to validate the effectiveness of the RSFM pretrained by our FedSense. These tasks include: scene classification (RESISC-45 [8], AID [55]); semantic segmentation (Inria [38], LoveDA [52]); object detection (DIOR-R [9], DOTA-v1.0 [56]); and change detection (LEVIR-CD+ [6], SECOND [62]). More details of the datasets are shown in Sec. C of the supplementary.

Evaluation Metrics. For scene classification, we employ overall accuracy (OA) as the evaluation metric. We use mean intersection over union (mIoU) for multi-class semantic segmentation tasks and IoU for binary segmentation tasks. For object detection, we use mean Average Precision (mAP) and F1 score for change detection tasks. For semantic change detection, we use the semantic change segmentation score (SCS) as the evaluation metric.

Implementation Details. To systematically evaluate the universality of our framework, we establish a comprehensive evaluation framework incorporating two mainstream SSL paradigms: contrastive learning through DINO [5], and masked image modeling via SimMIM [60]. All experiments employ the tiny version of Swin Transformer [33] (Swin-T) as the foundational backbone unless explicitly stated,

with consistent training protocols (100-round pre-training, AdamW optimizer, $1e-4$ base learning rate).

4.2. Comparison with State-of-the-Art Methods

In this section, we present a comparative analysis of performance on four downstream tasks using the RSFM pretrained by our FedSense and other state-of-the-art FSSL methods. The results are summarized in Tab. 1. We observe that our FedSense outperforms existing methods across almost all tasks, achieving an average improvement of 1.3% OA on scene classification, 1.1% mIoU on semantic segmentation, 0.9% mAP on object detection, and 0.8% F1 on change detection. Note that due to not using a public dataset in the original paper, we conduct experiments by averagely integrating public datasets into the federated pre-training dataset for a fair comparison with existing methods with the same number of total samples. To elaborate, we provide detailed comparisons with existing methods.

Our experimental results reveal several key observations. The randomly initialized model yields inferior performance across all downstream tasks, demonstrating Transformers’ inherent limitation in learning effective representations from limited labeled data. ImageNet supervised pre-training brings significant improvements (*e.g.*, +14.36% on RESISC-45), validating the importance of pre-training for vision transformers. However, the domain gap between natural images and top-down remote sensing views limits further performance gains, motivating our FL approach.

Among FSSL methods, most approaches (FedEMA, FedMKD) achieve moderate improvements over ImageNet pretraining, confirming that collaborative pretraining with distributed data can inject domain-specific knowledge. Notably, FedU² exhibits performance degradation compared to ImageNet pretraining, suggesting its vulnerability to data heterogeneity and insufficient utilization of pre-trained knowledge. Our FedSense consistently outperforms SSL-FL, achieving absolute gains of 1.12% and 1.37% on RESISC-45 and AID for scene classification. When using the DINO framework, our FedSense maintains advantages with 0.87% and 0.26% improvements, respectively.

For semantic segmentation, FedSense achieves 0.85% OA improvement on Inria and 1.07% gain on LoveDA over the best competitor. The marginal differences under the DINO framework (0.12%-0.35%) suggest that contrastive learning brings limited benefits for segmentation tasks that require precise pixel-level localization. In rotated object detection, FedSense with the SimMIM framework surpasses SOTA by 1.50% and 0.51% on DIOR-R and DOTA-v1.0, while maintaining competitive performance (+0.11% mAP) with DINO-based approaches.

Additionally, the 1-bit quantized version of FedSense demonstrates remarkable communication efficiency with minimal performance degradation (0.23% across tasks).

Framework	Method	RESISC-45	AID	Inira		LoveDA	DIOR-R	DOTA-v1.0	LEVIR-CD+			SECOND
		OA	OA	IoU	OA	mIoU	mAP	mAP	Precision	Recall	F1	SCS
Swin-Tiny	Random Init.	80.09	72.84	79.41	90.21	42.83	46.86	64.45	69.71	63.39	66.40	30.95
	ImageNet Sup.	94.45	96.36	80.65	93.16	51.52	64.37	76.90	73.26	71.44	72.34	34.79
SimMIM [60] (Swin-Tiny)	SSL-FL [61]	95.21	96.17	<u>81.33</u>	93.43	51.67	64.01	<u>76.82</u>	<u>74.31</u>	70.99	72.61	33.98
	FedEMA [68]	-	-	-	-	-	-	-	-	-	-	-
	L-DAWA [41]	-	-	-	-	-	-	-	-	-	-	-
	FedU ² [31]	-	-	-	-	-	-	-	-	-	-	-
	FedMKD [24]	-	-	-	-	-	-	-	-	-	-	-
	Our FedSense	96.33	97.54	81.66	94.28	52.74	65.51	77.33	74.82	71.67	73.21	35.43
	<u>96.01</u>	<u>96.68</u>	<u>80.97</u>	<u>93.45</u>	<u>51.91</u>	<u>64.53</u>	<u>76.76</u>	<u>73.98</u>	<u>71.59</u>	<u>72.77</u>	<u>34.87</u>	
DINO [5] (Swin-Tiny)	SSL-FL [61]	-	-	-	-	-	-	-	-	-	-	-
	FedEMA [68]	94.91	97.04	80.21	92.98	51.89	64.93	77.09	73.63	71.20	72.39	34.82
	L-DAWA [41]	94.23	96.27	80.82	93.33	51.68	65.08	77.24	<u>74.18</u>	71.82	72.98	34.97
	FedU ² [31]	93.85	95.76	80.13	92.59	51.34	64.03	75.88	73.30	70.67	71.96	33.27
	FedMKD [24]	95.34	97.17	<u>80.84</u>	<u>94.14</u>	51.99	65.44	<u>77.56</u>	74.36	73.05	<u>73.70</u>	35.11
	Our FedSense	96.21	97.43	81.96	94.23	<u>51.95</u>	<u>65.40</u>	77.64	74.98	73.13	74.04	<u>35.34</u>
	<u>95.88</u>	<u>97.23</u>	80.97	93.88	51.14	64.62	77.45	73.88	<u>73.09</u>	73.48	35.82	

Table 1. **Comparison results of our FedSense and previous SOTA methods.** The symbol (-) indicates unavailable results where methods are incompatible with specific SSL framework types. Results highlighted in **yellow** denote full-precision communication performance of our FedSense, while **blue** shading represents experiments with 1-bit communication-quantized transmission, demonstrating our method’s efficiency-accuracy trade-off. The best results are highlighted in **bold**, and the second-best results are underlined.

Quantized FedSense slightly outperforms full-precision baselines on the SECOND dataset (+0.06% SCS), which we conjecture stems from the error-compensated quantization acting as implicit regularization. This makes FedSense particularly suitable for bandwidth-constrained applications.

These results collectively demonstrate that our federated mutual-guidance learning framework effectively coordinates distributed clients to learn transferable representations while maintaining communication efficiency. The consistent improvements across diverse tasks validate FedSense’s ability to capture domain-specific patterns from unlabeled remote sensing data through federated collaboration.

4.3. Ablation Studies

In this part, we conduct ablation studies to analyze the system scalability, model scalability, effectiveness of components, and parameter analysis of our proposed FedSense.

System Scalability Analysis. To assess our framework’s adaptability to real-world distributed scenarios, we systematically investigate how model performance scales with increasing participants and training samples. As shown in Tab. 2, expanding from 2 to 10 collaborative clients (80K to 800K samples) yields consistent performance gains across all tasks. For RESISC-45 and DIOR-R, the performance improvements scale nearly linearly with client/sample quantities, suggesting these tasks particularly benefit from diverse perspectives in FedSense. However, LoveDA exhibits marginal gains despite 10× sample growth, indicating pixel-level tasks require more sophisticated feature aggregation beyond simple data scaling. These findings con-

firm our framework’s effectiveness in harnessing distributed resources.

#CL	#TS	RESISC-45	LoveDA	DIOR-R	LEVIR-CD+
		OA	mIoU	mAP	F1
2	80K	94.50	51.56	64.43	72.44
4	200K	95.23	51.86	64.75	72.68
8	650K	95.85	52.33	65.16	73.00
10	800K	96.33	52.74	65.51	73.21

Table 2. **Number of participants impact analysis.** #CL means the number of participants. #TS means the number of total participating samples.

Model Scalability Analysis. The impact of model capacity is systematically evaluated through progressively larger Swin Transformer variants, as detailed in Tab. 3. While all tasks benefit from increased model parameters, we observe distinct scaling patterns across task types. RESISC-45 shows diminishing returns, improving only by +1.21% from Swin-T to Swin-L, suggesting vision transformers approach saturation points for scene classification. Conversely, segmentation (LoveDA mIoU +3.22%) and detection (DIOR-R mAP +3.31%) exhibit near-linear improvements with model growth, indicating complex localization tasks inherently demand higher-capacity architectures. Notably, change detection (LEVIR-CD+ F1 +2.92%) demonstrates sustained sensitivity to model size, likely requiring deeper feature hierarchies to discern subtle temporal

changes. The 197M-parameter Swin-L achieves marginal gains (+0.42% mAP over Swin-B) compared to its 2.2× parameter increase, highlighting practical trade-offs between model capacity and computational costs.

Model	#Para.	RESISC-45	LoveDA	DIOR-R	LEVIR-CD+
		OA	mIoU	mAP	F1
Swin-T	28M	96.33	52.74	65.51	73.21
Swin-S	50M	96.80	54.12	67.03	74.65
Swin-B	88M	97.15	55.29	68.40	75.83
Swin-L	197M	97.54	55.96	68.82	76.13

Table 3. **Model size impact analysis.** # Para. means the number of model parameters.

Quantization Methods Comparison. We compare our proposed quantization method with uniform quantization and FedPAQ [42] on the RESISC-45 dataset. As shown in Tab. 4, our method outperforms uniform quantization and FedPAQ across all quantization bit-widths. Specifically, our method achieves 96.27% OA with 8-bit quantization, 96.13% OA with 2-bit quantization, and 96.01% OA with 1-bit quantization. Our FedSense outperforms FedPAQ by 0.14% OA with 8-bit quantization, 0.16% OA with 2-bit quantization, and 0.22% OA with 1-bit quantization. The results demonstrate the effectiveness of our proposed quantization method in enhancing the communication efficiency.

Method	32-bit	8-bit	2-bit	1-bit
Uniform Quant	96.33	96.01	95.87	95.56
FedPAQ [42]	96.33	96.13	95.97	95.79
Ours	96.33	96.27	96.13	96.01

Table 4. **Quantization bit-width comparison.**

SST	SCG	CSG	RESISC-45	LoveDA	DIOR-R	LEVIR-CD+
✓			94.52	51.87	64.83	72.43
✓	✓		94.70	51.85	64.91	72.55
✓		✓	95.25	51.97	65.12	72.78
✓	✓	✓	96.33	52.74	65.51	73.21

Table 5. **Effectiveness of proposed components.**

Effectiveness of Components. Our ablation study quantitatively validates the complementary nature of proposed components, as summarized in Tab. 5. The standalone SST mechanism achieves a 94.52% OA on RESISC-45, 51.87% mIoU on LoveDA, 64.83% mAP on DIOR-R, and 72.43% F1 on LEVIR-CD+. The SCG mechanism further boosts performance across tasks, with 0.18% OA, 0.02% mIoU, 0.08% mAP, and 0.12% F1 improvements. The CSG mechanism contributes to a 0.73% OA gain on RESISC-45, 0.29% mAP on DIOR-R, and 0.35% F1 on LEVIR-CD+.

The full model with all components achieves the best performance, demonstrating the effectiveness of our proposed components in enhancing FL for remote sensing tasks.

LE	Round	RESISC-45	LoveDA	DIOR-R	LEVIR-CD+
		OA	mIoU	mAP	F1
1	100	96.33	52.74	65.51	73.21
2	50	95.82	51.93	64.08	72.05
4	25	95.17	50.62	62.75	70.89
100	1 [†]	94.50	48.31	60.13	68.24

Table 6. **Local epochs (LE) impact analysis.** † means one-shot FL setting.

Parameter Analysis. The trade-off between local computation and communication frequency is systematically investigated through varying local epochs (LE), as shown in Tab. 6. Here, more local epochs indicate more discrepancy between local and global models, leading to potential client drift. It is worth noting that the one-shot setting (LE=100) suffers catastrophic performance collapse, confirming the inherent heterogeneity of remote sensing data demands periodic model synchronization. However, reducing communication rounds and developing one-shot FL (OSFL) is becoming widely adopted for only one round of communication. Our current design benefits from iterative refinement across rounds. Model distillation and meta-learning are adopted in OSFL and get promising results [65]. Additionally, recent studies on LLMs model merging like DARE [64] are similar to the OSFL setting, and are promising avenues for future work. We envision that our framework can be extended by integrating these techniques to help reconcile divergent models trained in a single pass. For example, model merging and distillation can be added to CSG for consolidation.

5. Conclusion & Future Work

FedSense is the first privacy-preserved pre-training framework for RSFMs, which enables multiple institutions to collaboratively pre-train RSFMs without sharing private data. We introduce federated mutual-guidance learning, which breaks the vicious cycle caused by remote sensing data heterogeneity and high communication overhead. Specifically, the SCG mechanism guides clients updates towards global-flatness optimal solutions. and the CSG mechanism injects local knowledge into the server by low-bit communication. Extensive experiments on four downstream tasks demonstrate the effectiveness of our FedSense in both full-precision and communication-reduced scenarios, showcasing remarkable communication efficiency and performance gains. In the future, we plan to extend the current framework to support the collaborative pre-training of multi-modal RSFMs with modality heterogeneity.

Acknowledgment

This work was supported by the National Natural Science Foundation of China under Grant 42371321.

References

- [1] Guillaume Astruc, Nicolas Gonthier, Clément Mallet, and Loic Landrieu. Anysat: One earth observation model for many resolutions, scales, and modalities. In *CVPR*, pages 19530–19540, 2025. 3
- [2] Favyen Bastani, Piper Wolters, Ritwik Gupta, Joe Ferdinando, and Aniruddha Kembhavi. Satlaspretrain: A large-scale dataset for remote sensing image understanding. In *ICCV*, pages 16772–16782, 2023. 6
- [3] Cosmin I Bercea, Benedikt Wiestler, Daniel Rueckert, and Shadi Albarqouni. Federated disentangled representation learning for unsupervised brain anomaly detection. *Nature Machine Intelligence*, 4(8):685–695, 2022. 2
- [4] Baris Büyüktas, Gencer Sumbul, and Begüm Demir. Federated learning across decentralized and unshared archives for remote sensing image classification: A review. *IEEE Geoscience and Remote Sensing Magazine*, pages 2–18, 2024. 1, 2
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021. 6, 7
- [6] Hao Chen and Zhenwei Shi. A spatial-temporal attention-based method and a new dataset for remote sensing image change detection. *Remote Sensing*, 12(10), 2020. 6, 3
- [7] Shengchao Chen, Guodong Long, Jing Jiang, and Chengqi Zhang. Federated foundation models on heterogeneous time series. In *AAAI*, pages 15839–15847, 2025. 3
- [8] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017. 6, 2
- [9] Gong Cheng, Jiabao Wang, Ke Li, Xingxing Xie, Chunbo Lang, Yanqing Yao, and Junwei Han. Anchor-free oriented proposal generator for object detection. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–11, 2022. 6, 2
- [10] Hongquan Cheng, Jie Zheng, Huayi Wu, Kunlun Qi, and Lihua He. A communication-efficient distributed deep learning remote sensing image change detection framework. *International Journal of Applied Earth Observation and Geoinformation*, 129:103840, 2024. 2
- [11] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *CVPR*, 2018. 6
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 4
- [13] Tao Fan, Hanlin Gu, Xuemei Cao, Chee Seng Chan, Qian Chen, Yiqiang Chen, Yihui Feng, Yang Gu, Jiaxiang Geng, Bing Luo, et al. Ten challenging problems in federated foundation models. *IEEE Transactions on Knowledge and Data Engineering*, 2025. 2
- [14] Jie Gui, Tuo Chen, Jing Zhang, Qiong Cao, Zhenan Sun, Hao Luo, and Dacheng Tao. A survey on self-supervised learning: Algorithms, applications, and future trends. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9052–9071, 2024. 2
- [15] Xin Guo, Jiangwei Lao, Bo Dang, Yingying Zhang, Lei Yu, Lixiang Ru, Liheng Zhong, Ziyuan Huang, Kang Wu, Dingxiang Hu, Huimei He, Jian Wang, Jingdong Chen, Ming Yang, Yongjun Zhang, and Yansheng Li. Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery. In *CVPR*, pages 27672–27683, 2024. 1, 2
- [16] Danfeng Hong, Bing Zhang, Xuyang Li, Yuxuan Li, Chenyu Li, Jing Yao, Naoto Yokoya, Hao Li, Pedram Ghamisi, Xiuping Jia, Antonio Plaza, Paolo Gamba, Jon Atli Benediktsson, and Jocelyn Chanussot. Spectralgpt: Spectral remote sensing foundation model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5227–5244, 2024. 2
- [17] Wenke Huang, Mang Ye, Zekun Shi, Guancheng Wan, He Li, Bo Du, and Qiang Yang. Federated learning for generalization, robustness, fairness: A survey and benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9387–9406, 2024. 2
- [18] Alex Iacob, Lorenzo Sani, Bill Marino, Preslav Aleksandrov, and Nicholas Donald Lane. Worldwide federated training of language models. *arXiv preprint arXiv:2405.14446*, 2024. 2
- [19] Sami Jaghouar, Jack Min Ong, Manveer Basra, Fares Obeid, Jannik Straube, Michael Keiblinger, Elie Bakouch, Lucas Atkins, Maziyar Panahi, et al. Intellect-1 technical report. *arXiv preprint arXiv:2412.01152*, 2024. 1
- [20] Shusen Jing, Anlan Yu, Shuai Zhang, and Songyang Zhang. FedSC: Provable federated self-supervised learning with spectral contrastive objective over non-i.i.d. data. In *ICML*, 2024. 2
- [21] Hansol Kim, Youngjun Kwak, Minyoung Jung, Jinho Shin, Youngsung Kim, and Changick Kim. Protofl: Unsupervised federated learning via prototypical distillation. In *ICCV*, pages 6470–6479, 2023. 3
- [22] Gihun Lee, Minchan Jeong, Sangmook Kim, Jaehoon Oh, and Se-Young Yun. Fedsol: Stabilized orthogonal learning with proximal restrictions in federated learning. In *CVPR*, pages 12512–12522, 2024. 3
- [23] Jingtao Li, Lingjuan Lyu, Daisuke Iso, Chaitali Chakrabarti, and Michael Spranger. MocoSFL: enabling cross-client collaborative self-supervised learning. In *ICLR*, 2023. 2
- [24] Mingyi Li, Xiao Zhang, Qi Wang, Tengfei Liu, Ruofan Wu, Weiqiang Wang, Fuzhen Zhuang, Hui Xiong, and Dongxiao Yu. Resource-aware federated self-supervised learning with global class representations. In *NeurIPS*, 2024. 2, 3, 7
- [25] Shenghui Li, Fanghua Ye, Meng Fang, Jiayu Zhao, Yun-Hin Chan, Edith C-H Ngai, and Thiemo Voigt. Synergizing foundation models and federated learning: A survey. *arXiv preprint arXiv:2406.12844*, 2024. 2
- [26] Xuyang Li, Danfeng Hong, and Jocelyn Chanussot. S2mae: A spatial-spectral pretraining foundation model for spectral remote sensing data. In *CVPR*, pages 24088–24097, 2024. 2
- [27] Yansheng Li, Wei Chen, Xin Huang, Zhi Gao, Siwei Li, Tao He, and Yongjun Zhang. Mfvnet: A deep adaptive fusion

- network with multiple field-of-views for remote sensing image semantic segmentation. *Science China Information Sciences*, 66(4):140305, 2023. 2
- [28] Yansheng Li, Jieyi Tan, Bo Dang, Mang Ye, Sergey A. Bartalev, Stanislav Shinkarenko, Linlin Wang, Yingying Zhang, Lixiang Ru, Xin Guo, Liangqi Yuan, Lei Yu, Jingdong Chen, Ming Yang, José Marcato Junior, and Yongjun Zhang. Unleashing the potential of remote sensing foundation models via bridging data and computility islands. *The Innovation*, page 100841, 2025. 1, 2
- [29] Yansheng Li, Linlin Wang, Tingzhu Wang, Xue Yang, Junwei Luo, Qi Wang, Youming Deng, Wenbin Wang, Xian Sun, Haifeng Li, et al. Star: A first-ever dataset and a large-scale benchmark for scene graph generation in large-size satellite imagery. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(3):1832–1849, 2025. 6
- [30] Yansheng Li, Yuning Wu, Gong Cheng, Chao Tao, Bo Dang, Yu Wang, Jiahao Zhang, Chuge Zhang, Yiting Liu, Xu Tang, et al. Meet: A million-scale dataset for fine-grained geospatial scene classification with zoom-free remote sensing imagery. *IEEE/CAA Journal of Automatica Sinica*, 12(5):1004–1023, 2025. 6
- [31] Xinting Liao, Weiming Liu, Chaochao Chen, Pengyang Zhou, Fengyuan Yu, Huabin Zhu, Binhui Yao, Tao Wang, Xiaolin Zheng, and Yanchao Tan. Rethinking the representation in federated unsupervised learning with non-iid data. In *CVPR*, pages 22841–22850, 2024. 3, 7
- [32] Hui Lin, Chao Zhang, Danfeng Hong, Kexin Dong, and Congcong Wen. Fedrscip: Federated learning for remote sensing scene classification using vision-language models. *IEEE Geoscience and Remote Sensing Magazine*, 2025. 1
- [33] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 2, 6
- [34] Senliang Lu, Yehang Chen, Yuan Chen, Peijun Li, Junqi Sun, Changye Zheng, Yujian Zou, Bo Liang, Mingwei Li, Qinggeng Jin, Enming Cui, Wansheng Long, and Bao Feng. General lightweight framework for vision foundation model supporting multi-task and multi-center medical image analysis. *Nature Communications*, 16(1):2097, 2025. 2
- [35] Siqi Lu, Junlin Guo, James R Zimmer-Dauphinee, Jordan M Nieusma, Xiao Wang, Parker VanValkenburgh, Steven A Wernke, and Yuankai Huo. Vision foundation models in remote sensing: A survey. *arXiv preprint arXiv:2408.03464*, 2025. 2
- [36] Ekdeep Lubana, Chi Ian Tang, Fahim Kawsar, Robert Dick, and Akhil Mathur. Orchestra: Unsupervised federated learning via globally consistent clustering. In *ICML*, pages 14461–14484, 2022. 3
- [37] Junwei Luo, Zhen Pang, Yongjun Zhang, Tingzhu Wang, Linlin Wang, Bo Dang, Jiangwei Lao, Jian Wang, Jingdong Chen, Yihua Tan, et al. Skysensegpt: A fine-grained instruction tuning dataset and model for remote sensing vision-language understanding. *arXiv preprint arXiv:2406.10100*, 2024. 2
- [38] Emmanuel Maggiori, Yuliya Tarabalka, Guillaume Charpiat, and Pierre Alliez. Can semantic labeling methods generalize to any city? the inria aerial image labeling benchmark. In *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. IEEE, 2017. 6, 3
- [39] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 1273–1282, 2017. 3
- [40] Yasar Abbas Ur Rehman, Yan Gao, Jiajun Shen, Pedro Porto Buarque de Gusmão, and Nicholas Lane. Federated self-supervised learning for video understanding. In *ECCV*, pages 506–522, 2022. 3
- [41] Yasar Abbas Ur Rehman, Yan Gao, Pedro Porto Buarque de Gusmão, Mina Alibeigi, Jiajun Shen, and Nicholas D. Lane. L-dawa: Layer-wise divergence aware weight aggregation in federated self-supervised visual representation learning. In *ICCV*, pages 16464–16473, 2023. 3, 7
- [42] Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, pages 2021–2031, 2020. 8
- [43] Lorenzo Sani, Alex Iacob, Zeyu Cao, Royson Lee, Bill Marino, Yan Gao, Dongqi Cai, Zexi Li, Wanru Zhao, Xinchu Qiu, and Nicholas D. Lane. Photon: Federated llm pre-training. *arXiv preprint arXiv:2411.02908*, 2024. 1
- [44] Lorenzo Sani, Alex Iacob, Zeyu Cao, Bill Marino, Yan Gao, Tomas Paulik, Wanru Zhao, William F Shen, Preslav Aleksandrov, Xinchu Qiu, et al. The future of large language model pre-training is federated. *arXiv preprint arXiv:2405.10853*, 2024. 2
- [45] Jun Sun, Tianyi Chen, Georgios B Giannakis, Qinmin Yang, and Zaiyue Yang. Lazily aggregated quantized gradient innovation for communication-efficient federated learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(4):2031–2044, 2020. 2
- [46] Xian Sun, Peijin Wang, Wanxuan Lu, Zicong Zhu, Xiaonan Lu, Qibin He, Junxi Li, Xuee Rong, Zhujun Yang, Hao Chang, Qinglin He, Guang Yang, Ruiping Wang, Jiwen Lu, and Kun Fu. Ringmo: A remote sensing foundation model with masked image modeling. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–22, 2023. 2
- [47] Jamie Tolan, Hung-I Yang, Benjamin Nosarzewski, Guillaume Couairon, Huy V Vo, John Brandt, Justine Spore, Sayantan Majumdar, Daniel Haziza, Janaki Vamaraju, et al. Very high resolution canopy height maps from rgb imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar. *Remote Sensing of Environment*, 300:113888, 2024. 1
- [48] Xin-Yi Tong, Gui-Song Xia, and Xiao Xiang Zhu. Enabling country-scale land cover mapping with meter-resolution satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 196:178–196, 2023. 6
- [49] Devis Tuia, Konrad Schindler, Begüm Demir, Xiao Xiang Zhu, Mrinalini Kochupillai, Sašo Džeroski, Jan N. van Rijn,

- Holger H. Hoos, Fabio Del Frate, Mihai Datcu, Volker Markl, Bertrand Le Saux, Rochelle Schneider, and Gustau Camps-Valls. Artificial intelligence to advance earth observation: A review of models, recent trends, and pathways forward. *IEEE Geoscience and Remote Sensing Magazine*, pages 2–25, 2024. 1
- [50] Ye Lin Tun, Chu Myaet Thwal, Le Quang Huy, Minh NH Nguyen, and Choong Seon Hong. Lw-fedssl: Resource-efficient layer-wise federated self-supervised learning. *arXiv preprint arXiv:2401.11647*, 2024. 3
- [51] Di Wang, Jing Zhang, Bo Du, Gui-Song Xia, and Dacheng Tao. An empirical study of remote sensing pretraining. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–20, 2023. 2
- [52] Junjue Wang, Zhuo Zheng, Ailong Ma, Xiaoyan Lu, and Yanfei Zhong. LoveDA: A remote sensing land-cover dataset for domain adaptive semantic segmentation. In *NeurIPS*, 2021. 6, 2
- [53] Lirui Wang, Kaiqing Zhang, Yunzhu Li, Yonglong Tian, and Russ Tedrake. Does learning from decentralized non-IID unlabeled data benefit from self supervision? In *ICLR*, 2023. 2
- [54] Yuan Wang, Huazhu Fu, Renuga Kanagavelu, Qingsong Wei, Yong Liu, and Rick Siow Mong Goh. An aggregation-free federated learning for tackling data heterogeneity. In *CVPR*, pages 26233–26242, 2024. 2
- [55] Gui-Song Xia, Jingwen Hu, Fan Hu, Baoguang Shi, Xiang Bai, Yanfei Zhong, Liangpei Zhang, and Xiaoqiang Lu. Aid: A benchmark data set for performance evaluation of aerial scene classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55(7):3965–3981, 2017. 6, 2
- [56] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *CVPR*, 2018. 6, 2
- [57] Weiyang Xie, Haowei Li, Jitao Ma, Yunsong Li, Jie Lei, Donglai Liu, and Leyuan Fang. Jointsq: Joint sparsification-quantization for distributed learning. In *CVPR*, pages 5778–5787, 2024. 2
- [58] Weiyang Xie, Jitao Ma, Tianen Lu, Yunsong Li, Jie Lei, Leyuan Fang, and Qian Du. Distributed deep learning with gradient compression for big remote sensing image interpretation. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–15, 2024. 2
- [59] Xingyu Xie, Zhijie Lin, Kim-Chuan Toh, and Pan Zhou. Loco: Low-bit communication adaptor for large-scale model training. *arXiv preprint arXiv:2407.04480*, 2024. 5
- [60] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. 6, 7
- [61] Rui Yan, Liangqiong Qu, Qingyue Wei, Shih-Cheng Huang, Liyue Shen, Daniel L. Rubin, Lei Xing, and Yuyin Zhou. Label-efficient self-supervised federated learning for tackling data heterogeneity in medical imaging. *IEEE Transactions on Medical Imaging*, 42(7):1932–1943, 2023. 3, 7
- [62] Kunping Yang, Gui-Song Xia, Zicheng Liu, Bo Du, Wen Yang, Marcello Pelillo, and Liangpei Zhang. Semantic change detection with asymmetric siamese networks. *arXiv preprint arXiv:2010.05687*, 2020. 6, 3
- [63] Yiwen Ye, Yutong Xie, Jianpeng Zhang, Ziyang Chen, Qi Wu, and Yong Xia. Continual self-supervised learning: Towards universal multi-modal medical data representation learning. In *CVPR*, pages 11114–11124, 2024. 2
- [64] Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin Li. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *ICML*, 2024. 8
- [65] Junyuan Zhang, Songhua Liu, and Xinchao Wang. One-shot federated learning via synthetic distiller-distillate communication. In *NeurIPS*, pages 102611–102633, 2024. 8
- [66] Mi Zhang, Bingnan Yang, Xiangyun Hu, Jianya Gong, and Zuxun Zhang. Foundation model for generalist remote sensing intelligence: Potentials and prospects. *Science Bulletin*, 69(23):3652–3656, 2024. 1
- [67] Tuo Zhang, Tiantian Feng, Samiul Alam, Dimitrios Dimitriadis, Sunwoo Lee, Mi Zhang, Shrikanth S Narayanan, and Salman Avestimehr. Gpt-fl: Generative pre-trained model-assisted federated learning. In *CVPR*, pages 1761–1770, 2025. 1
- [68] Weiming Zhuang, Yonggang Wen, and Shuai Zhang. Divergence-aware federated self-supervised learning. In *ICLR*, 2022. 3, 7