
GreenHyperSpectra: A multi-source hyperspectral dataset for global vegetation trait prediction

Eya Cherif*^{† 1,2,3}, Arthur Ouaknine*^{3,4}, Luke A. Brown⁵, Phuong D. Dao^{6,7,8}, Kyle R. Kovach⁹, Bing Lu¹⁰, Daniel Mederer¹, Hannes Feilhauer*^{1,2,12,13}, Teja Kattenborn*^{11,12}, and David Rolnick*^{3,4}

¹Institute for Earth System Science and Remote Sensing, Leipzig University, Germany

²Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Leipzig University, Germany

³Mila – Québec AI Institute, Canada

⁴McGill University, Canada

⁵School of Science, Engineering & Environment, University of Salford, UK

⁶Department of Agricultural Biology, Colorado State University, USA

⁷Graduate Degree Program in Ecology, Colorado State University, USA

⁸School of Global Environmental Sustainability, Colorado State University, USA

⁹Department of Forest and Wildlife Ecology, University of Wisconsin, USA

¹⁰Department of Geography, Simon Fraser University, Canada

¹¹Chair of Sensor-based Geoinformatics (geosense), University of Freiburg, Germany

¹²German Centre for Integrative Biodiversity Research (iDiv), Halle-Jena-Leipzig, Germany

¹³Helmholtz-Centre for Environmental Research (UFZ), Leipzig, Germany

Abstract

Plant traits such as leaf carbon content and leaf mass are essential variables in the study of biodiversity and climate change. However, conventional field sampling cannot feasibly cover trait variation at ecologically meaningful spatial scales. Machine learning represents a valuable solution for plant trait prediction across ecosystems, leveraging hyperspectral data from remote sensing. Nevertheless, trait prediction from hyperspectral data is challenged by label scarcity and substantial domain shifts (*e.g.* across sensors, ecological distributions), requiring robust cross-domain methods. Here, we present GreenHyperSpectra, a pretraining dataset encompassing real-world cross-sensor and cross-ecosystem samples designed to benchmark trait prediction with semi- and self-supervised methods. We adopt an evaluation framework encompassing in-distribution and out-of-distribution scenarios. We successfully leverage GreenHyperSpectra to pretrain label-efficient multi-output regression models that outperform the state-of-the-art supervised baseline. Our empirical analyses demonstrate substantial improvements in learning spectral representations for trait prediction, establishing a comprehensive methodological framework to catalyze research at the intersection of representation learning and plant functional traits assessment. We also share the dataset¹, code and pretrained model objects for this study [here](#).

1 Introduction

Plant functional traits are a fundamental component of biodiversity assessment, offering insights into plant productivity, ecological interactions, resilience, and adaptation to environmental change [17, 48, 97, 126]. Leaf traits such as leaf mass per area, as well as chlorophyll, nitrogen, and carbon content, are key to understanding plant growth dynamics and ecosystem processes such as carbon cycling and productivity [7, 9, 28, 83, 127, 128]. The monitoring of these traits is thus crucial for understanding ecosystem function and guiding biodiversity conservation strategies [18, 75, 125].

*Core Team

[†]Corresponding author: eya.cherif@uni-leipzig.de

¹GreenHyperSpectra dataset: <https://huggingface.co/datasets/Avatarr05/GreenHyperSpectra>

Initiatives such as the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services (IPBES; [38, 62]) have raised global awareness around the urgent need for monitoring functional traits and their diversity across spatial scales. However, we still lack efficient tools to track these functional traits in space and time. Hyperspectral remote sensing from airborne and satellite systems has emerged as a promising tool to bridge this gap, enabling non-destructive, scalable, and repeatable reflectance measurements that can be used to predict these traits [8, 24, 65, 93, 111]. Hyperspectral sensors measure radiation reflected from the ground across hundreds of narrow, contiguous spectral bands, spanning the visible to shortwave infrared (VNIR + SWIR) domains. These measurements are informative for trait prediction, as they are directly influenced by the chemical and structural characteristics of plant leaves and canopies [64]. Plant trait prediction from hyperspectral data is inherently a regression problem, and was initially often explored with Partial Least Squares Regression (PLSR, [50]) to link hyperspectral observations to individual traits [42, 92, 57, 109]. However, non-parametric machine learning methods, in particular deep learning, have recently been explored to offer greater flexibility in modeling complex, non-linear trait-spectral relationships and trait-trait interactions [24, 105, 54]. In this context, trait prediction is framed as a multi-output regression task within a multi-task learning framework, where the outputs, representing multiple plant traits, are inherently correlated.

Trait prediction poses fundamental machine learning challenges including heterogeneous target distributions requiring specialized multi-task methods [24], extreme label scarcity, and significant distributional biases (*e.g.* spatial, Figure 2). Current supervised approaches demonstrate limited cross-domain generalization due to training on sparse, non-representative datasets [29, 76]. Hyperspectral data further compound these challenges exhibiting substantial covariate shifts across acquisition conditions, sensor configurations and resolutions, radiometric calibrations and variable input modalities. To address these limitations, we introduce GreenHyperSpectra, a large spectral dataset designed to improve representation robustness against domain adaptation challenges while being collected from multiple ecosystems, instruments, spatial resolutions, and acquisition conditions. This dataset enables semi- and self-supervised learning applications, which take advantage of vast unlabeled spectral data, providing extensive coverage and variability to facilitate benchmarking.

Our contributions include: ① Building GreenHyperSpectra, a dataset for pretraining consisting of cross-domain samples and substantially expanding available datasets for representation learning; ② framing a suite of semi- and self-supervised methods for multi-output regression with one dimensional (1D) hyperspectral data; ③ comparing these methods with fully supervised baseline, highlighting the superior performance of the former, particularly in scenarios with limited labeled data; ④ Testing how well such methods generalize across variable inputs, representing the diversity of sensor settings (full-range (VNIR+SWIR) vs. half-range (VNIR-only), see Figure 1).

2 Related work

Despite increasing hyperspectral data availability, plant trait prediction is still largely constrained by the lack of large annotated datasets [29, 104]. Trait labels are costly and time-consuming to obtain, often requiring field sampling and laboratory analysis [27, 6]. Available datasets lack harmonization and often differ in sampling strategy, measurement assumptions and protocols. As a result, most labeled datasets are geographically and ecologically limited, with sparse coverage across space and time, as well as across ecosystems and acquisition conditions. To address this, previous studies explored synthetic datasets generated from Radiative Transfer Models (RTMs, [44]) that simulate canopy spectral responses under diverse conditions. In this context, it motivated hybrid approaches that combine RTM-based simulations with machine learning [30, 70, 73, 99, 114, 119].




Figure 1: Overview of the semi/self-supervised framework for multi-trait regression task.

However, several comparisons reveal important limitations, with models trained on real, multi-site datasets consistently outperforming those using only synthetic spectra [70, 73], highlighting a persistent domain gap between simulated and field data.

In this regard, there is a need of large-scale unlabeled spectral datasets from real measurements to pretrain models. To support this, a number of benchmarking efforts have been introduced [45, 117, 12]. However, most of these initiatives are built from a single sensor type, restricting spectral diversity and limiting generalization to new acquisition conditions. Moreover, these datasets often consist of full hyperspectral imagery, where spatially contiguous pixels are subject to spatial autocorrelation. This spatial redundancy reduces the spectral variability necessary for training models that generalize well in trait prediction task, which depend on spectral rather than spatial information [103]. They also typically include a broad mix of land cover types, including non-vegetated surfaces. This introduces inefficiencies, as vegetated pixels must be sampled through additional preprocessing.

Large-scale unlabeled datasets have been leveraged for self-supervised learning to improve natural image representations [23, 4, 72] as well as in remote sensing with multispectral data [106, 46, 124, 98, 10, 5, 107]. Semi- and self-supervised learning techniques are increasingly being explored to exploit hyperspectral data for image classification, segmentation, and super-resolution [2, 80, 120]. These methods are designed to better learn the spectral representation of hyperspectral remote sensing data to reduce the need for reference labels. Approaches such as masked autoencoders (MAE) [26, 60, 81, 87, 103], contrastive learning frameworks [22, 47, 55], generative networks (GAN) [3, 56, 69, 85, 130, 134], and autoencoders (AE) [1, 49, 103, 131] have been successfully applied for land cover classification. Whereas most attempts at trait prediction using hyperspectral data relied on fully supervised pipelines, notably PLSR [92, 96, 109, 121], Gaussian Process Regression (GPR) [99, 115, 113], ANN [88] and deep learning methods [24, 78, 95], they are cardinally constrained by label scarcity, hampering their ability to reliably generalize across ecosystems, sensor platforms, and acquisition conditions. Applying semi- and self-supervised methods to trait prediction remains largely unexplored and offers a compelling direction for investigation. Recent semi- and self-supervised applications in trait prediction include vision transformers for nitrogen estimation from simulated data [49] and Long short-term memory (LSTM) models for chlorophyll prediction with limited spectral bands [132] (*i.e.* VNIR-only). While semi- and self-supervised methods show promise for trait prediction, existing models remain constrained to single traits. Moreover, existing models are typically limited to specific sensor configurations and experimental conditions, limiting generalization across sensor modalities (*e.g.* full-range vs. half-range spectrometers), acquisition geometries, and vegetation types. This underscores the need for flexible approaches that can handle heterogeneous inputs while supporting transferable predictions in diverse real-world ecological scenarios. To the best of our knowledge, no semi- or self-supervised method addresses trait prediction via multi-output regression.

3 The GreenHyperSpectra dataset

We introduce a large-scale, multi-source hyperspectral dataset comprising over 140,000 vegetation canopy surface reflectance spectra captured across diverse continents, ecosystems, sensor platforms, spatial resolutions, and measurement geometries. Unlike existing benchmarks of hyperspectral data limited to single sensors or narrow ecological domains, our dataset features a substantially larger pretraining spectral dataset supporting semi- and self-supervised learning approaches.

Acquisition platforms and sensor diversity. We curated spectral data from multiple instruments across three primary platforms: proximal, airborne, and spaceborne (Figure 2). All data were processed to the level of at-surface reflectance. Proximal measurements were obtained using field




Figure 2: **Spatial coverage of the datasets.** Points represent sample locations of GreenHyperSpectra compared to the existing labeled dataset. GreenHyperSpectra data span diverse vegetation type and acquisition conditions.

spectrometers such as the ASD FieldSpec and SVC HR-1024i, typically positioned in a close range in nadir orientation to record top-of-canopy reflectance. Airborne data were acquired using high-spectral resolution sensors, including the AVIRIS-Next Generation, AVIRIS-Classic, NEON Airborne Observation Platform (AOP) and Specim AISAFenix instruments, which cover landscape-level vegetation scenes with variable viewing geometries and meter-scale spatial resolutions. Spaceborne acquisitions were collected from missions such as PRISMA, Hyperion, EMIT, and EnMAP, offering a larger scale observations at 30-60 m resolution with varying viewing geometry. Table 1 summarizes the platforms, spectral properties, and scene-level characteristics associated with each acquisition (more details see Appendix A). The pre-processing of spectra harmonization is described in Appendix A.

Platform	GSD	Spectral res.	#Samples
Proximal	<1 m	1–4 nm	5620
Airborne	1–20 m	3–7 nm	96699
Spaceborne	30–60 m	6–12 nm	36059

Table 1: Specifications of spectroscopy instruments with different platforms.

The multi-platform nature of our dataset introduces valuable reflectance signal variability through differences in spatial and spectral resolution, sun-sensor geometry, scene heterogeneity, background conditions and pre-processing from radiance to reflectance. This variability, often lacking in single-platform or synthetic datasets, is essential to develop generalizable models capable of scaling across diverse remote sensing contexts [24]. While satellite-based datasets such as SpectralEarth [12] provide temporally rich but sensor-specific imagery, our dataset uniquely incorporates multi-sensor observations across spaceborne, airborne, and proximal platforms.

Spatial and temporal coverage. The dataset includes samples from diverse biomes, with acquisitions spanning from 1992 to 2024, capturing broad ecological and climatic variability across a wide range of environments. Figure 2 maps the global spatial distribution of GreenHyperSpectra and a pool of previously aggregated datasets for plant trait prediction (see details in § 4). While the compiled labeled dataset is spatially limited, GreenHyperSpectra encompasses substantially broader spatial coverage and environmental heterogeneity, better representing real-world remote sensing operational conditions (more details in Appendix A).

4 Benchmarking methods and protocols

Trait-annotated dataset. For benchmarking different semi- and self-supervised methods, we use an existing aggregated dataset [24] comprising 7,900 canopy reflectance spectra with co-located measurements of seven functional plant traits: leaf mass per area (Cm) [g/cm^2], leaf protein content (Cp) [g/cm^2], equivalent water thickness (Cw) [cm], leaf total chlorophyll (Cab) [$\mu\text{g}/\text{cm}^2$], carotenoids (Car) [$\mu\text{g}/\text{cm}^2$] and anthocyanins (Anth) [$\mu\text{g}/\text{cm}^2$] content, and leaf area index (LAI) [m^2/m^2]. Trait values were obtained either through direct field measurements or via community-weighted means assigned at the pixel level based on ground-measured species composition. For the analysis, we treat Cp and nitrogen as equivalent due to their strong correlation, while acknowledging that they are not strictly the same. Additionally, we introduce a derived trait, carbon-based constituents (cbc), which is computed as the difference between Cm and Cp. These data were aggregated from 50 experiments and campaigns [51, 133, 21, 13, 82, 31, 15, 25, 14, 110, 66, 19, 40, 39, 123, 118, 121, 59, 77, 53, 52, 96, 32], covering diverse vegetation types such as forests, croplands, tundra, and pastures. Trait values were harmonized by converting mass-based traits to area-based units [67]. The pre-processing of the spectra is similar to that described in Appendix A. This aggregated labeled dataset serves as the reference to train and evaluate the regression models. To enhance training stability and better capture inter-trait correlations, we applied box-cox transformation [11] to the trait values [24] for all methods.

Data splitting and evaluation protocol. To ensure consistent representation of all contributing data sources during training, we divide GreenHyperSpectra into 20 non-overlapping subsets. In each split, the proportion of samples from any given data source matches that dataset’s overall contribution to the full merged dataset. This stratified splitting strategy maintains the natural diversity of vegetation types, sensors, and acquisition conditions, while preventing bias from individual sources by creating consistent and representative subsets suitable for semi- and self-supervised methods. For the labeled dataset, we define standardized train and validation splits using a 80/20 hold-out strategy. The 80% portion is combined with the pretraining spectral dataset for calibration, while the remaining 20% is fixed for all experiments and used to evaluate all methods. For out of distribution (OOD) evaluation experiments (detailed in §5), we perform cross validation across the 50 labeled datasets. Specifically,




Figure 3: **Overview of the semi- and self-supervised frameworks.** (3a) The semi-supervised regression GAN framework (SR-GAN): the generator maps a random noise z to synthetic samples $\hat{x}$, while the discriminator processes ① fake samples (x_{fake}), ② unlabeled real samples (x_{unlb}), and ③ labeled real data samples (x_{lb}) with associated traits (y), optimizing fake (L_{fake}), unlabeled (L_{unlb}), and labeled (L_{lb}) losses respectively. (3b) The RTM-based autoencoder (RTM-AE) predicts traits from labeled embeddings while reconstructing spectra ($x \rightarrow \hat{x}$, (L_{recon})). (3c). The 1D masked autoencoder framework (1D-MAE) reconstructs masked spectra through tokenization, (L_{recon}); the learned representations are then used for trait prediction (L_{lb}). **Abbreviations:** x_{fake} : generated fake spectra from the generator; x_{unlb} : unlabeled sample from GreenHyperSpectra; x_{lb} : spectra sample from the labeled data; L_{unlb} : unlabeled loss; L_{lb} : labeled loss; L_{recon} : reconstruction loss; L_{fake} : feature contrasting loss; RTM: radiative transfer model; AE: autoencoder; MAE: masked autoencoder.

from the 50 annotated sub-datasets, we hold out five datasets at a time for testing. The remaining datasets are used for training, with their data further split into 80% for training and 20% for validation.

Supervised baseline method. We consider a supervised CNN-based method [24, 70] as a baseline, selected for its state-of-the-art performance in multi-trait plant prediction from a sparse annotated dataset. It is built upon EfficientNet-B0 [100] specifically framed for 1D feature extraction. The network employs multi-output regression to simultaneously predict the seven plant traits, a strategy that demonstrably outperforms single-trait modeling approaches [24, 89].

Semi-supervised regression generative adversarial network (SR-GAN). We frame the SR-GAN framework [71] to address hyperspectral plant trait prediction. Our implementation employs a 1D convolutional GAN architecture designed specifically for spectral data processing. In this setup, the generator learns to produce synthetic reflectance spectra, while the discriminator simultaneously performs trait regression and learns discriminative feature representations. The training objective is formulated as a composite loss that encourages the discriminator to pull real spectral samples closer in the feature space, while pushing representations of generated (synthetic) samples further apart. This contrastive learning approach allows the model to leverage unlabeled data by learning




Figure 4: **Evaluation of trait prediction with variable-size labeled sets.** Validation performance (R^2) as a function of labeled data percentage used for training. The average R^2 performance across all traits is indicated by the dashed box. The higher R^2 , the better. For trait abbreviations, see Sec. 4.

informative spectral embeddings. The overall architecture of the SR-GAN framework is illustrated in Figure 3a and detailed formulations of the loss components are provided in Appendix B.1.

Radiative transfer model based autoencoder (RTM-AE). We introduce a version of the autoencoder framework proposed by [94], which replaces the decoder with a non-learnable RTM module to reconstruct spectra, thereby integrating physical constraints into the modeling process. Specifically, our implementation employs PROSAIL-PRO [43], constraining the latent space to correspond directly to plant traits. PROSAIL-PRO is an RTM that combines the leaf reflectance model (PROSPECT,[63]) with the canopy reflectance model (4SAIL,[112]). PROSPECT simulates leaf reflectance and transmittance based on biochemical composition and internal structure, while 4SAIL models the propagation of light through a vegetation canopy. Together, they simulate canopy spectral reflectance in the 400–2500 nm range using inputs such as chlorophyll content, leaf area index, and leaf angle. As previously mentioned regarding the gap between RTM-simulated and real-world spectra (§ 2), we address the inherent discrepancies between RTM-generated and observed spectra, primarily resulting from simplified geometric assumptions within the model, by implementing a learnable correction layer that refines the simulated output [94]. Our enhanced framework introduces three key improvements over the original design [94]: (1) incorporation of PROSAIL-PRO, (2) application of a supervised loss component targeting trait predictions, and (3) implementation of a composite reconstruction loss combining cosine similarity and mean absolute error to capture both spectral shape characteristics and amplitude information. The overall architecture is illustrated in Figure 3b and the specifications are detailed in Appendix B.2 .

Masked autoencoder (MAE). We adopt a MAE framework, originally designed for land cover classification [103], to predict plant traits with hyperspectral data. The model leverages self-supervised learning by reconstructing randomly masked spectral regions, enabling the extraction of meaningful representations from unlabeled hyperspectral signatures. Similarly to the RTM-AE, our adaptation incorporates a modified reconstruction objective that combines cosine similarity and mean squared error (MSE) with appropriate weighting, allowing the model to capture both spectral shape characteristics and amplitude information. For downstream trait prediction, we attach a multi-output regression head to the latent features and fine-tune the model using labeled data. The overall architecture is illustrated in Figure 3c and ablation studies for the MAE architecture are provided in Appendix B.3.

5 Experimental settings

This section describes experimental setups used to benchmark and evaluate the performance of models trained on semi- and self-supervised learning fashion for multi-trait plant prediction detailed in § 4. Our experimental framework is structured into four principal components to test the capabilities of

models across a range of scenarios reflecting critical use cases: comprehensive benchmarking using full-range (FR) spectra, benchmarking with half-range (HR) spectra, and assessment of OOD generalization capabilities, along with an ablation study on the design of the MAE models. Throughout these experiments, we maintain standardized data splits for both labeled and unlabeled datasets as described in § 4. Complete specifications regarding model hyperparameters, optimization settings, and implementation details are provided in Appendix B.

Full-range trait prediction. We assess all benchmark models using the full-range spectra spanning 400–2450 nm (1721 bands), encompassing visible through shortwave infrared wavelengths. Results are presented in § 6 with Table 2.

Sample sensitivity analysis. We examine the impact of label availability by simulating different levels of supervision and varying the amount of labeled data used for training from 20% to 100% while maintaining a consistent unlabeled dataset. Complementing this approach, we conduct experiments varying the quantity of unlabeled training data while maintaining fixed labeled data proportions to determine how unlabeled data volume influences model performance; note that in these experiments, we use only a subset of the full GreenHyperSpectra dataset (80,000 samples). Results are presented in § 6, with Figure 4 and 5.

Half-range trait prediction. A common constraint faced with satellite-based Earth observations is that many sensors do not cover the full spectrum. To evaluate model performance in this scenario, we replicate our benchmark procedure using only the half-range spectral subset spanning 400–900nm (500 bands). All models are trained on this spectral subset. Additionally, we implement an evaluation for the MAE architecture, where a model pretrained on full-range spectra is applied to half-range spectra inputs (this is possible only for the MAE models as the masking procedure means that they can accommodate variable input sizes). Results are presented in § 6 and Table 3.

OOD evaluation. To assess each model’s robustness to real-world distribution shifts, we perform a cross-dataset evaluation as described in § 4. We compute a macro-level performance metric by aggregating predictions across all held-out datasets. This setup reflects practical challenges in ecological monitoring applications, where spectral variability arises from differences in acquisition conditions, sensor platforms, or environmental contexts. Additionally, this approach ensures a broader coverage of trait value ranges, which often remain underrepresented when test sets are randomly sampled. Due to computational constraints, we conduct this evaluation using a single training run. To reduce the sensitivity of R^2 to unbalanced number of samples across the 50 aggregated datasets, we compute the macro-average over five random subsamples within each dataset, each constrained to the maximum number of 30 samples allowed per set, and report the mean and standard deviation of the resulting metrics. Results are presented in § 6 and Table 4.

Ablation studies on MAE. We conduct comprehensive ablation experiments across several dimensions to consistently evaluate design trade-offs of the MAE models in spectral representation learning. First, we explore architectural complexity through a grid search spanning transformer configurations with varying numbers of layers {6, 8, 10} and attention heads {4, 8, 16}. We select the configuration demonstrating optimal performance on the downstream trait prediction task for subsequent experiments. Considering this optimal architecture, we investigate alternative loss formulations for spectral reconstruction. Beyond conventional MSE, we examine hybrid approaches which incorporate cosine similarity loss weighted by a coefficient $\alpha \in \{1, 0.1, 0.01, 0\}$ to enhance capture of spectral shape specificity. Finally, we assess the effect of token granularity by varying patch sizes (10, 20, 40, and 430) used during spectral masking and reconstruction in the full-range scenario. These targeted ablations informed the final MAE configuration used in our other benchmark evaluations. All results of the aforementioned ablation studies are presented in Tables 16, 17 and 18 in Appendix B.3.

Evaluation metrics. For each experimental setting, we report the performance metrics averaged across three random seeds to measure the variability related to stochastic training effects. Our evaluation framework employs two complementary metrics: the coefficient of determination (R^2) and the normalized root mean square error (nRMSE). The nRMSE (in %) is computed by normalizing the root mean square error by the range of the traits observations (1–99% quantile), providing a scale-invariant measure of prediction error.




Figure 5: **Evaluation of trait prediction with variable-size unlabeled sets.** Validation performance (R^2) as a function of the percentage of unlabeled data used for training. The average R^2 performance is indicated by the dashed box. The higher R^2 , the better. For trait abbreviations, see Sec. 4.

	cab	cw	cm	LAI	cp	cbc	car	anth	average
R^2 ($\uparrow$)									
Supervised	0.517 (± 0.012)	0.621 (± 0.009)	0.667 (± 0.005)	0.561 (± 0.011)	0.651 (± 0.002)	0.679 (± 0.006)	0.544 (± 0.018)	0.454 (± 0.020)	0.587 (± 0.010)
SR_GAN	0.574 (± 0.008)	0.572 (± 0.013)	0.669 (± 0.008)	0.538 (± 0.005)	0.558 (± 0.005)	0.704 (± 0.004)	0.578 (± 0.006)	0.541 (± 0.041)	0.592 (± 0.011)
RTM_AE	0.584 (± 0.023)	0.658 (± 0.011)	0.679 (± 0.021)	0.552 (± 0.029)	0.671 (± 0.015)	0.689 (± 0.021)	0.566 (± 0.038)	0.337 (± 0.125)	0.592 (± 0.035)
MAE_FR_LP	0.462 (± 0.002)	0.514 (± 0.009)	0.577 (± 0.010)	0.470 (± 0.007)	0.464 (± 0.007)	0.611 (± 0.008)	0.412 (± 0.002)	0.217 (± 0.015)	0.466 (± 0.008)
MAE_FR_FT	0.515 (± 0.026)	<u>0.634</u> (± 0.014)	0.716 (± 0.027)	0.615 (± 0.016)	0.676 (± 0.010)	0.727 (± 0.034)	0.649 (± 0.012)	0.598 (± 0.025)	0.641 (± 0.020)
nRMSE ($\downarrow$)									
Supervised	17.341 (± 0.221)	13.103 (± 0.150)	10.671 (± 0.083)	17.487 (± 0.218)	10.229 (± 0.021)	10.634 (± 0.099)	13.647 (± 0.274)	16.465 (± 0.333)	13.697 (± 0.175)
SR_GAN	16.277 (± 0.161)	13.918 (± 0.208)	10.658 (± 0.127)	17.829 (± 0.091)	11.565 (± 0.074)	10.258 (± 0.076)	13.123 (± 0.095)	15.084 (± 0.662)	13.589 (± 0.187)
RTM_AE	16.096 (± 0.437)	12.443 (± 0.201)	10.492 (± 0.350)	17.556 (± 0.573)	9.989 (± 0.224)	10.495 (± 0.350)	13.297 (± 0.596)	18.091 (± 1.761)	13.557 (± 0.561)
MAE_FR_LP	18.297 (± 0.034)	14.830 (± 0.135)	12.052 (± 0.146)	19.094 (± 0.133)	12.746 (± 0.080)	11.747 (± 0.120)	15.501 (± 0.030)	19.727 (± 0.185)	15.499 (± 0.108)
MAE_FR_FT	17.374 (± 0.456)	<u>12.861</u> (± 0.256)	9.856 (± 0.462)	16.285 (± 0.333)	9.916 (± 0.155)	9.833 (± 0.599)	11.975 (± 0.202)	14.117 (± 0.438)	12.777 (± 0.363)

Table 2: **Evaluation of trait prediction with full-range (FR) samples.** Trait-wise performance (mean $\pm$ standard deviation) includes R^2 ($\uparrow$) and nRMSE ($\downarrow$) metrics. Competing methods are: fully supervised baseline (‘Supervised’); MAE with full-range training and linear probing (‘MAE_FR_LP’); MAE with full-range training and fine-tuning (‘MAE_FR_FT’); RTM-based auto-encoder (‘RTM_AE’); and semi-supervised regression GAN (‘SR-GAN’). In RTM-AE, cbc is not directly predicted but is derived from cm and cp estimates (cm – cp). We **bold** and underline best and second best scores respectively. Trait abbreviations are detailed in Sec. 4.

6 Results and discussions

Labeled and unlabeled data regimes. To assess each model’s sensitivity to the quantity of annotated samples, we analyze R^2 as a function of the proportion of available labeled and unlabeled data, as shown in Figures 4 and 5 respectively. The corresponding trends for nRMSE are presented in Appendix D. We observed that models leveraging unlabeled data through semi- and self-supervised methods consistently outperformed the fully supervised baseline, particularly in low-data regimes (20–40% labeled data). Notably, semi- and self-supervised methods achieved higher average R^2 and lower nRMSE scores across most traits as labeled data availability decreased (Fig. 4). This demonstrates that access to a large set of unlabeled spectra through GreenHyperSpectra substantially enhances model performance, leading to improved trait prediction accuracy. Interestingly, varying the size of the dataset for pretraining did not substantially impact performance. This suggests that the stratified splitting protocol (§ 4), ensuring consistent coverage across spectral sources, vegetation types, and acquisition conditions, plays a critical role in efficiently exploiting the available unlabeled data, even when subsampled.

	cab	cw	cm	LAI	cp	cbc	car	anth	average
R^2 ($\uparrow$)									
Sup_HR	0.277 (± 0.105)	0.072 (± 0.032)	0.197 (± 0.082)	0.048 (± 0.110)	0.197 (± 0.080)	0.219 (± 0.074)	0.126 (± 0.135)	0.166 (± 0.052)	0.163 (± 0.084)
SR-GAN_HR	0.496 (± 0.017)	0.336 (± 0.006)	0.356 (± 0.011)	0.428 (± 0.010)	0.371 (± 0.010)	0.381 (± 0.008)	0.455 (± 0.015)	0.598 (± 0.020)	0.427 (± 0.012)
RTM-AE_HR	0.582 (± 0.024)	0.450 (± 0.014)	0.472 (± 0.031)	0.541 (± 0.019)	0.546 (± 0.019)	0.471 (± 0.043)	0.491 (± 0.029)	0.538 (± 0.010)	0.511 (± 0.023)
MAE_FR_HR_LP	0.466 (± 0.008)	0.220 (± 0.009)	0.271 (± 0.006)	0.378 (± 0.004)	0.253 (± 0.008)	0.274 (± 0.003)	0.434 (± 0.004)	0.234 (± 0.013)	0.316 (± 0.007)
MAE_FR_HR_FT	0.578 (± 0.011)	0.553 (± 0.009)	0.655 (± 0.012)	0.540 (± 0.012)	0.612 (± 0.022)	0.642 (± 0.009)	0.512 (± 0.018)	0.433 (± 0.032)	0.566 (± 0.015)
MAE_HR_LP	0.493 (± 0.004)	0.221 (± 0.014)	0.247 (± 0.007)	0.435 (± 0.003)	0.280 (± 0.006)	0.279 (± 0.003)	0.375 (± 0.010)	0.376 (± 0.019)	0.338 (± 0.008)
MAE_HR_FT	0.518 (± 0.038)	0.392 (± 0.045)	0.397 (± 0.028)	0.567 (± 0.016)	0.478 (± 0.057)	0.402 (± 0.048)	0.418 (± 0.034)	0.547 (± 0.056)	0.465 (± 0.040)
nRMSE ($\downarrow$)									
Sup_HR	21.177 (± 1.575)	20.501 (± 0.333)	16.560 (± 0.831)	25.575 (± 1.465)	15.479 (± 0.757)	16.601 (± 0.774)	18.856 (± 1.492)	20.346 (± 0.634)	19.387 (± 0.982)
SR-GAN_HR	17.706 (± 0.297)	17.339 (± 0.073)	14.869 (± 0.131)	20.007 (± 0.175)	13.809 (± 0.108)	14.827 (± 0.092)	14.914 (± 0.211)	14.133 (± 0.355)	15.950 (± 0.180)
RTM-AE_HR	16.125 (± 0.460)	15.772 (± 0.199)	13.460 (± 0.395)	17.777 (± 0.370)	11.725 (± 0.342)	13.698 (± 0.558)	14.415 (± 0.413)	15.155 (± 0.163)	14.766 (± 0.350)
MAE_FR_HR_LP	17.742 (± 0.140)	17.242 (± 0.102)	16.930 (± 0.065)	19.352 (± 0.056)	18.377 (± 0.095)	17.136 (± 0.040)	15.197 (± 0.052)	19.545 (± 0.165)	17.690 (± 0.089)
MAE_FR_HR_FT	15.771 (± 0.202)	13.055 (± 0.126)	11.652 (± 0.205)	16.647 (± 0.213)	13.242 (± 0.370)	12.037 (± 0.154)	14.109 (± 0.255)	16.805 (± 0.469)	14.165 (± 0.249)
MAE_HR_LP	17.768 (± 0.070)	18.769 (± 0.163)	16.072 (± 0.071)	19.729 (± 0.050)	14.764 (± 0.065)	15.995 (± 0.036)	15.976 (± 0.125)	17.609 (± 0.262)	17.085 (± 0.106)
MAE_HR_FT	17.313 (± 0.692)	16.578 (± 0.621)	14.381 (± 0.337)	17.263 (± 0.314)	12.561 (± 0.692)	14.563 (± 0.589)	15.415 (± 0.453)	14.973 (± 0.948)	15.381 (± 0.581)

Table 3: Evaluation of trait prediction with half-range (HR) samples. Trait-wise performance (mean $\pm$ standard deviation) includes R^2 ($\uparrow$) and nRMSE ($\downarrow$) metrics. Competing methods are: supervised baseline with HR settings ('Sup_HR'); semi-supervised SR-GAN with HR settings ('SR-GAN_HR'); RTM-based autoencoder with HR settings ('RTM-AE_HR'); MAE pretrained on full-range and fine-tuned with linear probing ('MAE_FR_HR_LP'); MAE pretrained on full-range and fine-tuned ('MAE_FR_HR_FT'); MAE pretrained on HR and fine-tuned with linear probing ('MAE_HR_LP'); and MAE pretrained on HR and fine-tuned ('MAE_HR_FT'). In RTM-AE, cbc is not directly predicted but is derived from cm and cp estimates (cm – cp). We **bold** and underline best and second best scores respectively. Trait abbreviations are detailed in Sec. 4.

Full and half range spectra analyses. Trait-specific results, reported with R^2 and nRMSE scores, are summarized in Tables 2 and 3 for the full- and half-range experiments, respectively. Among all competing methods, the fine-tuned MAE (MAE-FR-FT) outperformed all other methods on most traits when trained and tested on full-range spectra, recording the highest R^2 values and lowest nRMSE scores. Compared to the fully supervised baseline, MAE-FR-FT led to an average improvement of 9% in R^2 and 6% in nRMSE. These results underscore the effectiveness of MAEs in learning meaningful spectral representations through masked spectral reconstruction. Pretrained MAE models also exhibited strong cross-spectral generalization, performing competitively on half-range data even when pretrained on full-range spectra and applied to half-range data (MAE-FR-HR-FT). It indicates a good feature transferability and adaptability across heterogeneous sensor configurations, particularly valuable for operational deployment with multi-source data streams.

The RTM-AE model, which introduces physical interpretability into the learned latent space, underperformed compared to MAE but consistently achieved the second best results for both full- and half-range experiments. It demonstrates that aligning latent representations with RTMs to enforce physically-constrained embeddings yields promising performance while simultaneously enhancing model explainability through semantically meaningful feature disentanglement and physics-informed representation learning.

To further assess robustness of the approaches under sensor noise such as illumination differences, or sensor-specific signal-to-noise characteristics, we additionally evaluated models' performances under additive Gaussian noise at inference time (details in Tables 26-30). Zero-mean noise with standard deviations of 0.01, 0.03, and 0.05 was added across all spectral bands. Results show that MAE-FR-FT and RTM-AE are substantially more robust than the supervised baseline and GAN. For instance, at $\sigma = 0.05$, MAE-FR-FT retains an R^2 of 0.331 compared to the baseline dropping to -0.065, and both MAE-FR-FT and RTM-AE exhibit smaller increases in nRMSE under spectral corruption, highlighting their resilience.

In the half-range setting, semi- and self-supervised methods also clearly outperformed the supervised baseline. Gains ranged between 100–200% in R^2 and 8–27% in nRMSE, reinforcing the value of leveraging spectral variability from GreenHyperSpectra even under reduced spectral coverage (Figures 12 and 13).

OOD evaluation. As shown in Table 4, the fine-tuned MAE (MAE-FR-FT) had the highest performance over all other methods across traits, achieving a slight improvement in R^2 relative to the supervised baseline (0.31 vs. 0.24), along with the lowest average nRMSE. Since many traits are not

	cab	cw	cm	LAI	cp	cbc	car	anth	average
	R^2 ($\uparrow$)								
Supervised	0.362 (± 0.048)	0.193 (± 0.053)	0.446 (± 0.049)	0.074 (± 0.031)	0.183 (± 0.041)	0.449 (± 0.052)	0.181 (± 0.045)	0.055 (± 0.079)	0.243 (± 0.050)
SR_GAN	0.300 (± 0.023)	0.350 (± 0.032)	0.507 (± 0.029)	-0.199 (± 0.111)	0.273 (± 0.037)	0.548 (± 0.026)	0.221 (± 0.064)	0.197 (± 0.175)	0.275 (± 0.062)
RTM-AE	0.272 (± 0.033)	0.193 (± 0.096)	0.453 (± 0.067)	0.019 (± 0.054)	0.192 (± 0.056)	-0.075 (± 0.008)	0.266 (± 0.056)	0.067 (± 0.252)	0.173 (± 0.078)
MAE_FR_LP	0.116 (± 0.028)	0.298 (± 0.028)	0.442 (± 0.039)	0.182 (± 0.059)	0.211 (± 0.032)	0.478 (± 0.044)	0.232 (± 0.020)	0.142 (± 0.153)	0.263 (± 0.050)
MAE_FR_FT	0.271 (± 0.030)	0.28 (± 0.102)	0.575 (± 0.041)	0.229 (± 0.041)	0.275 (± 0.068)	0.582 (± 0.044)	0.165 (± 0.044)	0.112 (± 0.234)	0.311 (± 0.076)
	nRMSE ($\downarrow$)								
Supervised	19.173 (± 0.695)	25.223 (± 13.109)	14.238 (± 0.496)	22.984 (± 0.475)	17.072 (± 0.669)	14.818 (± 0.560)	19.185 (± 0.581)	23.159 (± 1.644)	19.482 (± 2.278)
SR_GAN	20.098 (± 0.373)	22.394 (± 10.712)	13.445 (± 0.404)	26.075 (± 1.018)	16.108 (± 0.662)	13.438 (± 0.578)	18.698 (± 0.568)	21.360 (± 3.596)	18.952 (± 2.239)
RTM-AE	20.493 (± 0.443)	25.137 (± 13.001)	14.138 (± 0.616)	23.652 (± 0.668)	16.978 (± 0.802)	20.742 (± 0.741)	18.155 (± 0.670)	22.874 (± 3.780)	20.271 (± 2.590)
MAE_FR_LP	22.589 (± 0.406)	23.406 (± 11.682)	14.296 (± 0.474)	21.585 (± 0.763)	16.781 (± 0.677)	14.440 (± 0.641)	18.582 (± 0.322)	22.147 (± 3.725)	19.228 (± 2.336)
MAE_FR_FT	20.505 (± 0.334)	24.018 (± 13.502)	12.466 (± 0.425)	20.842 (± 0.548)	16.069 (± 0.860)	12.907 (± 0.426)	19.371 (± 0.414)	22.422 (± 4.166)	18.575 (± 2.584)

Table 4: **Cross-dataset generalization performance.** Models are trained on labeled data from all but five datasets (see Sec 4), and evaluated on held-out datasets to assess OOD generalization. Trait-wise performance includes R^2 ($\uparrow$) and nRMSE ($\downarrow$) metrics. In RTM-AE, cbc is not directly predicted but is derived from cm and cp estimates (cm – cp). We **bold** and underline best and second best scores respectively.

associated to single-band features but instead arise from complex interactions across multiple regions of the spectrum, MAE provides a strong prior: it enforces the learning of localized correlations and long-range dependencies within hyperspectral signals, by reconstructing both across adjacent tokens and distant tokens. This prior knowledge facilitates better generalization and more efficient fine-tuning with MAE-FR-FT for the downstream regression. However, this prior alone is not sufficient. When only linear probing is applied (MAE-FR-LP), the model retains general spectral trends leading to underperformance. The necessity of fine-tuning becomes evident in our feature attribution analysis (Fig. 15), where we compared gradient amplitudes across spectral bands for MAE-FR-LP, last block fine-tuning, and full fine-tuning MAE-FR-FT models. While MAE-FR-LP exhibited diffuse and noisy attributions across broad spectral regions, fine-tuning progressively reduced gradient variance, yielding sharper and more interpretable feature importance profiles. This indicates that fine-tuning allows the pretrained prior to be refined toward trait-relevant spectral dependencies, transforming general correlations into targeted representations that drive improved predictive performance.

Other competing methods, such as SR-GAN and RTM-AE, provided modest gains over the supervised baseline. The corresponding scatter plot of the observed and predicted trait values from the different methods is presented in Fig. 14 in the Appendix.

7 Conclusions and perspectives

In this study, we introduce **GreenHyperSpectra**, a large-scale cross-sensor and cross-ecosystem spectral dataset designed to train machine learning models for plant trait prediction from hyperspectral data. Leveraging GreenHyperSpectra as a pretraining resource, we demonstrated that models using MAE consistently outperformed all other benchmarked methods, including the fully supervised baseline, across a variety of settings. The adaptability of MAE models enables their application to multi-scale remote sensing platforms, including drone, airborne, and satellite imagery, paving the way to investigate how learned spectral features, derived from a heterogeneous spectral dataset, generalize across varying spatial resolutions. MAEs also serve as a strong foundation for advanced transfer learning architectures aimed at improving predictive performance. While we explore default sensor configurations (VNIR+SWIR and VNIR), extending pretrained encoders to other spectral ranges remains open for future work. The MAE architecture shows promise for cross-domain adaptation across heterogeneous sensing modalities through fine-tuning strategies. We contribute towards global pretraining datasets for spectral embeddings while highlighting critical biases affecting generalization. Despite improvements, run-to-run variance reveals challenges in learning stable representations from various ecological data distributions. Future research should expand multi-domain spectral datasets across biomes and sensing conditions to enhance transferability and address geographical and ecosystem-level biases in annotated data. Nevertheless, our pretrained models from GreenHyperSpectra will remain valuable as labeled data from underrepresented regions increases. This study confirms that semi- and self-supervised methods with large-scale pretraining are essential for advancing ecosystem monitoring.

Acknowledgments We thank all data owners for sharing the data either by request or through the public Ecological Spectral Information System (EcoSIS), Data Publisher for Earth and Environmental Science (PANGAEA) and DRYAD platforms. EC, AO and DR acknowledge support for this work from IVADO and the Canada CIFAR AI Chairs program, and computational support from Mila – Quebec AI Institute, including in-kind support from Nvidia Corporation. EC and HF acknowledge the financial support by the Federal Ministry of Education and Research of Germany and by the Sächsische Staatsministerium für Wissenschaft Kultur und Tourismus in the program Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI.

References

- [1] M. Ahmad, M. A. Alqarni, A. M. Khan, R. Hussain, M. Mazzara, and S. Distefano. Segmented and non-segmented stacked denoising autoencoder for hyperspectral band reduction. *Optik*, 180:370–378, 2019. [3](#)
- [2] M. Ahmad, S. Shabbir, S. K. Roy, D. Hong, X. Wu, J. Yao, A. M. Khan, M. Mazzara, S. Distefano, and J. Chanussot. Hyperspectral image classification—traditional to deep models: A survey for future prospects. *IEEE journal of selected topics in applied earth observations and remote sensing*, 15:968–999, 2021. [3](#)
- [3] T. Alipour-Fard and H. Arefi. Structure aware generative adversarial networks for hyperspectral image classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 13:5424–5438, 2020. [3](#)
- [4] M. Assran, Q. Duval, I. Misra, P. Bojanowski, P. Vincent, M. Rabbat, Y. LeCun, and N. Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15619–15629, 2023. [3](#)
- [5] G. Astruc, N. Gonthier, C. Mallet, and L. Landrieu. Anysat: One earth observation model for many resolutions, scales, and modalities, 2025. [3](#)
- [6] C. Baraloto, C. Timothy Paine, S. Patino, D. Bonal, B. Hérault, and J. Chave. Functional trait variation and sampling strategies in species-rich plant communities. *Functional Ecology*, 24(1):208–216, 2010. [2](#)
- [7] K. Berger, M. Machwitz, M. Kycko, S. C. Kefauver, S. Van Wittenberghe, M. Gerhards, J. Verrelst, C. Atzberger, C. Van der Tol, A. Damm, et al. Multi-sensor spectral synergies for crop stress detection and monitoring in the optical domain: A review. *Remote sensing of environment*, 280:113198, 2022. [1](#)
- [8] K. Berger, J. Verrelst, J.-B. Féret, Z. Wang, M. Woher, M. Strathmann, M. Danner, W. Mauser, and T. Hank. Crop nitrogen monitoring: Recent progress and principal developments in the context of imaging spectroscopy missions. *Remote Sensing of Environment*, 242:111758, 2020. [2](#)
- [9] F. J. Bongers, B. Schmid, H. Bruelheide, F. Bongers, S. Li, G. von Oheimb, Y. Li, A. Cheng, K. Ma, and X. Liu. Functional diversity effects on productivity increase with age in a forest biodiversity experiment. *Nature Ecology & Evolution*, 5(12):1594–1603, 2021. [1](#)
- [10] N. I. Bountos, A. Ouaknine, I. Papoutsis, and D. Rolnick. FoMo: Multi-Modal, Multi-Scale and Multi-Task Remote Sensing Foundation Models for Forest Monitoring. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(27):27858–27868, Apr. 2025. [3](#)
- [11] G. E. P. Box and D. R. Cox. An analysis of transformations. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26(2):211–243, 12 2018. [4](#)
- [12] N. A. A. Braham, C. M. Albrecht, J. Mairal, J. Chanussot, Y. Wang, and X. X. Zhu. Spectralearth: Training hyperspectral foundation models at scale. *arXiv preprint arXiv:2408.08447*, 2024. [3](#), [4](#)

- [13] P. Brodrick, R. Pavlick, M. Bernas, J. Chapman, R. Eckert, M. Helmlinger, M. Hess-Flores, L. Rios, F. Schneider, M. Smyth, et al. Shift: Aviris-ng l2a unrectified surface reflectance version 1. ornl daac, oak ridge, tennessee, usa, 2023. 4
- [14] L. A. Brown, H. Morris, A. MacLachlan, F. D'Adamo, J. Adams, E. Lopez-Baeza, E. Albero, B. Martínez, S. Sánchez-Ruiz, M. Campos-Taberner, et al. Hyperspectral leaf area index and chlorophyll retrieval over forest and row-structured vineyard canopies. *Remote Sensing*, 16(12):2066, 2024. 4
- [15] A. C. Burnett, S. P. Serbin, and A. Rogers. Source: sink imbalance detected with leaf-and canopy-level spectroscopy in a field-grown crop. *Plant, Cell & Environment*, 44(8):2466–2479, 2021. 4
- [16] P. Campbell. Eo-1 hyperion for vegetation function at mongu. <http://ecosis.org>, 2016. Ecological Spectral Information System (EcoSIS). 30
- [17] J. Cavender-Bares, J. A. Gamon, and P. A. Townsend. *Remote sensing of plant biodiversity*. Springer Nature, 2020. 1
- [18] J. Cavender-Bares, F. D. Schneider, M. J. Santos, A. Armstrong, A. Carnaval, K. M. Dahlin, L. Fatoyinbo, G. C. Hurtt, D. Schimel, P. A. Townsend, et al. Integrating remote sensing with ecology and evolution to advance biodiversity conservation. *Nature Ecology & Evolution*, 6(5):506–519, 2022. 1
- [19] S. Cerasoli, M. Campagnolo, J. Faria, C. Nogueira, and M. d. C. Caldeira. On estimating the gross primary productivity of mediterranean grasslands under different fertilization regimes using vegetation indices and hyperspectral reflectance. *Biogeosciences*, 15(17):5455–5471, 2018. 4
- [20] K. Chadwick et al. Unlocking ecological insights from sub-seasonal visible-to-shortwave infrared imaging spectroscopy: The shift campaign. *Ecosphere*, 16(3):e70194, 2025. 30
- [21] K. Chadwick, N. Queally, T. Zheng, J. Cryer, C. Vanden Heuvel, C. Villanueva-Weeks, C. Ade, L. Anderegg, Y. Angel, B. Baker, et al. Shift: Photosynthetic and leaf traits, santa barbara county, 2022. *ORNL DAAC*, 2023. 4
- [22] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmLR, 2020. 3
- [23] X. Chen, S. Xie, and K. He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9640–9649, October 2021. 3
- [24] E. Cherif, H. Feilhauer, K. Berger, P. D. Dao, M. Ewald, T. B. Hank, Y. He, K. R. Kovach, B. Lu, P. A. Townsend, et al. From spectra to plant functional traits: Transferable multi-trait models from heterogeneous and sparse data. *Remote Sensing of Environment*, 292:113580, 2023. 2, 3, 4, 5, 32
- [25] A. Chlus, E. L. Kruger, and P. A. Townsend. Mapping three-dimensional variation in leaf mass per area with imaging spectroscopy and lidar in a temperate broadleaf forest. *Remote Sensing of Environment*, 250:112043, 2020. 4
- [26] Y. Cong, S. Khanna, C. Meng, P. Liu, E. Rozi, Y. He, M. Burke, D. Lobell, and S. Ermon. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems*, 35:197–211, 2022. 3
- [27] J. H. Cornelissen, S. Lavorel, E. Garnier, S. Díaz, N. Buchmann, D. Gurvich, P. B. Reich, H. Ter Steege, H. Morgan, M. Van Der Heijden, et al. A handbook of protocols for standardised and easy measurement of plant functional traits worldwide. *Australian journal of Botany*, 51(4):335–380, 2003. 2

- [28] A. Damm, E. Paul-Limoges, E. Haghighi, C. Simmer, F. Morsdorf, F. D. Schneider, C. van der Tol, M. Migliavacca, and U. Rascher. Remote sensing of plant-water relations: An overview and future perspectives. *Journal of plant physiology*, 227:3–19, 2018. 1
- [29] M. F. Danilevicz, M. Gill, R. Anderson, J. Batley, M. Bennamoun, P. E. Bayer, and D. Edwards. Plant genotype to phenotype prediction using machine learning. *Frontiers in Genetics*, 13:822173, 2022. 2
- [30] M. Danner, K. Berger, M. Wocher, W. Mauser, and T. Hank. Efficient rtm-based training of machine learning regression algorithms to quantify biophysical & biochemical traits of agricultural crops. *ISPRS Journal of Photogrammetry and Remote sensing*, 173:278–296, 2021. 2
- [31] P. D. Dao, A. Axiotis, and Y. He. Mapping native and invasive grassland species and characterizing topography-driven species dynamics using high spatial resolution hyperspectral imagery. *International Journal of Applied Earth Observation and Geoinformation*, 104:102542, 2021. 4
- [32] P. D. Dao, Y. He, B. Lu, and A. Axiotis. Imaging spectroscopy reveals topographic variability effects on grassland functional traits and drought responses. *Ecology*, 106(3):e70006, 2025. 4
- [33] P. Dennison. Range creek utah species spectra. <http://ecosis.org>, 2018. Ecological Spectral Information System (EcoSIS). 30
- [34] P. Dennison. Rush valley utah sagebrush time series. <http://ecosis.org>, 2018. Ecological Spectral Information System (EcoSIS). 30
- [35] P. Dennison, C. Daughtry, M. Quemada, K. Roth, I. Numata, S. Meerdink, E. Wetherley, P. Gader, and D. Roberts. Fractional cover simulated vswir dataset version 2, noise & atmos. correction artifacts included. <http://ecosis.org>, 2019. Ecological Spectral Information System (EcoSIS). 30
- [36] P. Dennison and M. Gardner. Hawaii 2000 vegetation species spectra. <http://ecosis.org>, 2018. Ecological Spectral Information System (EcoSIS). 30
- [37] P. Dennison and D. Roberts. Santa monica mountains vegetation species spectra. <http://ecosis.org>, 2018. Ecological Spectral Information System (EcoSIS). 30
- [38] S. Díaz, S. Demissew, J. Carabias, C. Joly, M. Lonsdale, N. Ash, A. Larigauderie, J. R. Adhikari, S. Arico, A. Báldi, et al. The ipbes conceptual framework—connecting nature and people. *Current opinion in environmental sustainability*, 14:1–16, 2015. 2
- [39] M. Ewald, S. Skowronek, R. Aerts, K. Dolos, J. Lenoir, M. Nicolas, J. Warrie, T. Hattab, H. Feilhauer, O. Honnay, et al. Analyzing remotely sensed structural and chemical canopy traits of a forest invaded by prunus serotina over multiple spatial scales. *Biological Invasions*, 20:2257–2271, 2018. 4
- [40] M. Ewald, S. Skowronek, R. Aerts, J. Lenoir, H. Feilhauer, R. Van De Kerchove, O. Honnay, B. Somers, C. X. Garzón-López, D. Rocchini, et al. Assessing the impact of an invasive bryophyte on plant species richness using high resolution imaging spectroscopy. *Ecological Indicators*, 110:105882, 2020. 4
- [41] Y. Fang and L. Tian. The hyperspectra dataset for typical aquatic vegetation. <http://ecosis.org>, 2023. Ecological Spectral Information System (EcoSIS). 30
- [42] H. Feilhauer, G. P. Asner, R. E. Martin, and S. Schmidtlein. Brightness-normalized partial least squares regression for hyperspectral data. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 111(12-13):1947–1957, 2010. 2
- [43] J.-B. Féret, K. Berger, F. De Boissieu, and Z. Malenovsky. Prospect-pro for estimating content of nitrogen-containing leaf proteins and other carbon-based constituents. *Remote Sensing of Environment*, 252:112173, 2021. 6, 37

- [44] J.-B. Feret, C. François, G. P. Asner, A. A. Gitelson, R. E. Martin, L. P. Bidel, S. L. Ustin, G. Le Maire, and S. Jacquemoud. Prospect-4 and 5: Advances in the leaf optical properties model separating photosynthetic pigments. *Remote sensing of environment*, 112(6):3030–3043, 2008. 2
- [45] M. H. P. Fuchs and B. Demir. Hyspecnet-11k: A large-scale hyperspectral dataset for benchmarking learning-based hyperspectral image compression methods. In *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*, pages 1779–1782. IEEE, 2023. 3, 30
- [46] A. Fuller, K. Millard, and J. Green. Croma: Remote sensing representations with contrastive radar-optical masked autoencoders. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 5506–5538. Curran Associates, Inc., 2023. 3
- [47] A. Fuller, K. Millard, and J. Green. Croma: Remote sensing representations with contrastive radar-optical masked autoencoders. *Advances in Neural Information Processing Systems*, 36:5506–5538, 2023. 3
- [48] J. L. Funk, J. E. Larson, G. M. Ames, B. J. Butterfield, J. Cavender-Bares, J. Firn, D. C. Laughlin, A. E. Sutton-Grier, L. Williams, and J. Wright. Revisiting the holy grail: using plant functional traits to understand ecological processes. *Biological Reviews*, 92(2):1156–1173, 2017. 1
- [49] I. Gallo, M. Boschetti, A. U. Rehman, and G. Candiani. Self-supervised convolutional neural network learning in a hybrid approach framework to estimate chlorophyll and nitrogen content of maize from hyperspectral images. *Remote Sensing*, 15(19):4765, 2023. 3
- [50] P. Geladi and B. R. Kowalski. Partial least-squares regression: a tutorial. *Analytica chimica acta*, 185:1–17, 1986. 2
- [51] A. Gravel and E. Laliberte. Mapping canopy foliar functional traits in a mixed temperate forest using imaging spectroscopy. *bioRxiv*, pages 2024–11, 2024. 4
- [52] T. Hank, M. Locherer, K. Richter, W. Mauser, E. Consortium, et al. Neusling (landau ad isar) 2012-a multitemporal and multisensoral agricultural enmap preparatory flight campaign. 2016. 4
- [53] T. Hank, K. Richter, W. Mauser, E. Consortium, et al. Neusling (landau ad isar) 2009-an agricultural enmap preparatory flight campaign using the hmap instrument. 2015. 4
- [54] D. He, D. Kuhn, and L. Parida. Novel applications of multitask learning and multiple output regression to multiple genetic trait prediction. *Bioinformatics*, 32(12):i37–i43, 2016. 2
- [55] K. He, H. Fan, Y. Wu, S. Xie, and R. Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3
- [56] Z. He, H. Liu, Y. Wang, and J. Hu. Generative adversarial networks-based semi-supervised learning for hyperspectral image classification. *Remote Sensing*, 9(10):1042, 2017. 3
- [57] K. Helsen, L. Bassi, H. Feilhauer, T. Kattenborn, H. Matsushima, E. Van Cleemput, B. Somers, and O. Honnay. Evaluating different methods for retrieving intraspecific leaf trait variation from hyperspectral leaf reflectance. *Ecological Indicators*, 130:108111, 2021. 2
- [58] M. Herold, D. Roberts, M. Gardner, and P. Dennison. Urban reflectance spectra from santa barbara, ca. <http://ecosis.org>, 2004. Ecological Spectral Information System (EcoSIS). 30
- [59] I. Herrmann, A. Pimstein, A. Karnieli, Y. Cohen, V. Alchanatis, and D. J. Bonfil. Lai assessment of wheat and potato crops by ven μ s and sentinel-2 bands. *Remote Sensing of Environment*, 115(8):2141–2151, 2011. 4

- [60] D. Hong, Z. Han, J. Yao, L. Gao, B. Zhang, A. Plaza, and J. Chanussot. Spectralformer: Rethinking hyperspectral image classification with transformers. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–15, 2021. 3
- [61] J. Hu, R. Liu, D. Hong, A. Camero, J. Yao, M. Schneider, F. Kurz, K. Segl, and X. X. Zhu. Mdas: A new multimodal benchmark dataset for remote sensing. *Earth System Science Data*, 15(1):113–131, 2023. 30
- [62] B. IPBES et al. Global assessment report on biodiversity and ecosystem services of the intergovernmental science-policy platform on biodiversity and ecosystem services, 2019. 2
- [63] S. Jacquemoud and F. Baret. Prospect: A model of leaf optical properties spectra. *Remote sensing of environment*, 34(2):75–91, 1990. 6, 37
- [64] S. Jacquemoud and S. Ustin. *Leaf optical properties*. Cambridge University Press, 2019. 2
- [65] W. Jetz, J. Cavender-Bares, R. Pavlick, D. Schimel, F. W. Davis, G. P. Asner, R. Guralnick, J. Kattge, A. M. Latimer, P. Moorcroft, et al. Monitoring plant functional diversity from space. *Nature plants*, 2(3):1–5, 2016. 2
- [66] T. Kattenborn, F. E. Fassnacht, and S. Schmidtlein. Differentiating plant functional types using reflectance: which traits make the difference? *Remote Sensing in Ecology and Conservation*, 5(1):5–19, 2019. 4
- [67] T. Kattenborn, F. Schiefer, P. Zarco-Tejada, and S. Schmidtlein. Advantages of retrieving pigment content [$\mu\text{g}/\text{cm}^2$] versus concentration [%] from canopy reflectance. *Remote Sensing of Environment*, 230:111195, 2019. 4
- [68] R. Kokaly et al. Usgs spectral library version 7, no. 1035. Technical report, 2017. 30
- [69] T. Kwak and Y. Kim. Semi-supervised land cover classification of remote sensing imagery using cyclegan and efficientnet. *KSCE Journal of Civil Engineering*, 27(4):1760–1773, 2023. 3
- [70] D. Mederer, H. Feilhauer, E. Cherif, K. Berger, T. B. Hank, K. R. Kovach, P. D. Dao, B. Lu, P. A. Townsend, and T. Kattenborn. Plant trait retrieval from hyperspectral data: Collective efforts in scientific data curation outperform simulated data derived from the prosail model. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, 15:100080, 2025. 2, 3, 5
- [71] G. Olmschenk, Z. Zhu, and H. Tang. Generalizing semi-supervised generative adversarial networks to regression using feature contrasting. *Computer Vision and Image Understanding*, 186:1–12, 2019. 5, 33
- [72] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. HAZIZA, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski. DINOv2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2024. Featured Certification. 3
- [73] L. Parigi, G. Candiani, I. Gallo, P. Toscano, and M. Boschetti. Towards crop traits estimation from hyperspectral data: evaluation of neural network models trained with real multi-site data or synthetic rtm simulations. In *2024 19th Conference on Computer Science and Intelligence Systems (FedCSIS)*, pages 475–484. IEEE, 2024. 2, 3
- [74] T. Paz-Kagan, T. Caras, I. Herrmann, M. Shachak, and A. Karnieli. Multiscale mapping of species diversity under changed land-use imaging spectroscopy. <http://ecosis.org>, 2015. Ecological Spectral Information System (EcoSIS). 30
- [75] N. Pettorelli, N. A. Graham, N. Seddon, M. Maria da Cunha Bustamante, M. J. Lowton, W. J. Sutherland, H. J. Koldewey, H. C. Prentice, and J. Barlow. Time to integrate global climate change and biodiversity science-policy agendas. *Journal of Applied Ecology*, 58(11):2384–2393, 2021. 1

- [76] M. Pichler, V. Boreux, A.-M. Klein, M. Schleuning, and F. Hartig. Machine learning algorithms to infer trait-matching and predict species interactions in ecological networks. *Methods in Ecology and Evolution*, 11(2):281–293, 2020. 2
- [77] J. Pottier, Z. Malenovský, A. Psomas, L. Homolová, M. E. Schaepman, P. Choler, W. Thuiller, A. Guisan, and N. E. Zimmermann. Modelling plant species distribution in alpine grasslands using airborne imaging spectroscopy. *Biology letters*, 10(7):20140347, 2014. 4
- [78] R. Pullanagari, M. Dehghan-Shoar, I. J. Yule, and N. Bhatia. Field spectroscopy of canopy nitrogen concentration in temperate grasslands using a convolutional neural network. *Remote Sensing of Environment*, 257:112353, 2021. 3
- [79] N. Queally. California vegetation species image spectra. <http://ecosis.org>, 2018. Ecological Spectral Information System (EcoSIS). 30
- [80] P. Ranjan, R. Kumar, and A. Girdhar. Unlocking the potential of unlabeled data: Semi-supervised learning for stratification of hyperspectral images. In *2023 OITS International Conference on Information Technology (OCIT)*, pages 938–943. IEEE, 2023. 3
- [81] C. J. Reed, R. Gupta, S. Li, S. Brockman, C. Funk, B. Clipp, K. Keutzer, S. Candido, M. Uyttendaele, and T. Darrell. Scale-mae: A scale-aware masked autoencoder for multiscale geospatial representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4088–4099, 2023. 3
- [82] A. Rogers, S. Serbin, and K. Ely. Leaf mass area, leaf carbon and nitrogen content, kougark road and teller road, seward peninsula, alaska, 2016. Technical report, Next Generation Ecosystems Experiment-Arctic, Oak Ridge National Laboratory . . . , 2019. 4
- [83] C. Roscher, J. Schumacher, M. Gubsch, A. Lipowsky, A. Weigelt, N. Buchmann, B. Schmid, and E.-D. Schulze. Using plant functional traits to explain diversity–productivity relationships. *PloS one*, 7(5):e36760, 2012. 1
- [84] L. Roupioz et al. Multi-source datasets acquired over toulouse (france) in 2021 for urban microclimate studies during the camcatt/ai4geo field campaign. *Data in Brief*, 48:109109, 2023. 30
- [85] S. K. Roy, J. M. Haut, M. E. Paoletti, S. R. Dubey, and A. Plaza. Generative adversarial minority oversampling for spectral–spatial hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–15, 2021. 3
- [86] A. Savitzky and M. J. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, 1964. 28
- [87] L. Scheibenreif, M. Mommert, and D. Borth. Masked vision transformers for hyperspectral image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2166–2176, 2023. 3
- [88] M. Schlerf and C. Atzberger. Inversion of a forest reflectance model to estimate structural canopy variables from hyperspectral remote sensing data. *Remote sensing of environment*, 100(3):281–294, 2006. 3
- [89] M. Scutari, P. Howell, D. J. Balding, and I. Mackay. Multiple quantitative trait analysis using bayesian networks. *Genetics*, 198(1):129–137, 2014. 5
- [90] S. Serbin, S. DuBois, A. Jablonski, A. Shiklomanov, A. Desai, E. Kruger, and P. Townsend. Uw-bnl nasa hypsiri airborne campaign leaf and canopy spectra and trait data. <http://ecosis.org>, 2019. Ecological Spectral Information System (EcoSIS). 30
- [91] S. Serbin, W. Lieberman-Cribbin, K. Ely, and A. Rogers. Ngee arctic hr1024i canopy spectral reflectance, drained thaw lake basin, barrow, alaska, july 2015. <http://ecosis.org>, 2015. Ecological Spectral Information System (EcoSIS). 30

- [92] S. P. Serbin, A. Singh, A. R. Desai, S. G. Dubois, A. D. Jablonski, C. C. Kingdon, E. L. Kruger, and P. A. Townsend. Remotely estimating photosynthetic capacity, and its response to temperature, in vegetation canopies using imaging spectroscopy. *Remote Sensing of Environment*, 167:78–87, 2015. 2, 3
- [93] S. P. Serbin, J. Wu, K. S. Ely, E. L. Kruger, P. A. Townsend, R. Meng, B. T. Wolfe, A. Chlus, Z. Wang, and A. Rogers. From the arctic to the tropics: multibiome prediction of leaf mass per area using leaf reflectance. *New Phytologist*, 224(4):1557–1568, 2019. 2
- [94] Y. She, C. Atzberger, A. Blake, and S. Keshav. From spectra to biophysical insights: End-to-end learning with a biased radiative transfer model. *arXiv preprint arXiv:2403.02922*, 2024. 6
- [95] S. Shi, L. Xu, W. Gong, B. Chen, B. Chen, F. Qu, X. Tang, J. Sun, and J. Yang. A convolution neural network for forest leaf chlorophyll and carotenoid estimation using hyperspectral reflectance. *International Journal of Applied Earth Observation and Geoinformation*, 108:102719, 2022. 3
- [96] A. Singh, S. P. Serbin, B. E. McNeil, C. C. Kingdon, and P. A. Townsend. Imaging spectroscopy algorithms for mapping canopy foliar chemical and morphological traits and their uncertainties. *Ecological Applications*, 25(8):2180–2197, 2015. 3, 4
- [97] A. K. Skidmore, N. C. Coops, E. Neinavaz, A. Ali, M. E. Schaepman, M. Paganini, W. D. Kissling, P. Vihervaara, R. Darvishzadeh, H. Feilhauer, et al. Priority list of biodiversity metrics to observe from space. *Nature ecology & evolution*, 5(7):896–906, 2021. 1
- [98] D. Szwarcman, S. Roy, P. Fraccaro, Þorsteinn Elí Gíslason, B. Blumenstiel, R. Ghosal, P. H. de Oliveira, J. L. de Sousa Almeida, R. Sedona, Y. Kang, S. Chakraborty, S. Wang, C. Gomes, A. Kumar, M. Truong, D. Godwin, H. Lee, C.-Y. Hsu, A. A. Asanjan, B. Mujeci, D. Shidham, T. Keenan, P. Arevalo, W. Li, H. Alemohammad, P. Olofsson, C. Hain, R. Kennedy, B. Zadrozny, D. Bell, G. Cavallaro, C. Watson, M. Maskey, R. Ramachandran, and J. B. Moreno. Prithvi-eo-2.0: A versatile multi-temporal foundation model for earth observation applications, 2025. 3
- [99] G. Tagliabue, M. Boschetti, G. Bramati, G. Candiani, R. Colombo, F. Nutini, L. Pompilio, J. P. Rivera-Caicedo, M. Rossi, M. Rossini, et al. Hybrid retrieval of crop traits from multi-temporal prisma hyperspectral imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 187:362–377, 2022. 2, 3
- [100] M. Tan and Q. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 5
- [101] D. Thompson. Emit manually adjusted surface reflectance spectra. <http://ecosis.org>, 2021. Ecological Spectral Information System (EcoSIS). 30
- [102] D. Thompson. Emit manually adjusted vegetation reflectance spectra. <http://ecosis.org>, 2021. Ecological Spectral Information System (EcoSIS). 30
- [103] R. Thoreau, L. Risser, V. Achard, B. Berthelot, and X. Briottet. Toulouse hyperspectral data set: A benchmark data set to assess semi-supervised spectral representation learning and pixel-wise classification techniques. *ISPRS Journal of Photogrammetry and Remote Sensing*, 212:323–337, 2024. 3, 6
- [104] L. C. Todman, A. Bush, and A. S. Hood. ‘small data’ for big insights in ecology. *Trends in Ecology & Evolution*, 38(7):615–622, 2023. 2
- [105] N. L. Tsakiridis, K. D. Keramaris, J. B. Theocharis, and G. C. Zalidis. Simultaneous prediction of soil properties from vnir-swir spectra using a localized multi-channel 1-d convolutional neural network. *Geoderma*, 367:114208, 2020. 2
- [106] G. Tseng, R. Cartuyvels, I. Zvonkov, M. Purohit, D. Rolnick, and H. Kerner. Lightweight, pre-trained transformers for remote sensing timeseries, 2023. 3

- [107] G. Tseng, A. Fuller, M. Reil, H. Herzog, P. Beukema, F. Bastani, J. R. Green, E. Shelhamer, H. Kerner, and D. Rolnick. Galileo: Learning global and local features in pretrained remote sensing models, 2025. [3](#)
- [108] S. Unger, S. Vargas, J. May, and S. Oberbauer. Arctic moss spectral reflectance desiccation experiment from samples collected in northern alaska. <http://ecosis.org>, 2020. Ecological Spectral Information System (EcoSIS). [30](#)
- [109] S. L. Ustin, A. A. Gitelson, S. Jacquemoud, M. Schaepman, G. P. Asner, J. A. Gamon, and P. Zarco-Tejada. Retrieval of foliar information about plant pigment systems from high resolution spectroscopy. *Remote Sensing of Environment*, 113:S67–S77, 2009. [2](#), [3](#)
- [110] E. Van Cleemput, D. A. Roberts, O. Honnay, and B. Somers. A novel procedure for measuring functional traits of herbaceous species through field spectroscopy. *Methods in Ecology and Evolution*, 10(8):1332–1338, 2019. [4](#)
- [111] E. Van Cleemput, L. Vanierschot, B. Fernández-Castilla, O. Honnay, and B. Somers. The functional characterization of grass-and shrubland ecosystems using hyperspectral remote sensing: trends, accuracy and moderating variables. *Remote Sensing of Environment*, 209:747–763, 2018. [2](#)
- [112] W. Verhoef, L. Jia, Q. Xiao, and Z. Su. Unified optical-thermal four-stream radiative transfer theory for homogeneous vegetation canopies. *IEEE Transactions on geoscience and remote sensing*, 45(6):1808–1822, 2007. [6](#), [37](#)
- [113] J. Verrelst, L. Alonso, J. P. R. Caicedo, J. Moreno, and G. Camps-Valls. Gaussian process retrieval of chlorophyll content from imaging spectroscopy data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 6(2):867–874, 2012. [3](#)
- [114] J. Verrelst, G. Camps-Valls, J. Muñoz-Marí, J. P. Rivera, F. Veroustraete, J. G. Clevers, and J. Moreno. Optical remote sensing and the retrieval of terrestrial vegetation bio-geophysical properties—a review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 108:273–290, 2015. [2](#)
- [115] J. Verrelst, J. P. Rivera, J. Moreno, and G. Camps-Valls. Gaussian processes uncertainty estimates in experimental sentinel-2 lai and leaf chlorophyll content retrieval. *ISPRS journal of photogrammetry and remote sensing*, 86:157–167, 2013. [3](#)
- [116] G. Vivone, A. Garzelli, Y. Xu, W. Liao, and J. Chanussot. Panchromatic and hyperspectral image fusion: Outcome of the 2022 whispers hyperspectral pansharpening challenge, 2022. [30](#)
- [117] D. Wang, M. Hu, Y. Jin, Y. Miao, J. Yang, Y. Xu, X. Qin, J. Ma, L. Sun, C. Li, et al. Hypersigma: Hyperspectral intelligence comprehension foundation model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. [3](#)
- [118] D. R. Wang, E. J. Wolfrum, P. Virk, A. Ismail, A. J. Greenberg, and S. R. McCouch. Robust phenotyping strategies for evaluation of stem non-structural carbohydrates (nsc) in rice. *Journal of Experimental Botany*, 67(21):6125–6138, 2016. [4](#)
- [119] S. Wang, K. Guan, Z. Wang, E. A. Ainsworth, T. Zheng, P. A. Townsend, N. Liu, E. Nafziger, M. D. Masters, K. Li, et al. Airborne hyperspectral imaging of nitrogen deficiency on crop traits and yield of maize by machine learning and radiative transfer modeling. *International Journal of Applied Earth Observation and Geoinformation*, 105:102617, 2021. [2](#)
- [120] Y. Wang, C. M. Albrecht, N. A. A. Braham, L. Mou, and X. X. Zhu. Self-supervised learning in remote sensing: A review. *IEEE Geoscience and Remote Sensing Magazine*, 10(4):213–247, 2022. [3](#)
- [121] Z. Wang, A. Chlus, R. Geygan, Z. Ye, T. Zheng, A. Singh, J. J. Couture, J. Cavender-Bares, E. L. Kruger, and P. A. Townsend. Foliar functional traits from imaging spectroscopy across biomes in eastern north america. *New Phytologist*, 228(2):494–511, 2020. [3](#), [4](#)
- [122] B. Weinstein, S. Marconi, and E. White. Data for the neontreeevaluation benchmark (0.2.2). <https://doi.org/10.5281/zenodo.5914554>, 2022. Zenodo. [30](#)

- [123] M. Woher, K. Berger, M. Danner, W. Mauser, and T. Hank. Physically-based retrieval of canopy equivalent water thickness using hyperspectral data. *Remote Sensing*, 10(12):1924, 2018. 4
- [124] Z. Xiong, Y. Wang, F. Zhang, A. J. Stewart, J. Hanna, D. Borth, I. Papoutsis, B. L. Saux, G. Camps-Valls, and X. X. Zhu. Neural plasticity-inspired multimodal foundation model for earth observation, 2024. 3
- [125] H. Xu, Y. Cao, D. Yu, M. Cao, Y. He, M. Gill, and H. M. Pereira. Ensuring effective implementation of the post-2020 global biodiversity targets. *Nature Ecology & Evolution*, 5(4):411–418, 2021. 1
- [126] P. Yan, M. Fernández-Martínez, K. Van Meerbeek, G. Yu, M. Migliavacca, and N. He. The essential role of biodiversity in the key axes of ecosystem function. *Global Change Biology*, 29(16):4569–4585, 2023. 1
- [127] P. J. Zarco-Tejada, C. Camino, P. Beck, R. Calderon, A. Hornero, R. Hernández-Clemente, T. Kattenborn, M. Montes-Borrego, L. Susca, M. Morelli, et al. Previsual symptoms of xylella fastidiosa infection revealed in spectral plant-trait alterations. *Nature Plants*, 4(7):432–439, 2018. 1
- [128] P. J. Zarco-Tejada, A. Hornero, P. Beck, T. Kattenborn, P. Kempeneers, and R. Hernández-Clemente. Chlorophyll content estimation in an open-canopy conifer forest with sentinel-2a and hyperspectral imagery in the context of forest decline. *Remote sensing of environment*, 223:320–335, 2019. 1
- [129] S. Zesati, C. Tweedie, K. Huemmrich, P. Campbell, and M. Velez-Reyes. Above: Reflectance spectra of tundra plant communities across northern alaska. <https://doi.org/10.3334/ORNLDAAC/1685>, 2019. ORNL DAAC, Oak Ridge, Tennessee, USA. 30
- [130] Y. Zhan, K. Wu, W. Liu, J. Qin, Z. Yang, Y. Medjadba, G. Wang, and X. Yu. Semi-supervised classification of hyperspectral data based on generative adversarial networks and neighborhood majority voting. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 5756–5759. IEEE, 2018. 3
- [131] C. Zhao, X. Wan, G. Zhao, B. Cui, W. Liu, and B. Qi. Spectral-spatial classification of hyperspectral imagery based on stacked sparse autoencoder and random forest. *European journal of remote sensing*, 50(1):47–63, 2017. 3
- [132] R. Zhao, L. An, W. Tang, L. Qiao, N. Wang, M. Li, H. Sun, and G. Liu. Improving chlorophyll content detection to suit maize dynamic growth effects by deep features of hyperspectral data. *Field Crops Research*, 297:108929, 2023. 3
- [133] T. Zheng, Z. Ye, A. Singh, A. R. Desai, N. Krishnayya, M. G. Dave, and P. A. Townsend. Variability in forest plant traits along the western ghats of india and their environmental drivers at different resolutions. *Journal of Geophysical Research: Biogeosciences*, 129(3):e2023JG007753, 2024. 4
- [134] L. Zhu, Y. Chen, P. Ghamisi, and J. A. Benediktsson. Generative adversarial networks for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 56(9):5046–5063, 2018. 3

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: we summarize four key contributions in Section 1, which are addressed through our experimental design and results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See Section 7

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: This work does not include theoretical analysis.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We share the data and code and provide more details on the experiments in Appendix [B](#).

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We submit the code and data URL on the OpenReview platform.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Key training details, including model architecture, loss functions, data splits and major modifications, are summarized in the main text. Comprehensive descriptions of training procedures and hyperparameters are provided in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We evaluate all models using three random seeds and report the mean and standard deviation of the performance metrics (*i.e.* R^2 , nRMSE) to reflect variability due to random initialization. This is stated in the main text and reflected in relevant figures and tables (*e.g.* Figures 5 and 4). For the cross-dataset OOD evaluation, we report results from a single run due to computational constraints; however, we attempt to account for variability through controlled random sampling, as detailed in Section 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide detailed information on compute resources in Appendix B, including the type of GPUs, model-specific runtime estimates, and the number of trainable parameters per method (see Table 19).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have reviewed the NeurIPS ethical guidelines and ensured that the paper complies with them.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the positive societal impacts of our work in Section 1 and 7. While we did not identify any clear negative impacts, we remain open to addressing potential concerns if raised during the review process.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The data and models presented in this paper do not pose risks of misuse. All datasets are derived from publicly available, non-sensitive sources, and the models are focused on ecological applications.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All datasets used in this work are publicly available and properly credited in the paper. For each source included in GreenHyperSpectra, we provide citation information, access links in Table 7.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.

- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: The GreenHyperSpectra dataset introduced in this paper is documented in detail in Appendix A. The appendix includes information about data sources, sensor specifications, geographic and temporal coverage, preprocessing steps, and licensing terms for each included dataset. The asset is composed entirely of publicly available data, contains no personal or sensitive information, and respects the original licenses of the sources.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: The paper does not involve any crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core methods and experiments presented in this paper do not involve the use of large language models (LLMs).

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

Table of figures:

Figure	Description
1	Teaser illustration of the proposed semi-/self-supervised frameworks for multi-trait regression.
2	Comparison of GreenHyperSpectra and the labeled dataset, highlighting broader coverage in vegetation types and sensor diversity.
3a, 3b and 3c	Overview of the SR-GAN, RTM-AE, and MAE architectures for trait prediction.
4	Effect of increasing labeled data volume on R^2 performance.
5	Effect of increasing unlabeled pretraining data volume on R^2 performance.
6, 7,8 and 9	Dataset characteristics, focusing on spectral variability.
10 and 11	Complementary results to Figs. 4 and 5: nRMSE performance trends with increasing labeled and unlabeled data.
12 and 13	Complementary results to Table 3: Heatmaps of R^2 and nRMSE across traits in the half-range input settings to show the performance of MAE vs baseline.
14	Complementary results to Table 4: Observed vs. predicted plots, showing trait-wise calibration in the OOD setting.
15	Feature importance of MAE-based downstream regression as a function of fine-tuning depth (linear probing, final block, and full fine-tuning).

Table 5: Summary of figures and their descriptions.

Table of tables:

Table	Description
1	Summary of sensor and platform specifications in GreenHyperSpectra.
2	Trait-wise performance (R^2 and nRMSE) of all models under full-range input settings.
3	Trait-wise performance (R^2 and nRMSE) of all models under half-range (VNIR) input settings.
4	Trait-wise performance (R^2 and nRMSE) of all models under out-of-distribution (OOD) settings.
7 and 8	Dataset details: spectral data characteristics and trait distribution across sources.
9 and 10	Architecture and hyperparameters of SR-GAN.
11, 12 and 13	Architecture, hyperparameters, and RTM configuration of RTM-AE.
14 and 15	Architecture and hyperparameters of MAE.
16, 17 and 18	MAE ablation studies: effects of transformer depth, loss weighting, and token size on trait prediction (R^2 and nRMSE).
19	Model size, runtime, and GPU usage across methods.
20, 21, 22, 23, 24 and 25	Complement to Table 4 and Fig. 14: OOD model performance when one vegetation class is excluded from the test set.
26, 27, 28, 29 and 30	Robustness evaluation under additive Gaussian noise during inference, reported for all models (R^2 and nRMSE).

Table 6: Summary of tables and their descriptions.

A Details about the datasets

The data are publicly available [here](#).

Spectral preprocessing. For standardized cross-instrument comparison, all reflectance spectra were resampled to a uniform wavelength grid spanning the 400–2500 nm solar-reflective range. Spectral measurements were linearly interpolated to an interval of 1 nm, resulting in 2101 bands per sample. Regions of strong atmospheric water absorption, specifically 1351–1430 nm, 1801–2050 nm, and 2451–2500 nm, were removed to minimize noise and signal loss. The remaining bands were smoothed using a Savitzky-Golay filter with a 65 nm window [86]. After these steps, 1721 spectral bands were retained for analysis, providing a high-quality input space for training and evaluation.




Figure 6: Spectral reflectance across wavelengths. This plot shows the variation in canopy reflectance within GreenHyperSpectra across different data sources, highlighting differences due to acquisition conditions and sensor modalities. The colored ranges refer to the visible region.




Figure 7: **Sample distribution across vegetation types in GreenHyperSpectra.** The plot shows the number of samples in GreenHyperSpectra (blue) and the existing labeled (purple) for each vegetation type, highlighting class imbalance and the relative scarcity of labeled data in certain categories. The vegetation type information was retrieved from the [ESA WORLDCOVER](#) product

Dataset	Platform	Sensor	GSD	#Bands	Range (nm)	Year	#Samples	Processing	Land Cover	Source
DB1[33]	proximal	ASD FieldSpec Pro	N/A	2151	350–2500	2007 2009	31 49	Reflectance spectra	desert, shrubland	Link
DB2[36]	proximal	ASD FieldSpec FR	N/A	1063	352–2476	2000	792	Reflectance spectra	forest, shrubland: native-dominated Hawaiian forest types	Link
DB3[37]	proximal	ASD FieldSpec FR	N/A	1075	350–2498	1995 1996 1997 1998	226 93 132 10	Reflectance spectra	shrubland	Link
DB4[90]	proximal	SpecEvo PSM3500 ASD Fieldspec 3 ASD Fieldspec 4 SpecEvo_PSM3500	N/A	2151	350–2500	2013 2013 2013 2014	6 7 49 60	Reflectance spectra	Forest, Gra4Sand, Shrubland, Crops	Link
DB5[129]	proximal	SVC HR-1024i	1 m	2178	338–2516	2018	112	Reflectance spectra	tundra	Link
DB6[101]	proximal	ASD FieldSpec 3	N/A	2151	350–2500	2001	112	Reflectance spectra	urban vegetation	Link
DB7[35]	proximal	ASD FieldSpec Pro	N/A	210	400–2490		715	Reflectance spectra	crops, gra4Sand, forest, shrubland	Link
DB8[58]	proximal	ASD FieldSpec 3	N/A	1075	350–2498	2001	37	Reflectance spectra	urban vegetation	Link
DB9[68]	proximal	ASD FieldSpec 4 Hi-Res NG	N/A	2151	350–2500		87	Reflectance spectra	–	Link
DB10[68]	proximal	Beckman 5270	N/A	480	205–2976		19	Reflectance spectra	–	Link
DB11[41]	proximal	SVC HR-1024i	40 cm	995	346–2499		133	Reflectance spectra	aquatic vegetation	Link
DB12[108]	proximal	SVC HR-1024i	8 cm	994	338–2515		34	Reflectance spectra	tundra	Link
DB13[91]	proximal	SVC HR-1024i	N/A	2150	350–2500	2015	44	Reflectance spectra	coastal, wetland	Link
DB14[84]	proximal	ASD	N/A	2101	400–2500	2021	45	Reflectance spectra	urban vegetation	Link
DB15[34]	proximal	ADS FieldSpec Pro	N/A	2151	350–2500	2005	82	Reflectance spectra	shrubland, steppe	Link
DB16[129]	proximal	SVC HR-1024i	1 m	994	338–2516	2017	1660	Reflectance spectra	tundra	Link
DB17[102]	proximal	ASD FieldSpec 3	N/A	2151	350–2500	2001	490	Reflectance spectra	coastal, forest, shrubland	Link
DB18[79]	airborne	AVIRIS Classic	17–20 m	244	365–2496	2013 2014 2016	341 37 31	ACSR	Urban, chaparral, oak woodland, conifer forest	Link
DB19[68]	airborne	AVIRIS	17–20 m	224	365–2496		32		–	Link
DB20[74]	airborne	AisaFenix	1 m	360	400–2400	2014	22889	ACSR	Forest, Ecology, Land Cover, Agriculture	Link
DB21[16]	spaceborne	Hyperion	30 m	220	400–2500	2008 2009	25	ACSR	forest	Link
DB22[122]	airborne	NEON AOP	1 m	426	380–2510	2018	10322	ACSR	–	Link
DB23[61]	airborne	Hypex	–	368	417–2484	2018	9993	ACSR	crops	Link
DB24[122]	airborne	NEON AOP	1 m	426	380–2510	2019	16373	ACSR	–	Link
DB25[116]	spaceborne	PRISMA	30 m	69	400–2500	2021	2155	ACSR	–	Link
DB26[116]	spaceborne	PRISMA	30 m	63	400–2500	2021	10000	ACSR	–	–
DB27[45]	spaceborne	EnMAP	30 m	224	418–2445	2022	1890	ACSR	–	Link
DB28	airborne	NEON AOP	1 m	426	380–2510	2022	3959	ACSR	–	Link
DB29[84]	airborne	AisaFenix	1 m	420	382–2499	2021	31811	ACSR	urban vegetation	Link
DB30[20]	airborne	Aviris NG	20 m	425	380–2510	2022	911	ACSR	Mediterranean ecosystem	Link
DB31	spaceborne	EMIT	60 m	285	381–2492	2024	410	ACSR	–	Link ²
DB32	spaceborne	EnMAP	30 m	224	418–2445	2022– 2024	6653	ACSR	temperate forest	EnMAP ¹
DB33	spaceborne	EnMAP	30 m	224	418–2445	2022– 2024	1655	ACSR	Mediterranean ecosystem	EnMAP ¹
DB34	spaceborne	EnMAP	30 m	224	418–2445	2022– 2024	2088	ACSR	temperate gra4Sand	EnMAP ¹
DB35	spaceborne	EnMAP	30 m	224	418–2445	2022– 2024	4846	ACSR	tropical forest	EnMAP ¹
DB36	spaceborne	EnMAP	30 m	224	418–2445	2022– 2024	6337	ACSR	tropical savanna	EnMAP ¹

Table 7: Summary of data sources of GreenHyperSpectra. Technical detailed on the collected spectra and their corresponding sources. GSD = Ground Sampling Distance. ACSR = Atmospheri-




Figure 8: **t-SNE projection of reflectance spectra.** Each subplot shows the projected spectral signatures from a specific data source in GreenHyperSpectra, illustrating variability driven by differences in sensors, biomes, or acquisition conditions. For each source, spectra from GreenHyperSpectra are compared to those in the aggregated annotated dataset. Orange points represent labeled spectra, and blue points denote unlabeled samples.




Figure 9: **Spatial distribution of a subset from GreenHyperSpectra vs annotated data.** Points represent sample locations of the existing annotated data (left) and the GreenHyperSpectra subset (right). This subset, comprising 80,000 samples, was selected from the full dataset to ensure broad coverage of geographic regions and acquisition conditions. It is used for sample sensitivity analysis to assess the impact of data quantity on model performance, while enabling computationally efficient experimentation.

¹Contains modified EnMAP data © DLR [2024]. See [EnMAP Portal](#).

²Contains modified EMIT data, the original data used is licensed under the [Apache License, Version 2.0](#).

Table 8: **Descriptive statistics for plant traits in the aggregated annotated dataset[24]**. List of traits and units: leaf mass per area (g/cm^2) = Cm, leaf protein content (g/cm^2) = Cp, equivalent water thickness (cm) = Cw, leaf total chlorophyll content ($\mu\text{g}/\text{cm}^2$) = Cab, leaf carotenoid content ($\mu\text{g}/\text{cm}^2$) = Car, leaf anthocyanin content ($\mu\text{g}/\text{cm}^2$) = Anth, Leaf Area Index (m^2/m^2) = LAI and carbon-based constituents (g/cm^2) = cbc (Cm-Cp).

Trait	Count	Mean	Std	Min	25%	50%	75%	Max
cab	2593	39.1234	14.2312	4.4483	28.2500	38.0042	49.0675	229.4975
cw	2782	0.0160	0.0166	0.0000	0.0096	0.0130	0.0184	0.5138
cm	4062	0.0101	0.0083	0.0000	0.0051	0.0080	0.0117	0.0682
LAI	1656	3.4927	1.7178	0.0633	2.1944	3.4691	4.7743	8.7700
cp	3031	0.0009	0.0005	0.0000	0.0006	0.0008	0.0011	0.0050
cbc	3031	0.0104	0.0086	0.0000	0.0056	0.0078	0.0123	0.0671
car	1873	8.6925	2.8232	1.1826	6.9679	8.5176	10.2998	40.4432
anth	644	1.2730	0.4095	0.5610	0.9491	1.2345	1.5226	2.9811

B Details about Models

The code for accessing the dataset and benchmarking experiments can be found [here](#). The trained model objects are also available [here](#).

B.1 Semi-supervised regression generative adversarial network (SR-GAN)

To enable trait prediction from unlabeled hyperspectral spectra, we adopt a semi-supervised GAN framework [71]. The generator G learns to synthesize spectra that are indistinguishable from real data by optimizing $\mathcal{L}_{\text{gen}}$, while the discriminator D learns both to distinguish real from fake spectra and to regress plant trait values from real labeled data optimizing $\mathcal{L}_{\text{disc}}$. We adopt this notation for D outputs: f is an intermediate feature extractor and $D \circ f$ is final layer trait prediction. $\text{Dist}(\cdot, \cdot)$ denotes a distance metric (e.g., cosine or Euclidean).

Notations: x_{fake} : generated fake spectra from the generator; x_{unlb} : unlabeled sample from GreenHyperSpectra; x_{lb} : spectra sample from the labeled data;

Generator Matching Loss. The generator is trained to align the generated spectra with real spectra in the latent feature space:

$$\mathcal{L}_{\text{gen}} = \lambda_{\text{gen}} \cdot \text{Dist}(f(x_{\text{fake}}), f(x_{\text{unlabeled}})), \quad (1)$$

where λ_{gen} controls the influence of the generator loss.

Labeled Supervised Loss. For labeled spectra x_{labeled} with corresponding trait references y , we define a standard supervised regression loss:

$$\mathcal{L}_{\text{labeled}} = \lambda_{\text{labeled}} \cdot \text{MSE}(D \circ f(x_{\text{labeled}}), y), \quad (2)$$

where λ_{labeled} is a weighting coefficient, and MSE denotes the mean squared error between predicted and true traits.

Unlabeled Matching Loss. To regularize the feature space, we encourage the feature representations $f(\cdot)$ extracted by D from labeled and unlabeled real spectra to be similar:

$$\mathcal{L}_{\text{unlabeled}} = \lambda_{\text{unlabeled}} \cdot \lambda_{\text{srgan}} \cdot \text{Dist}(f(x_{\text{labeled}}), f(x_{\text{unlabeled}})), \quad (3)$$

where $\lambda_{\text{unlabeled}}$ and λ_{srgan} are scaling factors.

Fake Contrastive Loss. The discriminator is further trained to push away fake spectra $x_{\text{fake}} = G(z)$ from real ones in the feature space:

$$\mathcal{L}_{\text{fake}} = \lambda_{\text{fake}} \cdot \lambda_{\text{srgan}} \cdot \text{Dist}(f(x_{\text{unlabeled}}), f(x_{\text{fake}})), \quad (4)$$

where λ_{fake} weights the contrastive term.

Gradient Penalty. A gradient penalty is used to enforce Lipschitz continuity, stabilizing the training of the discriminator:

$$\mathcal{L}_{\text{GP}} = \lambda_{\text{GP}} \cdot \mathbb{E}_{\hat{x}} \left[\max(0, (\|\nabla_{\hat{x}} f(\hat{x})\|_2 - 1)^2) \right], \quad (5)$$

where $\hat{x}$ is an interpolated sample between x_{fake} and $x_{\text{unlabeled}}$, and λ_{GP} is the penalty coefficient.

Total Losses. The complete losses for the discriminator are defined as:

$$\mathcal{L}_{\text{disc}} = \mathcal{L}_{\text{labeled}} + \mathcal{L}_{\text{unlabeled}} + \mathcal{L}_{\text{fake}} + \mathcal{L}_{\text{GP}}, \quad (6)$$

Table 9: Architectural details of the convolutional GAN model used for spectral generation and trait regression.

Network	Layer	Description
Generator	Input	Latent vector $z \in \mathbb{R}^d$
	Fully Connected	Linear: $d \rightarrow 64 \times \frac{S}{4}$, reshaped to $(64, \frac{S}{4})$
	Transposed Conv1D	$(64, \frac{S}{4}) \rightarrow (64, S)$, kernel=16, stride=4, pad=6
	Residual Stack	Three residual blocks (dilations=1, 3, 9), LeakyReLU, skip connections
	Output Conv1D	Conv1D: $(64, S) \rightarrow (1, S)$, kernel=7, pad=3
	Output Activation	Tanh activation to constrain to $[-1, 1]$
Discriminator	Input	Spectral input $x \in \mathbb{R}^{1 \times S}$
	Conv1D Layer 1	SpectralNorm: $(1, S) \rightarrow (128, S/2)$, kernel=3, stride=2, pad=1
	BatchNorm + Activation	BatchNorm1D + LeakyReLU
	Conv1D Layer 2	SpectralNorm: $(128, S/2) \rightarrow (128, S/4)$
	Conv1D Layer 3 (output)	SpectralNorm: $(128, S/4) \rightarrow (128, S/8)$
	Adaptive Pooling	AdaptiveAvgPool1D (optional)
	Flatten	Flatten to $(128 \times S/8)$
	Dropout	Dropout $p = 0.4$
	Fully Connected (output)	Linear: $128 \times S/8 \rightarrow n_{\text{traits}}$

Table 10: **Training hyperparameters for the SR-GAN.** This table lists the default hyperparameters and optimization settings used during training for both generator and discriminator components.

Parameter	Value	Description
input shape	1720 or 500	Number of spectral input bands
latent dim	100	Generator latent vector size
n_lb	8	Number of predicted plant traits
batch size	128	Samples per batch
n_epochs	300	Total training epochs
learning rate G	1e-4	Generator optimizer learning rate
learning rate D	4*1e-4	Discriminator optimizer learning rate
optimizers	Adam (amsgrad=True)	Optimizer
weight decay	1e-4	L2 regularization
lambda_fk	1.0	Generator adversarial loss weight
lambda_un	10.0	Unsupervised feature loss weight
labeled_loss_multiplier	1.0	Supervised regression loss weight
matching_loss_multiplier	1.0	Real/fake match loss weight
contrasting_loss_multiplier	1.0	Contrastive loss weight
srgan_loss_multiplier	1.0	Contrastive loss weight
gradient penalty on	True	Enable gradient penalty
gradient_penalty_multiplier	10.0	Weight for GP term
augmentation	True	Data augmentation
contrasting_distance_function	CosineEmbeddingLoss	Real/fake separation
matching_distance_function	CosineEmbeddingLoss	Real-real alignment
labeled_loss_function	Huber loss	Regression loss for traits

B.2 Model Architecture Description for RTM-AE

Table 11: **Architecture of the RTM-AE model.**

Module	Layer	Description
Encoder	Input	Spectral input $x \in \mathbb{R}^{1 \times S}$
	Fully Connected 1	$S \rightarrow 64$, followed by LayerNorm and ReLU
	Fully Connected 2	$64 \rightarrow 32$, followed by LayerNorm and ReLU
	Fully Connected 3	$32 \rightarrow 16$, followed by LayerNorm and ReLU
	Trait Output Layer	$16 \rightarrow n_{\text{traits}}$
RTM Decoder	Non-learnable Module	PROSAIL-PRO: $n_{\text{traits}} \rightarrow \tilde{x} \in \mathbb{R}^{1 \times 2101}$
Correction Block	Fully Connected 1	$2101 \rightarrow 8404$, followed by ReLU
	Fully Connected 2	$8404 \rightarrow \hat{x} \in \mathbb{R}^{1 \times 2101}$
	Output	Corrected reflectance spectrum in $\mathbb{R}^S$

Table 12: **Training hyperparameters for the RTM-AE.** This table lists the default hyperparameters and optimization settings used during training.

Parameter	Value	Description
input shape	1720 or 500	Number of spectral input bands
latent dimension	8	Number of biophysical traits (latent features)
output spectrum length	2101	Number of bands in RTM-simulated output
batch size	128	Number of samples per training batch
training epochs	300	Set during experimental runs
learning rate	1e-4	Learning rate used for the Adam optimizer
weight decay	1e-4	L2 regularization term
optimizer	Adam (amsgrad=True)	Optimizer used
reconstruction loss	Cosine similarity + MAE	Match predicted vs. input spectra
label loss	Huber loss	Trait prediction loss on labeled samples
gradient stabilization	Enabled	Replace gradients when $\ \nabla\ < 10^{-5}$
RTM decoder	PROSAIL-PRO	Fixed physics-based decoder
leaf model	PROSPECTPro	Leaf optical model
canopy model	SAIL	Canopy RTM model

Table 13: **Parameter configuration for the PROSAIL-PRO model [43].** This table presents the default settings and corresponding notations for parameters used in the RTM block, which simulates leaf and canopy reflectance based on the PROSPECT-PRO [63] and 4SAIL [112] models.

Model	Variable	Notation (unit)	Range
PROSPECT-PRO	Chlorophyll content	C_{ab} ($\mu\text{g}/\text{cm}^2$)	Variable
	Carotenoid content	C_{ar} ($\mu\text{g}/\text{cm}^2$)	Variable
	Anthocyanin content	C_{anth} ($\mu\text{g}/\text{cm}^2$)	Variable
	Water content	C_w (g/m^2)	Variable
	Protein content	C_p (g/m^2)	Variable
	Carbon-based constituents	CBC (g/m^2) = $C_m - C_p$	Variable
	Brown pigment content	Brown (–)	0.25
	Structural coefficient	N_s (–)	1.5
4SAIL	Leaf area index	LAI (m^2/m^2)	Variable
	Average leaf inclination angle	LIDF (Beta index)	5
	Fraction of dry soil	psoil	0.8
	Hotspot	hspot	0.01
	Viewing zenith angle	tto ($^\circ$)	0
	Solar zenith angle	tts ($^\circ$)	30
	Relative azimuth angle	psi ($^\circ$)	0

B.3 Model Architecture Description for the 1D masked autoencoder (MAE)

B.3.1 Training description

Table 14: Architecture of the Masked Autoencoder (MAE)

Module	Layer	Description
Input	Input Spectrum	1D spectral vector $x \in \mathbb{R}^{1 \times S}$
	Patch Embedding	Tokens of size T via frozen Conv1D, $N = S/T$
	Positional Embedding	Fixed 1D sin-cos embeddings
	Masking	Random masking of 75% of tokens
Encoder	Transformer Blocks	d blocks with h -head attention, MLP ($4 \times d$)
	Normalization	LayerNorm
	Output	Latent representation $\in \mathbb{R}^{N \times d}$
Decoder	Linear Projection	Projects latent dim to decoder dimensions
	Token Restoration	Restore with mask tokens using <i>ids_restore</i>
	Transformer Blocks	d' blocks with h' -head attention
	Output Projection	Linear: decoder dim $\rightarrow$ token size T
Reconstruction	Spectrum Output	Reconstructed full spectrum $\in \mathbb{R}^{1 \times S}$

Table 15: **Training hyperparameters of the MAE.** This table lists the default hyperparameters and optimization settings used during training on the pretext task (spectra reconstruction).

Parameter	Value	Description
Input dimension (S)	1720 or 500	Number of spectral bands
Patch size (T)	10, 20, 40, 430	Token size (Ablation study)
Embedding dimension (d)	128	Latent dimension
Encoder depth (d)	4, 6, 8	Transformer blocks (Ablation study)
Encoder heads (h)	4, 8, 16	Attention heads (Ablation study)
Decoder depth (d')	4	Decoder transformer depth
Decoder heads (h')	4	Decoder attention heads
Mask ratio	0.75	Fraction of masked tokens
Loss function	WLoss * Cosine + MSE	Shape and amplitude combined loss
WLoss	1, 0.1, 0.001, 0	Weight for cosine term
MLP ratio	4.0	MLP expansion factor
Attention dropout	0.0	Dropout in attention
Projection dropout	0.0	Dropout in projections
Optimizer	AdamW	With AMSGrad
Learning rate	5e-4	Initial learning rate
Weight decay	1e-4	L2 regularization
Batch size	128	Samples per batch
Epochs	500	Total training epochs

B.3.2 Ablation study results

Table 16: **Ablation study on the effect of transformer depth and attention heads in the MAE model.** This table reports the R^2 and nRMSE values for MAE models evaluated on the trait prediction task, with varying transformer depths and numbers of attention heads. The highest R^2 scores and lowest nRMSE values are highlighted.

Depth	Heads	Final Val R^2	Final Val nRMSE
6	4	0.4492	15.73
6	8	0.4090	16.31
6	16	0.4327	15.96
8	4	0.4351	15.90
8	8	0.4356	15.90
8	16	0.3937	16.44
10	4	0.4060	16.30
10	8	0.4092	16.28
10	16	0.4692	15.43

Table 17: **Ablation study on the effect of cosine similarity loss weight in the MAE model.** This table presents the R^2 and nRMSE values for trait prediction as the cosine similarity loss weight (w_{loss}) is varied in the MAE objective. The highest R^2 scores and lowest nRMSE values are highlighted.

w_{loss}	Final Val R^2	Final Val nRMSE
1	0.5018	14.96
0.1	0.4907	15.15
0.01	0.4233	16.08
0.001	0.4627	15.55
0	0.4698	15.42

Table 18: **Ablation study on the effect of spectral token size in the MAE model.** This table reports the R^2 and nRMSE values for trait prediction as the spectral token (sequence) size is varied in the MAE architecture. The highest R^2 scores and lowest nRMSE values are highlighted.

Token Size	Final Val R^2	Final Val nRMSE
10	0.4542	15.68
20	0.5018	14.96
40	0.4744	15.35
430	0.2683	18.05

C Resource requirements

Table 19: **Comparison of methods by model size, runtime, and hardware usage.** This table summarizes the number of trainable parameters, average runtime, and GPU type used for each method, providing insights into their computational efficiency and resource requirements.

Method	# Trainable Parameters	Run Time	GPU
Supervised (EffNetB0)	6,998,280	~11h	Quadro RTX 8000
MAE (pretext task)	Encoder: 2,006,912 Decoder: 1,220,116 Total: 3,227,028	~20h	Quadro RTX 8000
MAE (downstream Linear Probing)	1,288	~15 min	Quadro RTX 8000
MAE (downstream Fine Tuning)	1,607,196	~15 min	Quadro RTX 8000
SR-GAN	Discriminator: 319,496 Generator: 2,920,130 Total: 3,239,626	~2.5d	Quadro RTX 8000
RTM-AE	35,437,289	~16h	NVIDIA L40S

D complements to results




Figure 10: **Evaluation of trait prediction with variable-size labeled sets.** Validation performance (nRMSE) as a function of the percentage of labeled data used for training. The average nRMSE performance across all traits is indicated by the dashed box. Lower nRMSE values indicate better predictive performance.




Figure 11: **Evaluation of trait prediction with variable-size unlabeled sets.** Validation performance (nRMSE) as a function of the percentage of labeled data used for training. The average nRMSE performance across all traits is indicated by the dashed box. Lower nRMSE values indicate better predictive performance.




Figure 12: **Trait-wise performance heatmaps in the half-range (HR) for MAE vs Baseline.** The heatmap displays the coefficient of determination (R^2 ; higher is better). Each cell represents the average performance across runs for a given trait-method combination.




Figure 13: **Trait-wise performance heatmaps in the half-range (HR) setting MAE vs Baseline.** The heatmap displays the normalized root mean square error (nRMSE; lower is better). Each cell represents the average performance across runs for a given trait-method combination.




Figure 14: **Observed vs. predicted trait values in the cross-dataset OOD setup.** Each subplot corresponds to a specific trait (rows) and method (columns), comparing predicted values to reference data. The black line indicates the 1:1 reference. R^2 and nRMSE values are reported in each plot to quantify predictive performance.

	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
Supervised	<u>0.2836</u>	0.1107	0.3728	0.1176	0.0795	0.3339	0.1784	0.0810	0.1947
SR_GAN	0.3108	0.3065	<u>0.4931</u>	-0.1324	<u>0.1924</u>	<u>0.4810</u>	0.0831	0.1022	<u>0.2296</u>
RTM_AE	0.1846	0.0577	0.4681	0.0790	0.0704	-0.1646	<u>0.1902</u>	-0.0186	0.1083
MAE_FR_LP	0.0542	<u>0.1803</u>	0.3857	<u>0.1673</u>	0.0100	0.4034	0.1864	0.0442	0.1789
MAR_FR_FT	0.2257	0.1337	0.5184	0.2239	0.2425	0.4837	0.2001	<u>0.0967</u>	0.2656
nRMSE ($\downarrow$)									
Supervised	<u>20.1711</u>	30.4925	16.2647	23.4458	18.6372	17.3007	17.7155	20.3835	20.5514
SR_GAN	19.7855	26.9263	<u>14.6236</u>	26.5084	<u>17.4572</u>	<u>15.2720</u>	18.7152	20.1473	<u>19.9294</u>
RTM_AE	21.5192	31.3877	14.9780	23.9535	18.7291	22.8767	<u>17.5886</u>	21.4592	21.5615
MAE_FR_LP	23.1762	<u>29.2746</u>	16.0959	<u>22.7759</u>	19.3277	16.3739	17.6299	20.7873	20.6802
MAR_FR_FT	20.9691	30.0953	14.2528	21.9882	16.9061	15.2323	17.4806	<u>20.2090</u>	19.6417

Table 20: **Cross-dataset generalization by vegetation type: Tundra.** Unlike the overall OOD results (Table 4), here we exclude samples from tundra during evaluation to assess its individual impact on model generalization. Trait-wise performance is reported using R^2 ($\uparrow$) and nRMSE ($\downarrow$). We highlight the best and second-best scores in **bold** and underline, respectively.

	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
Supervised	<u>0.3586</u>	0.0432	0.1471	0.1229	-0.0927	0.1366	<u>0.3344</u>	0.0810	0.1414
SR_GAN	0.3645	0.2524	<u>0.2372</u>	-0.1402	<u>0.0366</u>	<u>0.2824</u>	-0.0012	0.1022	<u>0.1417</u>
RTM_AE	0.1744	-0.0311	0.1994	0.1343	-0.0779	-0.3144	0.2837	-0.0186	0.0437
MAE_FR_LP	0.0412	<u>0.1122</u>	0.1686	<u>0.1624</u>	-0.2675	0.2253	0.2823	0.0442	0.0961
MAE_FR_FT	0.2188	0.0491	0.3057	0.2248	0.0710	0.2927	0.3348	<u>0.0967</u>	0.1992
nRMSE ($\downarrow$)									
Supervised	18.8612	37.3976	18.4678	23.8133	22.9505	19.2217	17.0518	20.3835	22.2684
SR_GAN	18.7736	33.0584	<u>17.4651</u>	27.1220	<u>21.5493</u>	<u>17.5235</u>	20.9137	20.1473	<u>22.0691</u>
RTM_AE	21.3980	38.8223	17.8924	23.6594	22.7941	23.7159	17.6905	21.4592	23.4290
MAE_FR_LP	23.0604	<u>36.0248</u>	18.2334	<u>23.2713</u>	24.7179	18.2073	17.7073	20.7873	22.7512
MAE_FR_FT	20.8151	37.2836	16.6626	22.3877	21.1619	17.3969	17.0472	<u>20.2090</u>	21.6205

Table 21: **Cross-dataset generalization by vegetation type: Forest.** Unlike the overall OOD results (Table 4), here we exclude samples from forest during evaluation to assess its individual impact on model generalization. Trait-wise performance is reported using R^2 ($\uparrow$) and nRMSE ($\downarrow$). We highlight the best and second-best scores in **bold** and underline, respectively.

	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
Supervised	0.2522	0.1019	0.3391	<u>-0.0108</u>	-0.1300	0.2478	<u>0.1712</u>	0.0810	0.1316
SR_GAN	<u>0.2811</u>	0.3282	<u>0.4731</u>	-0.0505	<u>-0.0075</u>	<u>0.3965</u>	0.0441	0.1022	<u>0.1959</u>
RTM_AE	0.2239	0.0859	0.4679	-0.1631	-0.1097	-0.2146	0.1450	-0.0186	0.0521
MAE_FR_LP	0.0786	<u>0.2257</u>	0.3792	-0.0207	-0.1404	0.3455	0.1641	0.0442	0.1345
MAR_FR_FT	0.2897	0.1377	0.5074	0.0678	0.0769	0.4326	0.1835	<u>0.0967</u>	0.2240
nRMSE ($\downarrow$)									
Supervised	20.3244	34.6730	16.6769	<u>24.8992</u>	20.4582	18.4717	<u>17.4666</u>	20.3835	21.6692
SR_GAN	<u>19.9301</u>	29.9879	<u>14.8919</u>	25.3229	<u>19.3170</u>	<u>16.5456</u>	18.7583	20.1473	<u>20.6126</u>
RTM_AE	20.7059	34.9801	14.9635	26.7101	20.2731	23.4729	17.7401	21.4592	22.5381
MAE_FR_LP	22.5613	<u>32.1942</u>	16.1629	25.0216	20.5521	17.2313	17.5413	20.7873	21.5065
MAR_FR_FT	19.8086	33.9740	14.3981	23.9122	18.4904	16.0433	17.3367	<u>20.2090</u>	20.5215

Table 22: **Cross-dataset generalization by vegetation type: Crops.** Unlike the overall OOD results (Table 4), here we exclude samples from crops during evaluation to assess its individual impact on model generalization. Trait-wise performance is reported using R^2 ($\uparrow$) and nRMSE ($\downarrow$). We highlight the best and second-best scores in **bold** and underline, respectively.

	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
Supervised	<u>0.3099</u>	0.1359	0.4383	0.1063	0.1549	0.4156	0.1811	0.0810	0.2279
SR_GAN	0.3318	0.3229	<u>0.5161</u>	-0.1941	<u>0.2496</u>	<u>0.5148</u>	0.0826	0.1022	<u>0.2408</u>
RTM_AE	0.2156	0.0706	0.5020	0.0455	0.1366	-0.1417	<u>0.1872</u>	-0.0186	0.1247
MAE_FR_LP	0.0860	<u>0.2016</u>	0.4433	<u>0.1270</u>	0.0973	0.4713	0.1868	0.0442	0.2072
MAR_FR_FT	0.2459	0.1419	0.5690	0.2074	0.3095	0.5471	0.2068	<u>0.0967</u>	0.2905
nRMSE ($\downarrow$)									
Supervised	20.0381	31.1490	14.8665	23.4980	17.7968	15.6488	17.7263	20.3835	20.1384
SR_GAN	19.7169	27.5735	<u>13.7989</u>	27.0932	<u>16.7698</u>	<u>14.2586</u>	18.7612	20.1473	<u>19.7649</u>
RTM_AE	21.3638	32.3043	13.9977	24.2834	17.9889	21.8728	17.6592	21.4592	21.3661
MAE_FR_LP	23.0615	<u>29.9417</u>	14.7998	<u>23.2242</u>	18.3933	14.8854	<u>17.6639</u>	20.7873	20.3446
MAR_FR_FT	20.9468	31.0410	13.0228	22.1287	16.0873	13.7758	17.4458	<u>20.2090</u>	19.3322

Table 23: **Cross-dataset generalization by vegetation type: Shrubland.** Unlike the overall OOD results (Table 4), here we exclude samples from shrubland during evaluation to assess its individual impact on model generalization. Trait-wise performance is reported using R^2 ($\uparrow$) and nRMSE ($\downarrow$). We highlight the best and second-best scores in **bold** and underline, respectively.

	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
Supervised	0.1565	0.2080	0.3163	0.2771	0.0615	0.3349	0.1825	NaN	0.2195
SR_GAN	0.1989	0.3005	<u>0.4212</u>	-0.1232	<u>0.1658</u>	<u>0.4556</u>	0.1364	NaN	<u>0.2222</u>
RTM_AE	0.0098	0.0615	0.4207	0.3146	0.0452	-0.1678	<u>0.2005</u>	NaN	0.1264
MAE_FR_LP	-0.1716	0.1244	0.3338	0.4612	-0.0172	0.3976	0.2543	NaN	0.1975
MAR_FR_FT	<u>0.0393</u>	<u>0.2813</u>	0.4818	<u>0.4239</u>	0.2123	0.4890	0.1801	NaN	0.3011
nRMSE ($\downarrow$)									
Supervised	<u>21.9248</u>	19.3998	16.7762	20.4917	18.6994	17.1365	20.1887	NaN	19.2310
SR_GAN	21.3651	18.2311	<u>15.4364</u>	25.5063	<u>17.6296</u>	<u>15.5046</u>	20.7497	NaN	<u>19.2033</u>
RTM_AE	23.7549	21.1173	15.4424	19.9577	18.8617	22.7073	19.9653	NaN	20.2581
MAE_FR_LP	25.8392	20.3969	16.5597	17.6901	19.4682	16.3092	19.2815	NaN	<u>19.3636</u>
MAR_FR_FT	23.3975	<u>18.4795</u>	14.6050	<u>18.2936</u>	17.1319	15.0207	20.2180	NaN	18.1638

Table 24: **Cross-dataset generalization by vegetation type: Grassland.** Unlike the overall OOD results (Table 4), here we exclude samples from grassland during evaluation to assess its individual impact on model generalization. Trait-wise performance is reported using R^2 ($\uparrow$) and nRMSE ($\downarrow$). We highlight the best and second-best scores in **bold** and underline, respectively.

	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
Supervised	<u>0.2861</u>	0.1110	0.3663	0.1023	0.0663	0.3336	0.1836	0.0810	0.1913
SR_GAN	0.3143	0.3069	<u>0.4736</u>	-0.1479	<u>0.1707</u>	<u>0.4667</u>	0.0895	0.1022	<u>0.2220</u>
RTM_AE	0.1849	0.0567	0.4620	0.0720	0.0592	-0.1583	<u>0.1924</u>	-0.0186	0.1063
MAE_FR_LP	0.0545	<u>0.1804</u>	0.3809	<u>0.1577</u>	-0.0022	0.4005	0.1858	0.0442	0.1752
MAR_FR_FT	0.2283	0.1339	0.5162	0.2179	0.2266	0.4849	0.2077	<u>0.0967</u>	0.2640
nRMSE ($\downarrow$)									
Supervised	<u>20.1602</u>	30.5654	16.0148	23.5725	18.9082	16.9249	17.6727	20.3835	20.5253
SR_GAN	19.7600	26.9896	<u>14.5977</u>	26.6337	<u>17.8197</u>	<u>15.1402</u>	18.6633	20.1473	19.9689
RTM_AE	21.5421	31.4855	<u>14.7563</u>	23.9671	18.9805	22.3141	<u>17.5763</u>	21.4592	21.5101
MAE_FR_LP	23.2023	<u>29.3485</u>	15.8300	<u>22.8332</u>	19.5902	16.0532	17.6486	20.7873	20.6617
MAR_FR_FT	20.9645	30.1712	13.9845	22.0142	17.2109	14.8817	17.4092	<u>20.2090</u>	19.6056

Table 25: **Cross-dataset generalization by vegetation type: Mix.** Unlike the overall OOD results (Table 4), here we exclude samples from mix during evaluation to assess its individual impact on model generalization. Trait-wise performance is reported using R^2 ($\uparrow$) and nRMSE ($\downarrow$). We highlight the best and second-best scores in **bold** and underline, respectively.

Noise	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 ($\uparrow$)									
0.01	0.552 $\pm$ 0.010	0.602 $\pm$ 0.021	0.677 $\pm$ 0.006	0.564 $\pm$ 0.030	0.659 $\pm$ 0.008	0.679 $\pm$ 0.008	0.588 $\pm$ 0.006	0.422 $\pm$ 0.062	0.593 $\pm$ 0.019
0.03	0.392 $\pm$ 0.072	0.387 $\pm$ 0.095	0.426 $\pm$ 0.059	0.285 $\pm$ 0.094	0.404 $\pm$ 0.118	0.448 $\pm$ 0.065	0.374 $\pm$ 0.101	0.298 $\pm$ 0.032	0.377 $\pm$ 0.079
0.05	-0.031 $\pm$ 0.214	-0.123 $\pm$ 0.212	-0.011 $\pm$ 0.060	-0.181 $\pm$ 0.230	-0.083 $\pm$ 0.173	0.009 $\pm$ 0.059	-0.127 $\pm$ 0.220	0.028 $\pm$ 0.045	-0.065 $\pm$ 0.152
nRMSE ($\downarrow$)									
0.01	16.704 $\pm$ 0.181	13.409 $\pm$ 0.354	10.534 $\pm$ 0.100	17.332 $\pm$ 0.620	10.165 $\pm$ 0.121	10.674 $\pm$ 0.127	12.975 $\pm$ 0.093	16.935 $\pm$ 0.906	13.591 $\pm$ 0.313
0.03	19.436 $\pm$ 1.164	16.617 $\pm$ 1.326	14.021 $\pm$ 0.727	22.154 $\pm$ 1.436	13.392 $\pm$ 1.361	13.983 $\pm$ 0.834	15.952 $\pm$ 1.323	18.669 $\pm$ 0.422	16.778 $\pm$ 1.074
0.05	25.244 $\pm$ 2.650	22.482 $\pm$ 2.104	18.624 $\pm$ 0.555	28.434 $\pm$ 2.725	18.078 $\pm$ 1.438	18.753 $\pm$ 0.563	21.385 $\pm$ 2.152	21.969 $\pm$ 0.507	21.871 $\pm$ 1.587

Table 26: **Supervised: Noise robustness analysis.** Model performance under different noise intensities (0.01, 0.03, 0.05). Trait-wise R^2 (higher is better) and nRMSE (lower is better) are reported as mean $\pm$ standard deviation.

Noise	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 (†)									
0.01	0.539 ± 0.018	0.544 ± 0.006	0.630 ± 0.010	0.522 ± 0.012	0.528 ± 0.017	0.659 ± 0.007	0.533 ± 0.009	0.497 ± 0.012	0.556 ± 0.011
0.03	0.329 ± 0.047	0.292 ± 0.031	0.318 ± 0.060	0.375 ± 0.025	0.264 ± 0.047	0.348 ± 0.096	0.283 ± 0.054	0.108 ± 0.118	0.290 ± 0.060
0.05	-0.038 ± 0.164	-0.071 ± 0.122	-0.195 ± 0.300	0.065 ± 0.212	-0.246 ± 0.280	-0.196 ± 0.364	-0.080 ± 0.103	-0.468 ± 0.317	-0.154 ± 0.233
nRMSE (‡)									
0.01	16.943 ± 0.323	14.367 ± 0.097	11.272 ± 0.147	18.143 ± 0.236	11.958 ± 0.218	11.008 ± 0.115	13.810 ± 0.133	15.811 ± 0.185	14.164 ± 0.182
0.03	20.421 ± 0.708	17.887 ± 0.390	15.294 ± 0.664	20.773 ± 0.442	14.897 ± 0.478	15.175 ± 1.130	17.108 ± 0.656	21.016 ± 1.414	17.821 ± 0.735
0.05	25.367 ± 1.969	21.970 ± 1.256	20.147 ± 2.468	25.295 ± 2.775	19.220 ± 2.021	20.449 ± 3.018	20.982 ± 1.001	26.895 ± 3.012	22.541 ± 2.190

Table 27: **SR_GAN: Noise robustness analysis.** Model performance under different noise intensities (0.01, 0.03, 0.05). Trait-wise R^2 (higher is better) and nRMSE (lower is better) are reported as mean ± standard deviation.

Noise	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 (†)									
0.01	0.576 ± 0.028	0.633 ± 0.020	0.659 ± 0.018	0.549 ± 0.029	0.654 ± 0.012	0.666 ± 0.022	0.551 ± 0.040	0.312 ± 0.140	0.575 ± 0.038
0.03	0.527 ± 0.037	0.595 ± 0.004	0.620 ± 0.032	0.515 ± 0.043	0.602 ± 0.015	0.636 ± 0.029	0.499 ± 0.056	0.219 ± 0.164	0.527 ± 0.048
0.05	0.377 ± 0.023	0.439 ± 0.082	0.456 ± 0.054	0.457 ± 0.042	0.456 ± 0.068	0.456 ± 0.066	0.350 ± 0.055	-0.258 ± 0.880	0.342 ± 0.159
nRMSE (‡)									
0.01	16.245 ± 0.531	12.892 ± 0.346	10.817 ± 0.281	17.615 ± 0.567	10.246 ± 0.172	10.879 ± 0.350	13.536 ± 0.616	18.420 ± 1.940	13.831 ± 0.600
0.03	17.153 ± 0.667	13.543 ± 0.066	11.411 ± 0.486	18.302 ± 0.782	10.987 ± 0.214	11.356 ± 0.451	14.282 ± 0.813	19.619 ± 2.134	14.582 ± 0.702
0.05	19.690 ± 0.362	15.899 ± 1.198	13.647 ± 0.686	19.385 ± 0.771	12.828 ± 0.818	13.884 ± 0.863	16.284 ± 0.692	24.058 ± 8.329	16.959 ± 1.715

Table 28: **RTM_AE: Noise robustness analysis.** Model performance under different noise intensities (0.01, 0.03, 0.05). Trait-wise R^2 (higher is better) and nRMSE (lower is better) are reported as mean ± standard deviation.

Noise	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 (†)									
0.01	0.518 ± 0.003	0.432 ± 0.005	0.587 ± 0.005	0.416 ± 0.006	0.438 ± 0.009	0.591 ± 0.006	0.443 ± 0.008	0.170 ± 0.017	0.449 ± 0.007
0.03	0.153 ± 0.029	0.357 ± 0.026	0.486 ± 0.004	0.395 ± 0.011	0.253 ± 0.028	0.476 ± 0.017	0.250 ± 0.009	0.100 ± 0.025	0.309 ± 0.019
0.05	-0.677 ± 0.039	0.084 ± 0.066	0.387 ± 0.002	0.359 ± 0.018	0.043 ± 0.031	0.372 ± 0.018	-0.215 ± 0.008	0.001 ± 0.030	0.044 ± 0.027
nRMSE (‡)									
0.01	16.858 ± 0.048	14.712 ± 0.058	12.747 ± 0.069	18.752 ± 0.092	15.937 ± 0.133	12.861 ± 0.094	15.070 ± 0.107	20.346 ± 0.203	15.910 ± 0.101
0.03	22.348 ± 0.391	15.658 ± 0.318	14.222 ± 0.053	19.085 ± 0.182	18.367 ± 0.349	14.551 ± 0.230	17.494 ± 0.107	21.187 ± 0.292	17.864 ± 0.240
0.05	31.452 ± 0.367	18.672 ± 0.675	15.524 ± 0.029	19.645 ± 0.280	20.783 ± 0.352	15.928 ± 0.236	22.266 ± 0.076	22.317 ± 0.338	20.823 ± 0.294

Table 29: **MAE_FR_LP: Noise robustness analysis.** Model performance under different noise intensities (0.01, 0.03, 0.05). Trait-wise R^2 (higher is better) and nRMSE (lower is better) are reported as mean ± standard deviation.

Noise	cab	cw	cm	LAI	cp	cbc	car	anth	Average
R^2 (†)									
0.01	0.583 ± 0.012	0.645 ± 0.024	0.787 ± 0.006	0.648 ± 0.014	0.667 ± 0.026	0.781 ± 0.007	0.584 ± 0.037	0.447 ± 0.031	0.643 ± 0.020
0.03	0.440 ± 0.028	0.536 ± 0.007	0.659 ± 0.009	0.530 ± 0.021	0.496 ± 0.035	0.647 ± 0.012	0.461 ± 0.029	0.264 ± 0.029	0.504 ± 0.021
0.05	0.242 ± 0.045	0.388 ± 0.023	0.477 ± 0.008	0.411 ± 0.014	0.276 ± 0.060	0.460 ± 0.009	0.258 ± 0.027	0.136 ± 0.073	0.331 ± 0.032
nRMSE (‡)									
0.01	15.677 ± 0.227	11.628 ± 0.388	9.141 ± 0.121	14.564 ± 0.290	12.265 ± 0.478	9.401 ± 0.147	13.025 ± 0.575	16.594 ± 0.462	12.787 ± 0.336
0.03	18.163 ± 0.461	13.302 ± 0.099	11.571 ± 0.146	16.823 ± 0.378	15.083 ± 0.526	11.947 ± 0.200	14.830 ± 0.406	19.154 ± 0.380	15.109 ± 0.325
0.05	21.131 ± 0.621	15.270 ± 0.291	14.346 ± 0.105	18.836 ± 0.219	18.078 ± 0.748	14.772 ± 0.123	17.400 ± 0.320	20.738 ± 0.879	17.571 ± 0.413

Table 30: **MAE_FR_FT: Noise robustness analysis.** Model performance under different noise intensities (0.01, 0.03, 0.05). Trait-wise R^2 (higher is better) and nRMSE (lower is better) are reported as mean ± standard deviation.




Figure 15: **Feature importance of MAE-based downstream regression.** Results are shown for (top) linear probing (*MAE_FR_LP*), (middle) fine-tuning the last block (*MAE_FR_9B_FT*), and (bottom) full fine-tuning (*MAE_FR_FT*). The blue lines indicate the importance scores across spectral bands, while the orange line shows a reference of a vegetation spectra.