

Hybrid modelling of leaf traits

Integrating neural networks with radiative transfer theory

Sun, Peng; van Bodegom, Peter M.; Timmermans, Joris; Liu, Shuwen; Wu, Jin; Visser, Marco D.

DOI

[10.1016/j.rse.2025.114958](https://doi.org/10.1016/j.rse.2025.114958)

Publication date

2025

Document Version

Final published version

Published in

Remote Sensing of Environment

Citation (APA)

Sun, P., van Bodegom, P. M., Timmermans, J., Liu, S., Wu, J., & Visser, M. D. (2025). Hybrid modelling of leaf traits: Integrating neural networks with radiative transfer theory. *Remote Sensing of Environment*, 329, Article 114958. <https://doi.org/10.1016/j.rse.2025.114958>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Hybrid modelling of leaf traits: Integrating neural networks with radiative transfer theory

Peng Sun^{a,*}, Peter M. van Bodegom^a, Joris Timmermans^b, Shuwen Liu^c, Jin Wu^{c,d}, Marco D. Visser^a

^a Institute of Environmental Sciences, University of Leiden, Van Steenis Building, Einsteinweg 2, 2333CC Leiden, the Netherlands

^b Faculty of Civil Engineering and Geosciences, Delft University of Technology, Delft, the Netherlands

^c School of Biological Sciences and Institute for Climate and Carbon Neutrality, The University of Hong Kong, Pokfulam, Hong Kong, China

^d State Key Laboratory of Agrobiotechnology (CUHK), Shatin, Hong Kong, China

ARTICLE INFO

Editor: Jing M. Chen

Keywords:

Remote sensing
Spectral analysis
Hyperspectral
Leaf traits
Deep learning
Radiative transfer model

ABSTRACT

The rapid growth of remotely sensed earth observation data presents clear opportunities for monitoring complex ecosystem change and answering fundamental ecological questions. However, large-scale automated monitoring of ecosystems faces challenges. Data-driven models require extensive datasets and often lack generalizability when training data are unrepresentative, while process-driven models, such as radiative transfer models (RTMs), can be imprecise due to gaps in knowledge, or simplified representation of physical processes. To enhance the prediction of plant functional traits and simultaneously discover where process-driven models can be improved, we explore the potential of Physics-Informed Neural Networks (PINNs) as a hybrid approach that combines the strengths of both methodologies at the leaf scale. In contrast to data augmentation approaches, our implementation directly integrates the widely-used PROSPECT5B model into the architecture of an autoencoder framework. Our results show that our PINNs approach is able to outperform data-driven techniques even when trained on very limited training data (i.e. 17 % training vs 83 % validation). We also identified weak points in the PROSPECT5B model by progressively replacing individual components of PROSPECT5B with convolutional neural networks. Our case study indicates that especially Prospect's generalized "plate model" could be refined to improve predictive ability. Hence, our framework provides a self-diagnostic capability and identifies areas for improvement in process-driven models and their components. Thus, we conclude that PINNs 1) improve data-driven predictive accuracy while maintaining physical consistency with minimal training data while 2) being able to identify limitations in process-driven models. Hence, we believe our framework could serve as a new standard for evolving and improving radiative transfer models.

1. Introduction

The rapid expansion of Earth observation data, combined with advances in machine learning and computational power, has enabled new opportunities for large-scale ecological monitoring and analysis (Borowiec et al., 2022; Christin et al., 2019; Cord et al., 2017; Miralles et al., 2023; Persello et al., 2022). However, linking reflected radiation to ecosystem functional traits remains challenging, as this relationship is shaped by complex and scale-dependent processes, including vegetation dynamics that vary across time and space (Fensholt et al., 2015). As Timmermans and Daniel Kissling (2023) note, there is a pressing need for tailored satellite remote sensing products that bridge the gap

between detailed physical process understanding and practical data-driven insight.

Data-science techniques such as machine learning provide an ideal approach for analyzing these datasets, given their computational efficiency and ability to handle complex situations (Pérez-Cutillas et al., 2023). However, despite the substantial volumes of Earth observation data, the availability of representative ground-based validation datasets is often limited, which may be reflected in the variability and inconsistency in upscaled trait maps (Meyer and Pebesma, 2022; Todman et al., 2023; Dechant et al., 2024). Moreover, machine learning models, including deep, convolutional, and recurrent neural networks, tend to struggle to provide accurate predictions when the training data do not

* Corresponding author.

E-mail addresses: sunpeng18@mails.ucas.ac.cn, p.sun@cml.leidenuniv.nl (P. Sun).

<https://doi.org/10.1016/j.rse.2025.114958>

Received 21 January 2025; Received in revised form 1 August 2025; Accepted 5 August 2025

Available online 9 August 2025

0034-4257/© 2025 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

encompass the full spectrum of scenarios (Coveney et al., 2016; Todman et al., 2023). Consequently, current data-driven methods face criticism for their extensive data and representation requirements, lack of generalizability to out-of-sample scenarios, or physically inconsistent results (Meyer and Pebesma, 2022; Verrelst et al., 2015; Willard et al., 2022).

In contrast, traditional mechanistic models, such as radiative transfer models (RTMs) offer some advantages due to their inherent generalizability. RTMs simulate physical and biological interactions to help understand how vegetation reflects and transmits radiation (Medvigy et al., 2009; Yang et al., 2020), and as they are based on underlying physical laws, they ensuring at least a basic consistency with real-world phenomena. However, these models also have limitations, such as a lack of precision or computational feasibility at desired resolutions, mostly due to simplification, incomplete knowledge of certain processes or the difficulty in capturing site-specific parameters for model parameterization (Gastellu-Etchegorry et al., 2012; Gumiere et al., 2020; Hatfield et al., 2021; Martínez-Ferrer et al., 2022). These limitations arise partly from their reliance on prior knowledge, such as assumptions needed to address unidentifiable parameters, and partly from an incomplete understanding of the system's complex behaviour.

Given that neither purely mechanistic nor purely data-driven approaches suffice for predicting complex phenomena, especially with limited data, researchers are increasingly exploring hybrid models that integrates scientific knowledge and machine learning (ML). Such hybrid modelling is gaining momentum in various fields (Cavanagh et al., 2021; Das and Tesfamariam, 2022; Lam et al., 2023; Raissi et al., 2019; Reichstein et al., 2019), including remote sensing (Rao et al., 2020). Hybrid models attempt to leverage the strengths of both approaches to address the challenges posed by purely mechanistic or data-driven methods.

A notable form of these models is Physics-Informed Neural Networks (PINNs), which integrate machine learning with physical laws (Raissi et al., 2019) and can be a powerful tool for solving problems that require adherence to physical laws. PINNs offer several advantages: they are 1) computationally efficient, 2) have greater predictive accuracy than purely mechanistic models, 3) generally require less data for good out-of-sample performance compared to the typical data-driven, and 4) offer improved interpretability over neural networks. Despite their potential, PINNs are still in their infancy within remote sensing, and it remains unclear whether they can help overcome the limitations currently faced by radiative transfer models and machine learning retrieval algorithms.

While physics-informed approaches are gaining traction in remote sensing, the full integration of physical models into neural network architectures remains extremely rare. García-Soria et al. (2024) make an important contribution by improving uncertainty estimation for hybrid models aimed at retrieving canopy nitrogen content from imaging spectroscopy data. Their work reflects a growing trend in hybrid retrieval frameworks, though it does not structurally embed physical equations to constrain the models. Dehghan-Shoar et al., 2024 present an innovative physics-guided approach, fitting surrogate models derived from radiative transfer model outputs to constrain predictions. While this offers a flexible way to incorporate physical knowledge, it treats physics externally. In contrast, PINNs embed the governing equations directly into the learning process, enabling the model to enforce physical consistency throughout. Zerah et al. (2024) take a further step by embedding radiative transfer equations into a variational autoencoder for large-scale PROSAIL inversion from Sentinel-2 imagery. While their results show strong predictive performance, they also exhibit saturation effects for several key leaf traits—suggesting that a bottleneck may lie in the underlying leaf-scale modelling and upscaling rather than in the machine learning approach itself. Together, these studies underscore the need to improve the physical modelling of vegetation at the leaf level, in parallel with the development of scalable hybrid architectures.

Addressing this gap, we developed a PINN that directly embeds the

biophysical equations of the PROSPECT5B leaf reflectance model into a neural network. This allows high-fidelity leaf-scale trait retrieval under data-limited conditions, while also enabling diagnostic insights into where and how current process-based models may require refinement. We focus on the remote sensing of leaf biochemical traits due to their broad relevance and the availability of well-characterized spectral databases for training and validation. PROSPECT5B (Féret et al., 2008), one of the most widely used radiative transfer models for leaf reflectance, provides a strong foundation for this integration. We evaluate whether PINNs can (1) improve the accuracy of leaf trait prediction from hyperspectral data in limited-data settings, (2) outperform mechanistic models in computational efficiency, and (3) enhance model interpretability by identifying components of the physical model that may benefit from further development. These advantages present a promising path forward—not only for trait prediction, but also for improving the physical understanding underlying radiative transfer modelling in ecological remote sensing.

2. Data and methods

2.1. Overview

To assess the potential of hybrid modelling (McGreivy and Hakim, 2024) for ecological remote sensing in trait prediction from leaf optical data, we assembled eight datasets comprising field-measured leaf reflectance spectra and corresponding functional traits, collected using spectroradiometers, for model training and validation. These datasets encompass a large environmental gradient, including various continents, climates, vegetation types and land cover (Table 1). Our approach involves training a PINN model on the two original calibration datasets used for PROSPECT5B. We then tested the model's predictive capability on the six remaining datasets (see Table 1). The performance of the PINN was compared to the widely used Partial Least Squares Regression (PLSR) method in remote sensing. Additionally, we evaluated the PINN against RTM model inversion using Markov Chain Monte Carlo (MCMC). Finally, we built three replacement PINNs where components of PROSPECT5B were replaced with artificial neural nets. These replacements allowed identifying which components of the RTM model benefit most from integration with neural networks, emphasizing areas where the mechanistic RTMs can be substantially improved.

2.2. RTMs, design and training of the PINNs

2.2.1. Data

We compiled data from eight sources featuring full-range leaf spectra (400–2400 in 1 nm steps) along with corresponding traits: C_w (EWT, Equivalent Water Thickness), C_m (Carbon mass per Area, or LMA, Leaf Mass per Area), C_{ab} (Chlorophyll $a + b$ content), C_{ar} (Carotenoid content) and the number of cell-air interfaces (N) (Ely et al., 2019; Féret et al., 2008; Helsen et al., 2021; Serbin et al., 2019). The data were sourced from various global sites, although most datasets originate from Europe and the United States, they do represent diverse environmental conditions (Fig. 1, Table 1). An additional “unlabelled” dataset (Meireles et al., 2020) on leaf spectra without corresponding leaf traits (Schweiger et al., 2018) was included to enhance the model's robustness (see 2.1.4). All leaf reflectance spectra were acquired using spectroradiometers equipped with a leaf clip attached to a plant probe. Detailed information on these datasets is provided in Table 1. These datasets are publicly available through the Ecological Spectral Information System (ECOSIS, <https://ecosis.org/>) and figshare, with individual data sources listed in Table S.1.

2.2.2. Radiative transfer theory

To integrate physical process information into the PINNs, we used the PROSPECT5B family of leaf RTMs, originally introduced by Jacquemoud and Baret (1990). Specifically, we focused on the

Table 1

The region, country, sample number, coordinates, Ecosystem Type and climate conditions in each leaf scaled data source.

Data Name	Country	traits	Samples	Equipment	Lat (°)	Long (°)	Ecosystem Type	Climate Condition (Köppen Classification)
Training sets								
1 lopex	Italy	N, Cab, Car, Cw & Cm	320	lambda-19 spectrophotometer	45.8179	8.6078	Leaf samples from over 50 species of trees, crops, and plants, including conifers and broadleaf species	Temperate humid climate
2 anger	France	N, Cab, Car, Cw & Cm	276	ASD FieldSpec III (ASD, Inc. Boulder, CO, USA)	47.4712	0.5518	Leaf samples from over 40 plant species, including trees and crops	Temperate oceanic climate
Unlabelled datasets								
3 Meireles		None	16,765	Multiple, see original source			Leaf samples from over 40 plant species, including trees and crops	Temperate oceanic climate
Validation sets								
4 nasa_fft	the United States	N, Cab, Car, Cw & Cm	682	ASD FieldSpec III (ASD, Inc. Boulder, CO, USA)	Multiple		Various tree species from northern temperate and boreal forests	Humid continental climate
5 bel_nal	Belgium and Japan	Cw & Cm	994	ASD FieldSpec III (ASD, Inc. Boulder, CO, USA)	Multiple		<i>Rosa rugosa</i> (wrinkled rose) in Hokkaido and 4 plant species from Belgian dune grasslands	Subarctic humid climate
6 ifgg	Germany	Cw & Cm	739	ASD FieldSpec III (ASD, Inc. Boulder, CO, USA)	51.4783	−0.298	Over 40 grassland species, including herbaceous plants and grasses	Temperate oceanic climate
7 bel	Belgium	Cw & Cm	256	(SVC HR-1024TM/i, Spectra Vista Corporation, NY, USA)	50.5039	4.4699	36 species growing in <i>Rosa rugosa</i> -invaded coastal grassland communities	Temperate oceanic climate
8 bnl	the United States	Cw & Cm	184	(SVC HR-1024TM/i, Spectra Vista Corporation, NY, USA)	40.8643	−72.8752	8 crop species, including structural and biochemical leaf trait data	Temperate humid climate

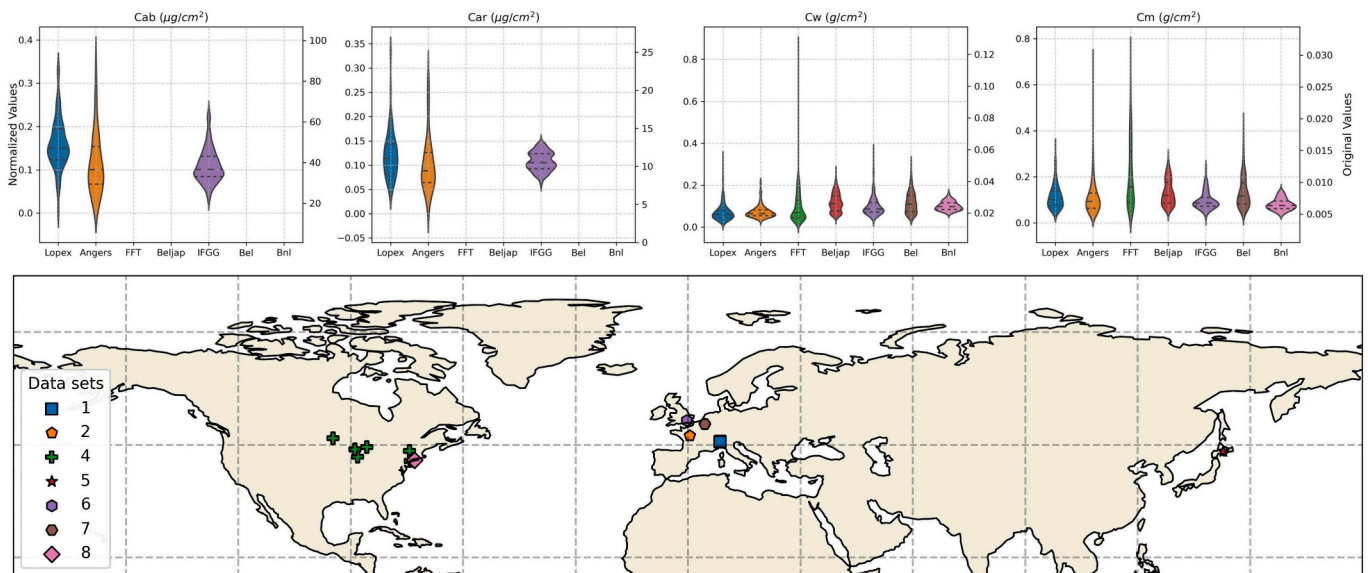


Fig. 1. The distribution of leaf biochemical trait values across sites (top panels) and locations of seven data sources used in this study (bottom panel). Cab = chlorophyll a + b content, Car = carotenoid content, Cw = canopy water content (equivalent water thickness), and Cm = carbon mass per area. Note: dataset 3 contains unlabelled spectra and thus does not have a single geo-location (see original source).

PROSPECT5B model (Féret et al., 2008), which effectively reconstructs leaf directional-hemispherical reflectance and transmittance across the 400–2500 nm wavelength range using six leaf optical parameters (hereafter traits): Chlorophyll a + b (C_{ab} ; $\mu\text{g}/\text{cm}^2$), Carotenoids (C_{ar} ; $\mu\text{g}/\text{cm}^2$), Brown pigments (C_{br} , Arbitrary units), Equivalent Water Thickness (C_w , g/cm^2), Leaf Mass per Area (C_m , g/cm^2) and the number of cell-air interfaces (N). The PROSPECT5B model comprises three physical modules:

1. Leaf boundary transmissivity function ($T_{av}(\alpha, n_1(\lambda), n_2(\lambda))$): This function, based on Fresnel's equations and effective diffractive indices ($n_1(\lambda), n_2(\lambda)$), calculates the average transmissivity at air-leaf interfaces by integrating over all incident angles up to the maximum angle α .
2. Plate element transmission function ($\tau(\lambda)$): This function determines light transmission through individual leaf plates by applying a hemispherical integration of the Beer-Lambert law to calculate absorption by leaf biochemical traits.

3. Generalized plate model (pm($\alpha, n, 1, n, 2, k$)): This model simulates leaf reflectance and transmission using N stacked plates, following principles established by Stokes (1862) and Allen et al. (1969).

We next briefly introduce these elements to foster a clear understanding of the model. For comprehensive details, please refer to the original papers (Féret et al., 2008; Jacquemoud and Baret, 1990) and textbooks (Jacquemoud and Ustin, 2019):

The PROSPECT5B models adapts Allen et al. (1969)'s generalized plate model to include leaf microtopography at the epidermal cell scale. Although incident light is generally considered perpendicular to the leaf blade, the model accounts for the leaf's micro-morphology by simulating incident radiation striking a flat plane at angles ranging from 0 to α , with α typically set to $\sim 40^\circ$ for most leaves (Féret et al., 2008). After the first plate, light is assumed to be isotropic and diffuse, incident from 0 to 90° . The reflection of a single elementary plate is given by:

$$R(n(\lambda), k(\lambda), \alpha) = r_{12} + \frac{T_{av}(\alpha, 1, n(\lambda)) T_{av}(90, n(\lambda), 1) r_{21} \tau(k(\lambda))^2}{1 - r_{21}^2 \tau(k(\lambda))^2} \quad (1)$$

Similarly, transmittance is given by:

$$T(n(\lambda), k(\lambda), \alpha) = \frac{T_{av}(\alpha, 1, n(\lambda)) T_{av}(90, n(\lambda), 1) \tau(k(\lambda))}{1 - r_{21}^2 \tau(k(\lambda))^2} \quad (2)$$

Here, $T_{av}(\alpha, 1, n(\lambda))$ represents the average transmissivity at the air-to-leaf interface, integrated from 0 to angle α . $T_{av}(90, n(\lambda), 1)$ denotes the average transmissivity of the leaf-to-air interface, integrated hemispherically, assuming light passing through a plate is diffuse. The effective refractive index of the plate is given by $n(\lambda)$ while the reflectivity of an interface is given by $r_{12} = 1 - T_{av}(\alpha, 1, n(\lambda))$ and $r_{21} = 1 - T_{av}(90, n(\lambda), 1)$, assuming no discernible absorption at the interface. The function $\tau(k(\lambda))$ predicts transmission through a leaf plate while considering the absorption by various leaf biochemical traits. It is defined as:

$$\tau(k(\lambda)) = (1 - k(\lambda)) e^{-k(\lambda)} + k(\lambda)^2 E_1(k(\lambda)) \quad (3)$$

This function represents the hemispherical integration of the Beer-Lambert law, with $E_1()$ being the exponential integral. Here, $k(\lambda)$ is an empirically derived absorption coefficient for a single plate of unit depth. It is computed as the linear combination of five individual leaf biochemical constituents, distributed across each of the N plates:

$$k(\lambda) = \frac{1}{N} \sum_{i=1}^5 C_i \times k_i(\lambda) \quad (4)$$

Here, C_i represents the mass content per unit area of the i th biochemical constituent, and $k_i(\lambda)$ is its specific absorption coefficient (SAC). Values of $k_i(\lambda)$ have been empirically calibrated using measurements of C_i . This calibration was done by inverting the PROSPECT5B model and minimizing the summed squared error between the modelled and measured leaf reflectance and transmission, using datasets 1 & 2 in Table 1 (Féret et al., 2008).

Finally, the total reflectance of a stack comprising N plates can be calculated, considering $N-1$ plates positioned beneath the first, with the first plate incorporating micro topographical leaf structure:

$$R(N, n(\lambda), k(\lambda), \alpha) = \frac{R_a(ab^{N-1} - a^{-1}b^{1-N}) + (T_a T_{90} - R_a R_{90})(b^{N-1} - b^{1-N})}{ab^{N-1} - a^{-1}b^{1-N} - R_{90}(b^{N-1} - b^{1-N})} \quad (5)$$

With transmittance given by:

$$T(N, n(\lambda), k(\lambda), \alpha) = \frac{T_a(a - a^{-1})}{ab^{N-1} - a^{-1}b^{1-N} - R_{90}(b^{N-1} - b^{1-N})} \quad (6)$$

Here, T_a , T_{90} , R_a and R_{90} are calculations from eqs. 1 and 2, with integration angles α and 90° , respectively. The symbols a and b are composite algebraic functions of transmittance (T_{90}, T_a) and reflectance

(R_a , R_{90}) given in previous studies and textbooks e.g. (Jacquemoud and Ustin, 2019).

2.2.3. PINN: Integrating neural networks within leaf radiative transfer models

PINNs are characterized by integrating physical laws into their architecture and training processes. This can be achieved either through inclusion of differential equations within the network or by using a loss function that penalizes deviations from physical laws, typically measured as residuals between the network's predictions and the physical equations (Cuomo et al., 2022; Willard et al., 2022). We implemented a PINN that uses an Encoder-decoder architecture (autoencoder) similar to Zerah et al. (2024).

This design condenses input data into an efficient lower-dimensional form (encoder) and then reconstructs it back to the original format (decoder). This generally allows the autoencoder to capture the most salient features of the data (Meng et al., 2017). Leaf RTMs, such as PROSPECT5B, naturally fit within this framework, as they essentially take encoded information (leaf trait parameters) and decode these into leaf spectra. We capitalized on this natural fit and constructed an encoder that acts as an RTM-emulator, extracting traits from spectral data, to feed to the decoder. The decoder then incorporates (elements of) PROSPECT5B, and uses the traits predicted by the encoder to recreate the spectra in a physically constrained way (Fig. 2).

We selected PROSPECT5B as a base model to enable a focused examination of structural components that have remained unchanged across PROSPECT versions. Later models, including PROSPECT-D and PROSPECT-PRO (Féret et al., 2021), introduce altered or additional biochemical inputs with recalibrated spectral absorption coefficients (SACs; Eq. 4). While these extensions improve ecological interpretability, their gains in predictive accuracy appear modest across global datasets (Féret et al., 2021; Visser et al., 2025). PROSPECT5 thus remains a relatively simple yet competitive version, allowing us to isolate the impact of structural modifications to the radiative transfer formulation without the added complexities of empirical SAC recalibration.

2.2.3.1. Autoencoder design. We constructed four basic PINN formulations (Fig. 2). Each starts with an encoder ($E(R)$), which takes a set leaf reflectance spectra R ($n \times k$ matrix with n samples over k wavelengths) and reduces it to a set of $n \times 5$ traits (hereafter the trait-encoding). Predicted traits are then fed to the decoder ($D(T)$), which attempts to recreate the original spectra. We built four different versions of the decoder; an RTM-only version using PROSPECT5B as the decoder $D_{rtm}(T)$, and three alternative decoders where components of PROSPECT5B were replaced with artificial neural nets (see 2.1.3). These alternatives include a Convolutional Neural Network (CNN) replacement for the leaf transmission ($D_{\tau}(T)$), a generalized plate model function $D_{pm}(T)$, and a combination of both transmission and plate models $D_{tpm}(T)$. Replacing individual decoder sub-modules reveals which parts of PROSPECT5B leave the most room for improvement, highlighting discrepancies and guiding future refinement. Note that the CNN decoder replacements here are not viewed as ideal physically grounded modules, but rather as diagnostic tools: a flexible functional approximator to assess the impact of replacing radiative transfer components.

2.2.3.2. Encoder design. Our encoder consists of two linear layers with sigmoid activation functions. Unlike traditional neural networks that treat all input data uniformly during training, we incorporated simplified attention mechanisms, as described by Ghaffarian et al. (2021), with small differences in our implementation: we use the Hadamard product between the input data and its transformed version, rather than the normalized dot product. These mechanisms enable the encoder to discern and prioritize the most informative elements relevant to each trait prediction, enhancing the efficiency of the encoding process, allowing identification and concentration on the most pertinent parts of

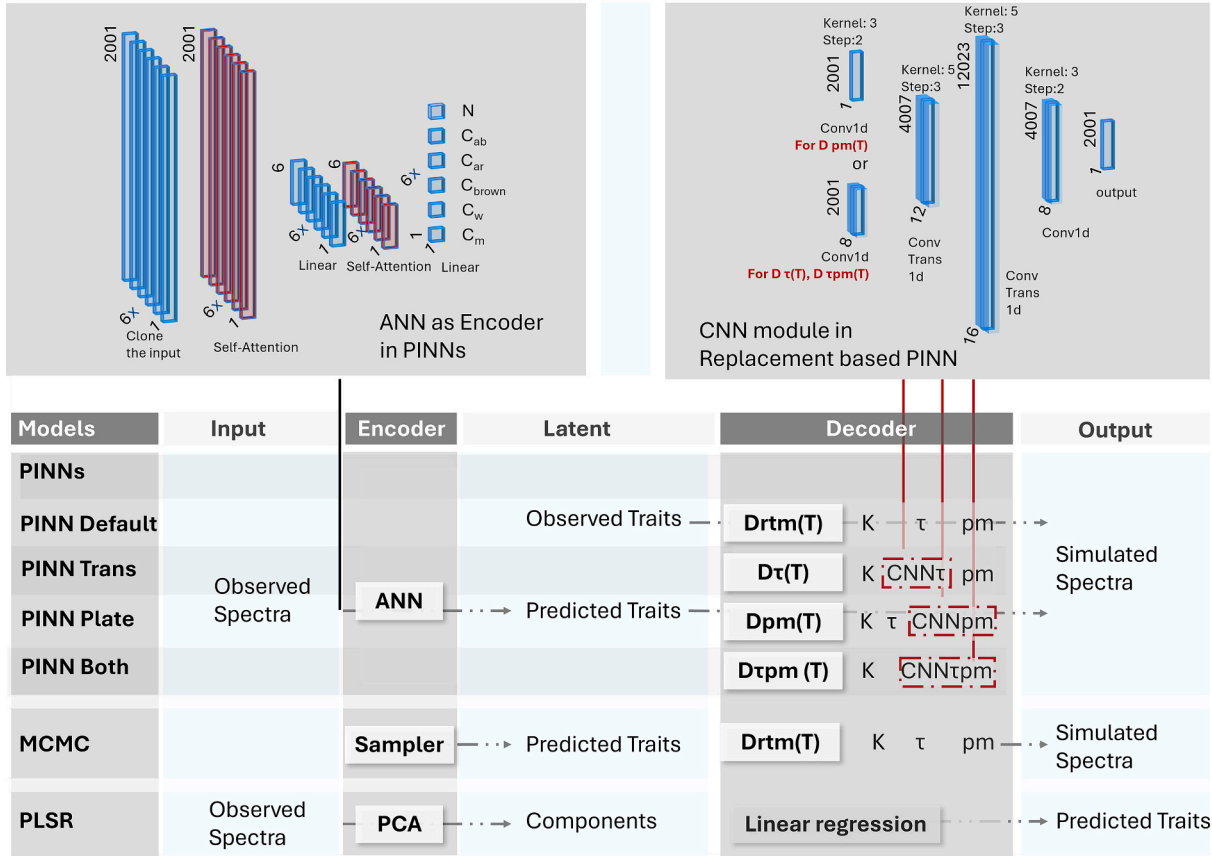


Fig. 2. Overview of the model architecture, module replacement strategy (top panels), and comparison frameworks (bottom panel). We illustrate four decoder variants: the original radiative transfer model (PROSPECT5B) and three hybrid versions where specific components—transmission (τ), plate model (pm), or both—are replaced by neural networks. These decoders are interchangeable within the full PINN framework. K denotes the absorption coefficient matrix; τ , the transmission function; pm , the plate model. *Latent* refers to the compressed feature representation learned by the encoder. In PLSR, *Components* are PCA-derived; in PINN, *ANN* denotes the encoder network.

Top-left panel: The encoder applies six parallel attention-like heads. Each head consists of a linear layer with sigmoid activation generating soft gates, followed by element-wise multiplication with the input, and a second linear-sigmoid layer reducing each gated signal to a scalar. The outputs are concatenated and passed through a final linear-sigmoid layer that reweights the combined features via element-wise multiplication. Unlike standard attention mechanisms, this design omits query-key-value projections, normalization, residuals, and positional encodings.

Top-right panel: The replacement decoder module substitutes RTM components with a learnable upsampling architecture. It starts with replication padding to maintain boundary conditions, followed by two transpose convolution layers (from 8/1 to 12 channels with kernel size 3/stride 2; then to 16 channels with kernel size 5/stride 3). Feature dimensions are reduced via two 1D convolutions (16 \rightarrow 8 \rightarrow 1). All intermediate layers use ReLU activations; the output uses a sigmoid. Slanted numbers indicate channel depth; vertical numbers indicate feature width.

the input data.

2.2.3.3. Decoder design. Our decoder architecture adheres to the general structure of the PROSPECT5B model. The simplest implementation of our decoder directly employs the PROSPECT5B model as a “default” RTM-decoder. For this, we refactored the original Fortran-based PROSPECT5B code into CUDA-enabled GPU code that interfaces with Torch. This integration allows us to embed radiative transfer theory directly into the architecture and training of an autoencoder model.

2.2.4. Training and loss functions

The PINN autoencoder is trained using a semi-supervised approach, incorporating multiple loss functions that utilize Mean Square Error (MSE) as the optimization objective, leveraging both labelled and unlabelled data, facilitating improved model robustness and accuracy in predicting leaf traits from spectral data:

$$L_{\text{total}} = L_{\text{Reconstruction}} + L_{\text{traits reconstruction}} + L_{\text{cycle consistency}} \quad (7)$$

where,

$$L_{\text{Reconstruction}} = L_{\text{encoder reconstruction}} + L_{\text{decoder reconstruction}} + L_{\text{module replacement}} + L_{\text{unlabeled spectra reconstruction}} \quad (8)$$

We adopted a simple loss weighting scheme to prioritize comparability across (RTM-replacement) module variants: all loss terms are weighted equally (1.0), except for the cycle-consistency loss, which is down-weighted to 0.1.

2.2.5. Reconstruction losses

These losses assess the accuracy in predicting leaf spectra and have three components: $L_{\text{encoder reconstruction}} = \text{MSE}(R, D(E(R)))$ evaluates how well the decoder reconstructs the predicted R from the encoded traits $E(R)$. This loss ensures the basic functionality of the encoder-decoder pipeline for spectral reconstruction. $L_{\text{decoder reconstruction}} = \text{MSE}(R, D(T))$ evaluates how well the decoder reconstructs the spectra from the original paired traits T . It is informative in concert with $L_{\text{module replacement}} = \text{MSE}(R, D_i(T))$ and is used when evaluating replacement modules in the RTM. Their fit assesses the robustness of the CNN by evaluating the decoder’s reconstruction performance when applied to an alternative representation of the decoder (see section 2.3: *Model component*

replacement evaluation). $L_{\text{unlabelled spectra reconstruction}} = \text{MSE}(\mathbf{R}_-, D(E(\mathbf{R}_-)))$ measures the reconstruction quality of the decoder when applied to the encoded unlabelled data $E(\mathbf{R}_-)$. This loss highlights the model's ability to generalize to unlabeled data for spectral reconstruction, and further prevents overfitting to the training data.

2.2.6. Traits reconstruction loss

$L_{\text{traits reconstruction}} = \text{MSE}(T, E(R))$ quantifies the accuracy of the encoder by measuring the loss between the predicted traits and observed traits T , and can be seen as conceptually equivalent to a prior in Bayesian inference (Shiklomanov et al., 2016). It serves as a critical evaluation of the encoder's predictive capabilities to represent the data.

2.2.7. Cycle consistency loss

$L_{\text{cycle consistency}} = \text{MSE}(E(\mathbf{R}_-), E(D(E(\mathbf{R}_-))))$ evaluates the consistency of the encoder-decoder cycle by assessing how well the model preserves information through a round-trip transformation. This loss compares the encoder's output on the reconstructed spectra $E(D(E(\mathbf{R}_-)))$ to its original input $E(\mathbf{R}_-)$. This is a typical loss function in any autoencoder framework, which ensures the encoder-decoder system retains a reasonable data integrity level when transforming back and forth.

Neural-network optimisation was performed using AdamW ($\beta_1 = 0.9$, $\beta_2 = 0.999$, weight decay = 0.01). The encoder uses a fixed learning rate of 1×10^{-4} , while the decoder uses 1×10^{-5} for the 'PINN-default', 'PINN-both', and 'PINN-plate' configurations and 1×10^{-4} for 'PINN-tran'. Training proceeds over 20,000 iterations, using mini-batches of size 420 sampled from 20 labelled and 400 unlabelled spectra. Convergence is assessed via stabilisation of the total training loss. Across replicates, training was stable and consistent. Representative loss curves are shown in Fig. S8.

2.3. Assessing accuracy of PINN by cross comparison

To evaluate the performance of the PINNs in trait inversion from observed spectral data, we compared our trained PINNs to two benchmark methods: PLSR (Gerlach et al., 1979) and a Bayesian RTM inversion via MCMC (Shiklomanov et al., 2016), using the independent validation datasets (Fig. 2). The PINNs and PLSR models were trained on only 596 of 3451 spectra (~17.3 %), specifically from the LOPEX and ANGERS datasets—the same empirical datasets used to originally calibrate PROSPECT5B (Féret et al., 2008)—ensuring that all methods (RTMs, PLSRs, and PINNs) were trained with an equal amount of labelled data. All remaining datasets were withheld for testing, enabling evaluation of out-of-distribution generalization.

2.3.1. Partial least squares regression

The PLSR model was trained and validated on the same datasets as the PINN. We applied Repeated Double Cross-Validation (RdCV), which is a method designed to provide a robust assessment, reducing the risk of overfitting, as a standard cross-validation method (Dechant et al., 2017; Liu et al., 2023; Yan et al., 2021). RdCV involves repeatedly splitting a dataset into training and testing sets in a two-tier 'double' cross-validation approach, as described by (Filzmoser et al., 2009). In the first tier, we divided the data 100 times (ratio 25:75 for validation vs training) to optimize model complexity. Within each of these splits, a second inner 2-fold cross-validation was applied to evaluate model error. This process resulted in 100 distinct models, each optimized for its specific validation samples. The final ensemble output was derived by averaging these 100 models. RdCV allows for a more comprehensive evaluation of model reliability. Notably, all the validation samples for our RdCV were sampled in the training set.

2.3.2. Bayesian model inversion

We employed a Bayesian framework for spectral inverse modelling (Shiklomanov et al., 2016), which identifies the posterior distribution of

a parameter set that minimizes the discrepancy between the reflectance predicted by the PROSPECT5B model and measured reflectance on all validation datasets. The joint posterior distributions of all parameters were sampled using Monte-Carlo Markov Chain (MCMC) methods with an Affine Invariant MCMC Ensemble sampler implemented in the EMCEE packages (Goodman and Weare, 2010). This algorithm, with adaptive scaling, automatically adjusts step sizes for improved exploration in various parameter directions, particularly in high-dimensional parameter spaces, enabling better consideration of complex parameter relationships. We assumed a Gaussian likelihood function with an additional variance parameter σ^2 . Each inversion was initialized with random parameter values from predefined prior distributions and executed in parallel for 20 independent chains. As recommended by the developers of EMCEE (Foreman-Mackey et al., 2013), we assessed "convergence" (the production of independent samples from the posterior) through the integrated autocorrelation time (τ) for each parameter. We ensured that the sampler was run for more than $50 \times \tau$ steps and discarded the first $5 \times \tau$ as burn-in, retaining effectively independent samples by thinning by τ .

2.3.3. Validation by cross comparison

We used the independent validation leaf reflectance spectra to predict leaf optical traits using each method (Fig. 2). For the PINN and PLSR models, which were trained on the same datasets, this meant that traits were predicted by the models on the independent validation datasets (Table 1). For the Bayesian RTM inversion, this meant that PROSPECT5B was directly inverted on the validation datasets. We evaluated the accuracy of each method's predictions by calculating the Root Mean Squared Error (RMSE). Note that during training, the PINN model additionally used unlabelled spectral data, leveraging it for semi-supervised learning to enhance its performance.

To facilitate an intuitive comparison of prediction performance across different traits, we normalized the traits data, capping the maximum value at three times the size of the observed traits in the Lopex dataset.

To evaluate model uncertainty, we used five different random seeds for the PINN model. The PLSR model produced 100 parallel prediction results which we used to quantify uncertainty, while for the MCMC method we calculated the mean and standard deviation (SD) of the posterior distribution for each parameter to quantify uncertainty.

Considering that inference speed is crucial for real-world applications, especially in big data scenarios, we assessed its computation time on the validation set and compared the performance with the above benchmark methods. Speed benchmarks were done on computational hardware consisting of an NVIDIA RTX 3090 GPU and an Intel(R) Core (TM) i7-10700KF CPU @ 3.80GHz.

2.4. Model component replacement evaluation

Following the PINN default, we developed three variations of leaf RTM decoders, which we refer to as "module replacements". These modifications were intentionally designed to be minimal, ensuring they seamlessly interact and align with the overarching sequence of algorithms and biophysical mathematical principles inherent in the PROSPECT5B model. Specifically, we substituted the transmission and generalized plate model components, as well as their combinations (transmission + plate) with CNNs, assessing them against sequences of their inputs (see Eq. (3) & (5)). The rationale behind using CNNs is their ability to share weights across wavelengths through convolutional kernels (i.e. filters): the same filter, with consistent weights and biases, is uniformly applied throughout the input field. In essence, CNNs execute a uniform mathematical operation across all wavelengths, which is analogous to physical equations, which also apply a consistent mathematical operation across all wavelengths. This facilitates a direct comparison and physical interpretation alongside the original physical models. If our CNN-based replacements enhance the model's performance, it suggests

an opportunity to refine and possibly advance the foundational principles of these existing models.

To reflect the reduction or increase of physical constraints, we introduced a dynamic parameter that adjusts the contribution of physical laws in the model. This allows the model to either strictly adhere to the physical constraints or to relax them in favour of more flexible, data-driven predictions. Specifically, the final output of the module as a weighted sum of two components: one derived from the physical equations and the other from a CNN. This weight is determined by a learnable hyperparameter (physical ratio, w), which ranges from 0 to 1 for the physical component, with the CNN component weighted as $1 - \text{physical ratio}$. This design effectively merges physics-based approaches with machine learning techniques, providing a quantitative measure of how much the model deviates from physical laws to improve predicted power. When $w = 0$ it indicates that the physical unit is entirely replaced by the CNN, while $w = 1$ signifies the full physical model. Recall that the loss functions $L_{\text{decoder reconstruction}}$ and $L_{\text{module replacement}}$ are informative here.

2.5. Assessing deviations between physical models in PROSPECT5B and CNN replacements

We examined the impact of replacing and adjusting modules within the PINN's decoder to understand how these changes influenced its predictive performance compared to the original PROSPECT5B decoder. When these modifications led to improved performance, we visualized these changes to elucidate how the CNN-based replacements affected the calibration and physical predictions of the modules. This approach allowed us to identify areas where predictions diverged. For each scenario, we conducted a visual comparison of the response curves various inputs, highlighting discrepancies between the predictions of the PINN and the original model components. These visual comparisons serve as an initial step toward targeted refinements of the PROSPECT5B model.

3. Results

3.1. Assessing the accuracy of PINN

The PINN default, which employs an autoencoder with PROSPECT5B as the decoder (D_{rtm}), demonstrated superior predictive accuracy compared to the PLSR model and the Bayesian RTM model inversion. This predictive improvement is particularly visible in the reduction of error across all validation datasets (Fig. 3). The Observed-Predicted plots in Fig. 3 illustrate the lower error in trait predictions achieved by PINNs. When analyzing specific leaf traits, PINNs outperformed both PLSR (RMSE = 0.034 vs 0.038, reduced by 12 %) and MCMC in nearly all cases (RMSE = 0.034 vs 0.045, reduced by 24 %). Only for C_w , PINNs and PLSR exhibited comparable performance, both of which surpassed the performance of the MCMC model inversion. Especially for C_m , the MCMC model inversion showed extreme predicted values (RMSE = 0.042), which means that the inversion capabilities of pure RTMs are limited for some sites. PLSR predicts some values below 0, showing that data-driven approaches may predict physical impossibilities (Fig. 4). Some methods also consistently over or underestimated trait values (i.e. regression lines consistently off the 1 to 1 line in Fig. 4). This is best reflected in the MAE scores in Fig. 4, which are a more direct measure of bias than RMSE. The PINN Default showed the lowest MAE over all traits (compare e.g. the MAE in Fig. 4 c & d to g & h or k & l).

Across the different validation datasets, PINNs consistently outperformed MCMC-based RTM inversion methods, and outperformed PLSR for C_{ab} , C_{ar} , and C_m at nearly all datasets. The main exception was C_w , where PINNs outperformed PLSR in only 2 of the 5 validation datasets. To maintain clarity, we present summary results in the main text; full site-specific performance metrics across all traits, datasets, and methods are provided in Fig. S5.

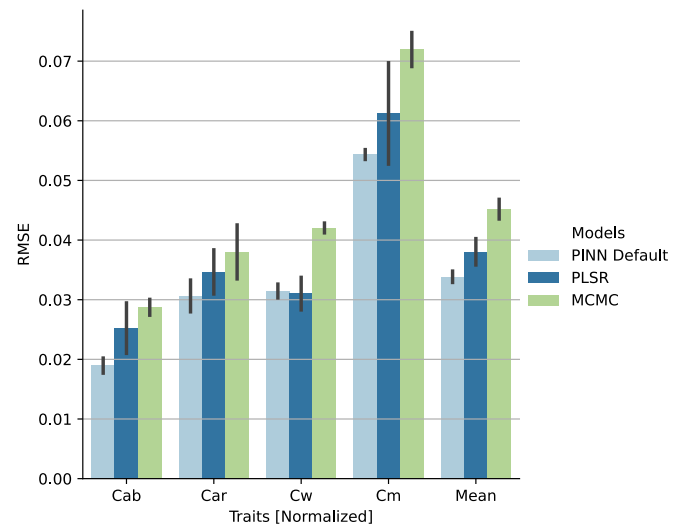


Fig. 3. Root Mean Square Error (RMSE) in leaf trait predictions across five validation datasets for PINN, PLSR, and MCMC approaches. All models were trained on 596 of 3451 spectra (~17.3 %) and validated on the remaining 2855 spectra (~82.7 %). “PINN Default” uses Prospect5B as the decoder (denoted D_{rtm}). “PLSR” refers to Partial Least Squares Regression, and “MCMC” to a Bayesian inversion of Prospect5B. Note that while the MCMC inversion does not require training, Prospect5B was originally calibrated using the same dataset applied in training PINNs and PLSR (Féret et al., 2008). PINNs consistently outperform the benchmark methods across trait.

3.2. Model component replacement evaluation

Substituting the leaf transmission, the generalized plate model, or both with CNNs ($D_{\text{r}}(T)$, $D_{\text{pm}}(T)$ and $D_{\text{rpm}}(T)$, respectively) did not consistently enhance predictions for specific traits. Instead, these substitutions produced varied outcomes of overall RMSE reduction (see Fig. S.1 and Fig. S.2 for spectra reconstruction and Fig. S.3 for pairwise relationships of PINNs predictions). The PINN Plate with $D_{\text{pm}}(T)$ gave the overall best improvement (for Cab, Car, Cw and Cm). Generally, replacing only the generalized plate model with a CNN resulted in a noticeable performance boost (decrease of RMSE, Fig. 5 & Table S.2). However, no single substitution led to uniform improvements across all metrics—the observed RMSE range across different decoder versions encompassed the RMSE of the default decoder (PROSPECT5B). In general, all of the substitution PINNs and especially $D_{\text{pm}}(T)$ showed notable improvements in predicting leaf chlorophyll $a + b$ (C_{ab}) and carotenoid (C_{ar}). Enhancements from $D_{\text{r}}(T)$ and the $D_{\text{rpm}}(T)$ came at the expense of reduced predictive accuracy for leaf water content (C_w), as the $D_{\text{r}}(T)$ showed a 2.1 % increase in RMSE, while $D_{\text{rpm}}(T)$ showed a 6.2 % increase compared to $D_{\text{r}}(T)$ with the $D_{\text{rpm}}(T)$ demonstrating more significant improvements in leaf dry-mass (C_m ; (0.053, equalling a 2.2 % reduction in RMSE relative to $D_{\text{r}}(T)$). Specifically, $D_{\text{pm}}(T)$ gave predictions of C_w more concentrated along the 1:1 line, whereas for the prediction of $D_{\text{r}}(T)$, there was a noticeable deviation in the C_m predictions for certain plots (Fig. 6).

3.3. Assessing deviations between physical models in PROSPECT5B and CNN replacements

We assessed the differences in predictions between the simplified physical models in PROSPECT5B and their CNN replacements. Our goal was to derive insights from the CNN predictions that could help determine components for potential improvement for PROSPECT5B. Upon integrating CNNs into the decoder, we observed that the PINN predominantly favoured the physical models, assigning them over 61 % of the weight (0.620 for $D_{\text{pm}}(T)$ and 0.617 for $D_{\text{rpm}}(T)$, except for the transmission module replacement ($D_{\text{r}}(T)$), which received a 25.6 %

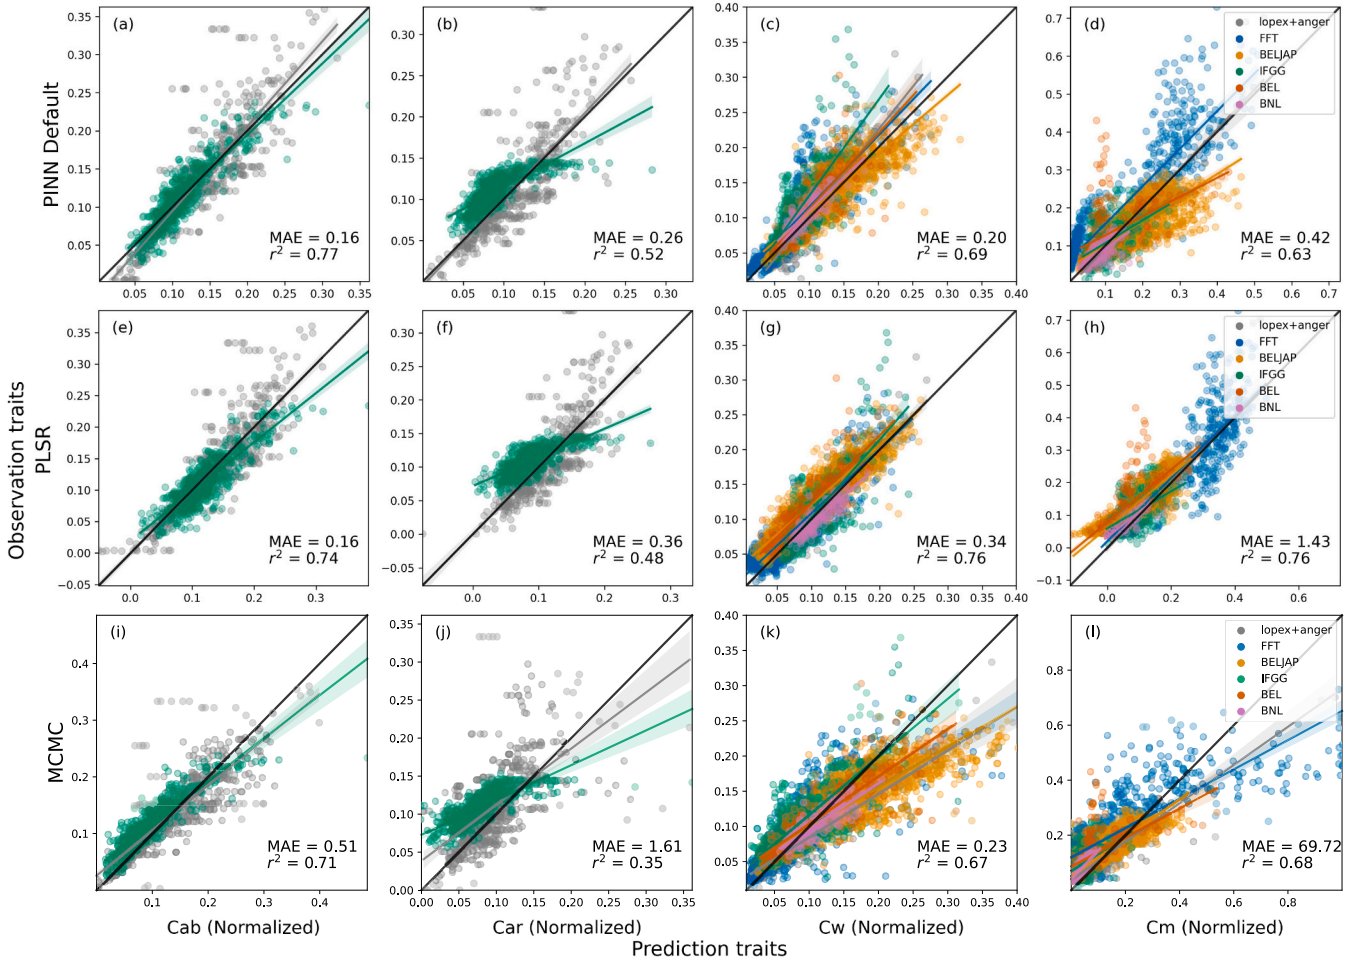


Fig. 4. Observed (normalized) trait values against predicted values for training (grey dots) and validation datasets (coloured dots) for PINN vs benchmark models of PLSR and MCMC. Here, r^2 represents the coefficient of determination, and MAE represents the Mean Absolute Error.




Fig. 5. Change in RMSE for trait predictions when replacing modules of Prospect5B with CNN-based surrogates, relative to the PINN Default. Traits include chlorophyll (Cab), carotenoids (Car), equivalent water thickness (Cw), and carbon mass per area (Cm). Replacing the transmission function (PINN Tran) yielded no improvement, while substituting the plate model (PINN Plate) led to the largest gains—particularly for Car. Replacing both components (PINN Both) did not improve performance overall.

weighting.

By replacing the transmission model (Eq. (3)) in the PROSPECT5B physics-based decoder with a CNN ($D_r(T)$), we observed the most significant deviation from the PROSPECT5B baseline. This substitution indicated a slower decrease in electromagnetic transmission for an equivalent total transmission coefficient (k , as defined by Eq. (4)) (Fig. 7).

Replacing the plate model (Eq. (5)) in the PROSPECT5B physics-based decoder with a CNN ($D_{pm}(T)$) also showed a notable shift from the PROSPECT5B standard. This alteration produced more distinct predictions, showing a stronger response to shifts in the number of leaf cell-air interfaces (N) or leaf diffractive index. The observed deviation suggests differences in the representation of light scattering or internal leaf structure, although the analytical complexity of Eq. (5) precludes straightforward analysis of potential effects.

3.4. Computational performance of PINN default design

PINNs demonstrated a substantial speed advantage over traditional model inversion techniques. Once trained by the encoder, they could predict traits across all 2855 validation spectra in 447 microseconds. This is more than 4.8 million times faster than the MCMC inversion, which takes 25 days to complete ($2,160,000 \text{ s} / 0.000447 \text{ s}$). This remarkable speed makes PINNs computationally comparable to the PLSR model with 100 ensembles, which takes 3.85 s, and makes PINNs 8.6 thousand times faster in inference ($3.85 \text{ s} / 0.000447 \text{ s}$). Note that this speed up is enabled by refactoring the RTM code for CUDA




Fig. 6. Observed (normalized) trait values against predicted values for training (grey dots) and validation datasets (coloured dots) for different configurations of PINN. Here, r^2 represents the coefficient of determination, and MAE represents the mean absolute error.




Fig. 7. Total leaf transmittance versus the total absorption, accounting for the density of all leaf biochemical traits. The graph displays predicted transmittance lines for different models: Prospect5B (blue), and the final PINN model used in the decoder ($D_{\tau}(T)$; red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

compatibility, a step that was needed for the integration of deep learning models (such as CNNs) and RTMs (Luebbe, 2008), which enables parallel processing across thousands of GPU cores.

4. Discussion

In this study, we trained Physics-Informed Neural Networks (PINNs) to predict leaf traits from spectral data. PINNs yielded higher predictive accuracy than Radiative Transfer Model (RTM) inversion and Partial Least Squares Regression (PLSR), particularly in reducing bias (MAE), even when trained on a limited dataset ($\sim 17\%$). They also enabled prediction at speeds comparable to PLSR while vastly outpacing MCMC inversion. We further tested how replacing specific PROSPECT5B modules with CNNs affected performance and found that substituting the generalized plate model ($D_{\text{pm}}(T)$) provided the largest gain. Not all replacements improved performance equally, underscoring that trait- and module-specific factors shape potential benefits. In the following discussion, we will delve deeper into the reasons behind PINNs' performance, compare across different PINN configurations and classical methods, evaluate computational efficiency, and explore how PINNs may help inform the future evolution of RTMs.

4.1. Explaining accuracy improvements of PINNs

Our experiments revealed that the PINN default, utilizing an autoencoder with PROSPECT5B as the decoder (D_{rtm}), provided better

accuracies than both the PLSR model and the physical model inversion using MCMC methods, significantly reducing bias in all leaf trait predictions (PLSR and MCMC showed 12.4 % and 33.7 % higher RMSE values in mean trait estimates compared to the PINN default; Figs. 3 & 4, Table S2). This section highlights the major differences and similarities between the PINNs, PLSR, and physical model inversion methods, emphasizing the distinct flexibility and robustness of PINNs.

The semi-supervised autoencoder architecture of PINNs is conceptually similar to PLSR as both methods handle numerous, highly correlated predictor variables with relatively few independent samples by constructing new lower-dimensional predictor variables. However, PLSR's dimensionality reduction is linear via Principal Component Analysis (PCA), while our NN encoder uses non-linear activation functions, stacked convolutional and pooling layers in the encoding. This means that the encoder's dimensionality reduction can be seen as a hierarchical non-linear extension of PLSR's PCA dimension reduction (Goodfellow et al., 2016). Hence, the first clear benefit of our PINN approach is that while PLSR constructs these components as linear combinations of the original variables, the autoencoder used here may construct non-linear combinations, offering greater flexibility but increasing the risk of overfitting (Goodfellow et al., 2016).

In PINNs, the risk of overfitting is offset by the forced integration of the RTM modules in the decoder. This is the second demonstrated benefit of the PINN approach. Typically, effective cross-validation to assess and mitigate overfitting requires large ground truth observational datasets. Classically, the availability of such large ground truth datasets is limited (Todman et al., 2023), which typically leads to sacrificing most data to training of the model. This risk comes at a cost of reduced generality, especially for global studies with vastly larger scales of prediction or extrapolation compared to training. In contrast, a physical model as decoder forces adherence to physical rules, which by itself helps prevent retention of irrelevant information in the encodings. This allowed us to train the PINN on only 17 % of the (ground truth) data, thus offering a promising solution to criticisms on overfitting and extrapolation issues in remote sensing (Meyer et al., 2018).

In contrast to data-driven approaches, RTM inversion methods often simplify complex phenomena, leading to "ill-posedness" where multiple parameter combinations can produce similar spectral characteristics, reducing predictive precision (Atzberger, 2004; Darvishzadeh et al., 2008). Our modular PINN framework enables selective replacement of specific components or equations (e.g. the plate model). This allows us to retain physical constraints where appropriate while improving flexibility and model capacity - offering one pathway to mitigate problems caused by ill-posedness, incomplete mechanistic assumptions, and limited calibration datasets. Nevertheless, our implementation - like all standard PINNs - inherits both the strengths and limitations of the embedded physical model. For instance, physical models like PROSPECT5B rely on in situ observed and lab-measured data for calibrating empirical constants such as specific absorption coefficients (SAC, Eq. 4) (Féret et al., 2008; Jacquemoud and Baret, 1990). When calibrated on limited datasets, these models may retain site or region specificity, making extrapolation difficult (Alberts and Bilionis, 2023; Martin and Schaub, 2022; Willard et al., 2022). Hence, radiative transfer models like PROSPECT5B are also not immune to "overfitting" problems, and may still retain some site or region specificity, which may complicate extrapolation (Alberts and Bilionis, 2023; Martin and Schaub, 2022; Willard et al., 2022).

In our specific application, we believe that the default PINNs' ability to overcome the above limitations is for a large part enable by their use of unlabelled data in a self-reflective learning process typical of autoencoders. Unlabelled spectral data—as leaf spectra without corresponding trait measurements—capture broader environmental variation, including rare or underrepresented conditions. This helps the encoder generalize beyond the potential site-specific calibration domain of the RTM by refining trait predictions based on how well the reconstructed spectra match the inputs. In doing so, the model not only learns

to respect the constraints of the RTM but also becomes more robust to site- or context-specific biases in the physical model. Hence, the autoencoder framework offers a practical compromise coming in PINNs: combining a rigid physical structure with data-driven flexibility of using unlabeled spectra.

We note that an important future direction is to explicitly incorporate measurement uncertainty into the PINN framework. Zerah et al. (2024) recently addressed this by combining PINNs with a variational autoencoder, but their approach should be extended by integrating physically grounded noise models (Huang et al., 2013). Spectral reflectance measurements are inherently noisy, with variance arising from both signal-dependent and signal-independent sources. Incorporating sensor-informed noise characterizations and explicitly modelling uncertainty within the PINN architecture has benefits as it allows distinguishing true trait variability from measurement artefacts. We believe that this would enhance robustness, particularly under low signal-to-noise conditions or across heterogeneous sensor platforms.

4.2. Advancing radiative transfer modelling using PINNs

In addition to evaluating predictive performance, we used submodule replacement as a diagnostic tool to identify which components of PROSPECT5B most constrain model accuracy. These replacements help pinpoint where simplified assumptions in the RTM may be limiting performance and suggest pathways for physical model refinement. In addition, the outputs of the CNN replacements remain physically interpretable: Fig. 7 shows how predicted transmission varies with absorption by biochemical leaf traits (Eq. 3), revealing a clear departure from the simplified hemispherical integration in Beer's law. While revisions to the original equations are beyond the scope of this study, this divergence suggests concrete directions that - if proven robust - could inspire improvements in RTM structure, which we discuss further below.

We show that especially the replacement of the plate module (Dr(T)) improved overall predictive accuracy across all four traits. For instance, the CNN replacement predicted lower asymptotic reflectance values as refractive index increased (Fig. 8, Fig. S.4), indicating a more gradual sensitivity to changes in internal leaf structure - essentially predicting that increasing leaf thickness - as approximated by the number of air-cell interfaces (N) - doesn't always increase reflectance. This suggests that the original plate model—based on the Fresnel equations and perhaps the assumptions of layer homogeneity—may oversimplify the heterogeneous optical properties of real leaves and how these change with leaf thickness. Empirical studies on cross-sections of leaves, such as those by (Hughes et al., 2007), show that primary and secondary (helper) photosynthetic pigments have a varied spatial distribution, with some pigment types concentrated on the abaxial side and others on the adaxial side. Moreover, the internal structure of the leaf, and hence its heterogeneity, varies strongly across species (Onoda et al., 2011). A distinct heterogeneous nature of leaves, with pigments concentrated in outer layers, is one possible explanation for a stronger saturating effect with increasing air-cell interfaces (Fig. 8. c).

Finally, our PINNs rely on empirically calibrated specific absorption coefficient (SAC) in PROSPECT5B from previous studies. This reliance poses a challenge for extrapolation by the PINN, especially given the variability of SAC across different types of vegetation and conditions (e.g. differences in leaf structure or chemical composition of dry matter). It is possible that the original calibration of these coefficients in earlier studies may have resulted in overfitting, limiting their applicability in broader contexts. Recalibrating the SAC was beyond the scope of this study, which focused on exploring alternative, often overlooked methods for improving PROSPECT5B. We believe that our focus here was relevant as recalibration efforts for Prospect have been conducted multiple times in the past (Féret et al., 2021; Féret et al., 2017; Féret et al., 2008; Jacquemoud and Baret, 1990). In addition, keeping the SACs constant allows us to isolate the impact of structural modifications to the radiative transfer formulation without the added complexities of

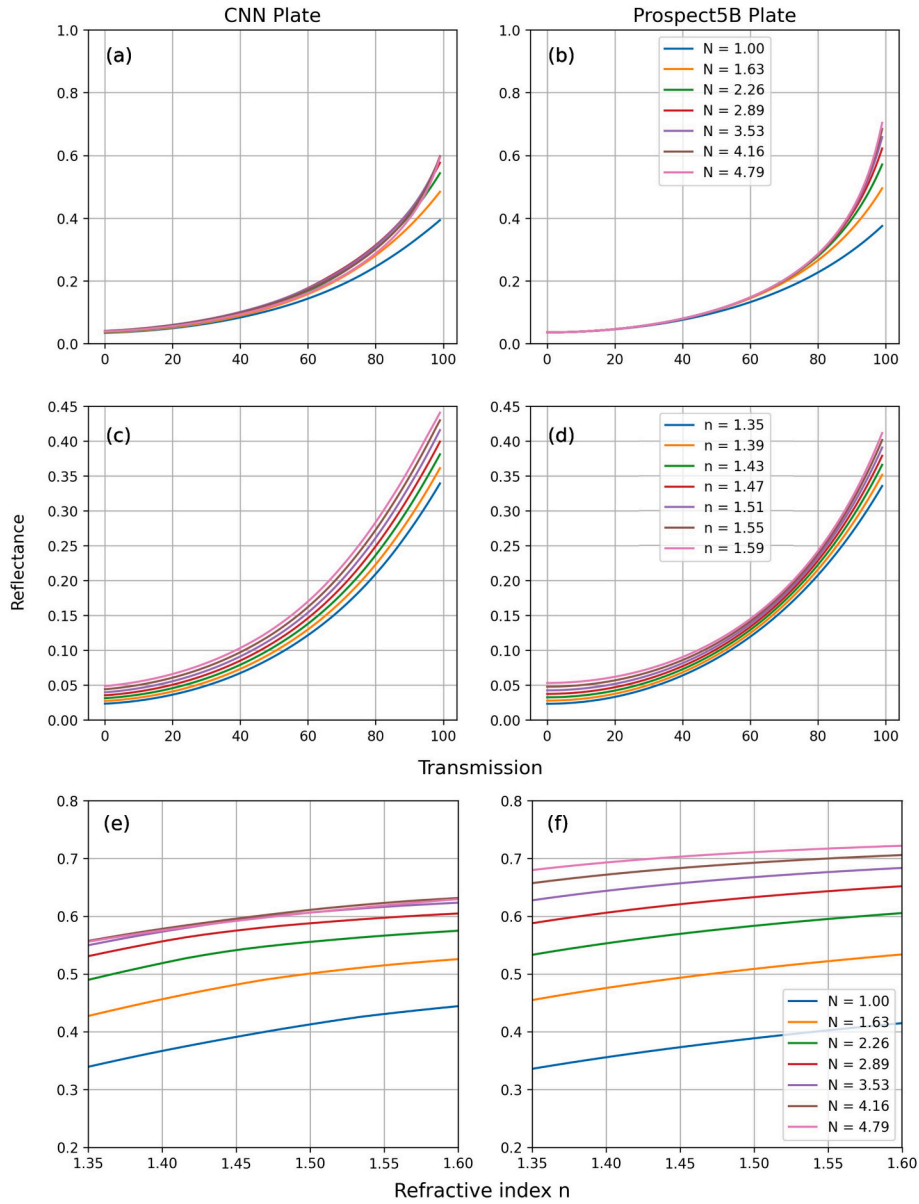


Fig. 8. Transmission vs. reflectance response curves for the generalized plate model and CNN replacements in the decoder (pure CNN, and the Prospect5B plate model Dpm(T)), using the designed transmission and fixed default values for incident angle as input. Panels and lines show response curves for a range of air-cell interfaces in the leaf (N), at a fixed refractive index $n = 1.46$ (a & b), and at a range of refractive indices (n) at a single air-cell interface ($N = 1$, c & d). Finally, panels e and f show the predicted leaf reflectance as a function of n , at different values of N . In general, the CNN plate model shows a more complex non-linear behaviour compared to PROSPECT, showing a stronger interaction N and transmission (a vs b), being more responsive to changes in n over transmission (c vs d) and a stronger saturating effect (in response N) for reflectance over n (e vs f). We used the standard α values as in PROSPECT5 to calculate all inputs of the plate model.

empirical SAC recalibration. Nevertheless, a follow-up study with a more flexible PINN that allows reasonably varying absorption coefficients should be possible. We expect that this will further enhance the capabilities of the PINN models introduced here, as well as the relevance and predictive ability of the replacement modules.

4.3. The evolution of RTM inspired by CNN integration

Building on the diagnostic results of Section 4.2, we propose that physics-informed neural architectures such as PINNs offer a promising framework for evolving radiative transfer models (RTMs). Rather than viewing RTMs as static pipelines, a modular approach enables selective refinement of specific components based on empirical evidence—preserving physical interpretability while adapting where traditional formulations fall short. Many RTM sub-models were originally designed for

analytical tractability, using fixed coefficients or simplified equations to approximate complex biophysical processes. While effective in controlled settings, these simplifications often limit generalizability across vegetation types and conditions. Our results suggest that targeted neural replacements can relax such constraints and highlight which assumptions may require revision.

This modular strategy allows for progressive hybridization. Rather than abandoning physical insight, neural replacements should be introduced only where they improve performance or reveal limitations—making them useful both for prediction and model development. This is especially valuable in complex or poorly understood domains, such as chlorophyll fluorescence, where standard formulations (e.g. Fluspect) use fixed parameters that may not generalize across environmental conditions (Campbell et al., 2008; Vilfan et al., 2016). CNN-based modules could flexibly capture such dynamics, either as

direct replacements or as learnable corrections. Our results suggest that the first venue to start such advance is by replacing the plate model.

Beyond leaf-level modelling, this approach generalizes to more complex models like PROSAIL, where canopy scattering, hotspot effects, and viewing geometry introduce additional layers of simplification (Dehghan-Shoar et al., 2024; Zérah et al., 2024). CNN replacements could be selectively introduced at this level to bridge gaps between canopy- and leaf-scale processes. Even atmospheric radiative transfer components could benefit from learnable corrections based on satellite or tower-based data, allowing a more realistic representation of layer-specific interactions.

We do note that as more RTM components are replaced, interactions between them may introduce complex emergent behaviours that require coordinated training strategies. Future work should develop structured frameworks for multi-component learning, along with hyperparameter optimization strategies tailored to PINNs. While hyperparameter tuning is critical in all neural networks, PINNs introduce additional challenges by requiring a balance between data fitting and physical consistency. This makes them particularly sensitive to architectural choices, loss weighting, and training dynamic. Automated tools for adaptive hyperparameter tuning – particularly those designed for physics-informed settings – could yield substantial performance gains (Jin et al., 2023; Wang et al., 2021).

5. Conclusion

Recent work highlights several major challenges in remote sensing, particularly focusing on the uncritical use of purely data-driven models (Dechant et al., 2017; Meyer and Pebesma, 2022; Todman et al., 2023). These challenges include the difficulty of working with limited datasets, the risk of models lacking generalizability to out-of-sample scenarios, and the potential for producing physically inconsistent results due to the unrepresentativeness of training data. Our study demonstrates that PINNs can help address these issues by embedding mechanistic models within neural networks, supporting consistency with known physical processes while improving predictive accuracy. They also offer a means to test and improve mechanistic models by identifying structural limitations. This integration provides a practical path toward more reliable trait retrieval and a stronger foundation for theory-informed remote sensing science.

CRedit authorship contribution statement

Peng Sun: Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Peter M. van Bodegom:** Writing – review & editing, Supervision, Formal analysis, Conceptualization. **Joris Timmermans:** Writing – review & editing, Conceptualization. **Shuwen Liu:** Writing – review & editing. **Jin Wu:** Writing – review & editing. **Marco D. Visser:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Peng Sun reports financial support was provided by China Scholarship Council. Shuwen Liu reports financial support was provided by National Natural Science Foundation of China. Jin Wu reports financial support was provided by National Natural Science Foundation of China. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Peng Sun reports financial support was provided by China Scholarship Council (#202104910033). Marco Visser acknowledges the CML Impact fund of Leiden University. Shuwen Liu and Jin Wu reports financial support was provided by National Natural Science Foundation of China (#31922090) and the Innovation and Technology Fund (funding support to State Key Laboratory of Agrobiotechnology). If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.rse.2025.114958>.

Data availabilityThe link to the -dataset is provided in Supplementary Material 1 (Table S1).

The link to the -dataset is provided in Supplementary Material 1 (Table S1).

References

- Alberts, A., Bilonis, I., 2023. Physics-informed information field theory for modeling physical systems with uncertainty quantification. *J. Comput. Phys.* 486, 6424. <https://doi.org/10.1016/j.jcp.2023.112100>.
- Allen, W.A., Gausman, H.W., Richardson, A.J., Thomas, J.R., 1969. Interaction of isotropic light with a compact plant leaf. *J. Opt. Soc. Am.* 59, 1376. <https://doi.org/10.1364/JOSA.59.001376>.
- Atzberger, C., 2004. Object-based retrieval of biophysical canopy variables using artificial neural nets and radiative transfer models. *Remote Sens. Environ.* 93, 53–67. <https://doi.org/10.1016/j.rse.2004.06.016>.
- Borowiec, M.L., Dikow, R.B., Frandsen, P.B., McKeen, A., Valentini, G., White, A.E., 2022. Deep learning as a tool for ecology and evolution. *Methods Ecol. Evol.* 13, 1640–1660. <https://doi.org/10.1111/2041-210X.13901>.
- Campbell, P.K.E., Middleton, E.M., Corp, L.A., Kim, M.S., 2008. Contribution of chlorophyll fluorescence to the apparent vegetation reflectance. *Sci. Total Environ.* 404, 433–439. <https://doi.org/10.1016/j.scitotenv.2007.11.004>.
- Cavanagh, H., Mosbach, A., Scalliet, G., Lind, R., Endres, R.G., 2021. Physics-informed deep learning characterizes morphodynamics of Asian soybean rust disease. *Nat. Commun.* 12, 6424. <https://doi.org/10.1038/s41467-021-26577-1>.
- Christin, S., Hervet, É., Lecomte, N., 2019. Applications for deep learning in ecology. *Methods Ecol. Evol.* 10, 1632–1644. <https://doi.org/10.1111/2041-210X.13256>.
- Cord, A.F., Brauman, K.A., Chaplin-Kramer, R., Huth, A., Ziv, G., Seppelt, R., 2017. Priorities to advance monitoring of ecosystem services using earth observation. *Trends Ecol. Evol.* 32, 416–428. <https://doi.org/10.1016/j.tree.2017.03.003>.
- Coveney, P.V., Dougherty, E.R., Highfield, R.R., 2016. Big data need big theory too. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* 374, 20160153. <https://doi.org/10.1098/rsta.2016.0153>.
- Cuomo, S., Di Cola, V.S., Giampaolo, F., Rozza, G., Raissi, M., Piccialli, F., 2022. Scientific machine learning through physics-informed neural networks: where we are and what's next. *J. Sci. Comput.* 92, 88. <https://doi.org/10.1007/s10915-022-01939-z>.
- Darvishzadeh, R., Skidmore, A., Atzberger, C., van Wieren, S., 2008. Estimation of vegetation LAI from hyperspectral reflectance data: effects of soil type and plant architecture. *Int. J. Appl. Earth Obs. Geoinf.* 10, 358–373. <https://doi.org/10.1016/j.jag.2008.02.005>.
- Das, S., Tesfamariam, S., 2022. State-of-the-art review of design of experiments for physics-informed deep learning. *arXiv Prepr.* <https://doi.org/10.48550/arXiv.2202.06416>.
- Dechant, B., Cuntz, M., Vohland, M., Schulz, E., Doktor, D., 2017. Estimation of photosynthesis traits from leaf reflectance spectra: correlation to nitrogen content as the dominant mechanism. *Remote Sens. Environ.* 196, 279–292. <https://doi.org/10.1016/j.rse.2017.05.019>.
- Dechant, B., Kattge, J., Pavlick, R., Schneider, F.D., Sabatini, F.M., Moreno-Martínez, Á., Butler, E.E., van Bodegom, P.M., Vallicrosa, H., Kattenborn, T., Boonman, C.C.F., Madani, N., Wright, I.J., Dong, N., Feilhauer, H., Peñuelas, J., Sardans, J., Aguirre-Gutiérrez, J., Reich, P.B., Leitão, P.J., Cavender-Bares, J., Myers-Smith, I.H., Durán, S.M., Croft, H., Prentice, I.C., Huth, A., Rebel, K., Zaehle, S., Šimová, I., Díaz, S., Reichstein, M., Schiller, C., Brulheide, H., Mahecha, M., Wirth, C., Malhi, Y., Townsend, P.A., 2024. Intercomparison of global foliar trait maps reveals fundamental differences and limitations of upscaling approaches. *Remote Sens. Environ.* 311, 114276. <https://doi.org/10.1016/j.rse.2024.114276>.
- Dehghan-Shoar, M.H., Kereszturi, G., Pullanagari, R.R., Orsi, A.A., Yule, I.J., Hanly, J., 2024. A physically informed multi-scale deep neural network for estimating foliar

- nitrogen concentration in vegetation. *Int. J. Appl. Earth Obs. Geoinf.* 130, 103917. <https://doi.org/10.1016/j.jag.2024.103917>.
- Ely, K.S., Burnett, A.C., Lieberman-Cribbin, W., Serbin, S.P., Rogers, A., 2019. Spectroscopy can predict key leaf traits associated with source-sink balance and carbon-nitrogen status. *J. Exp. Bot.* 70, 1789–1799. <https://doi.org/10.1093/jxb/erz061>.
- Fensholt, R., Hiron, S., Tagesson, T., Ehammer, A., Ivits, E., Rasmussen, K., 2015. Global-scale mapping of changes in ecosystem functioning from earth observation-based trends in total and recurrent vegetation. *Glob. Ecol. Biogeogr.* 24, 1003–1017. <https://doi.org/10.1111/geb.12338>.
- Féret, J.B., François, C., Asner, G.P., Gitelson, A.A., Martin, R.E., Bidet, L.P.R., Ustin, S.L., le Maire, G., Jacquemoud, S., 2008. PROSPECT-4 and 5: advances in the leaf optical properties model separating photosynthetic pigments. *Remote Sens. Environ.* 112, 3030–3043. <https://doi.org/10.1016/j.rse.2008.02.012>.
- Féret, J.B., Gitelson, A.A., Noble, S.D., Jacquemoud, S., 2017. PROSPECT-D: towards modeling leaf optical properties through a complete lifecycle. *Remote Sens. Environ.* 193, 204–215. <https://doi.org/10.1016/j.rse.2017.03.004>.
- Féret, J.-B., Berger, K., de Boissieu, F., Malenkovsky, Z., 2021. PROSPECT-PRO for estimating content of nitrogen-containing leaf proteins and other carbon-based constituents. *Remote Sens. Environ.* 252, 112173. <https://doi.org/10.1016/j.rse.2020.112173>.
- Filzmoser, P., Liebmann, B., Varmuza, K., 2009. Repeated double cross validation. *J. Chemom.* 23, 160–171. <https://doi.org/10.1002/cem.1225>.
- Foreman-Mackey, D., Hogg, D.W., Lang, D., Goodman, J., 2013. emcee: The MCMC Hammer. *Publ. Astron. Soc. Pacific* 125, 306–312. <https://doi.org/10.1086/670067>.
- García-Soria, J.L., Morata, M., Berger, K., Pascual-Venteo, A.B., Rivera-Caicedo, J.P., Verrelst, J., 2024. Evaluating epistemic uncertainty estimation strategies in vegetation trait retrieval using hybrid models and imaging spectroscopy data. *Remote Sens. Environ.* 310, 114228. <https://doi.org/10.1016/j.rse.2024.114228>.
- Gastellu-Etchegorry, J.P., Grau, E., Lauret, N., 2012. DART: a 3D model for remote sensing images and radiative budget of earth surfaces, in: modeling and simulation in engineering. InTech. <https://doi.org/10.5772/31315>.
- Gerlach, R.W., Kowalski, B.R., Wold, H.O.A., 1979. Partial least-squares path modelling with latent variables. *Anal. Chim. Acta* 112, 417–421. [https://doi.org/10.1016/S0003-2670\(01\)85039-X](https://doi.org/10.1016/S0003-2670(01)85039-X).
- Ghaffarian, S., Valente, J., Van Der Voort, M., Tekinerdogan, B., 2021. Effect of attention mechanism in deep learning-based remote sensing image processing: a systematic literature review. *Remote Sens.* 13, 1–22. <https://doi.org/10.3390/rs13152965>.
- Goodfellow, I., Yoshua, B., Aaron, C., 2016. *Deep Learning*. MIT press, Cambridge.
- Goodman, J., Wear, J., 2010. Ensemble samplers with affine invariance. *Commun. Appl. Math. Comput. Sci.* 5, 65–80. <https://doi.org/10.2140/camcos.2010.5.65>.
- Gumiere, S.J., Camporese, M., Botto, A., Lafond, J.A., Paniconi, C., Gallichand, J., Rousseau, A.N., 2020. Machine learning vs. physics-based modeling for real-time irrigation management. *Front. Water* 2, 8. <https://doi.org/10.3389/frwa.2020.00008>.
- Hatfield, P.W., Gaffney, J.A., Anderson, G.J., Ali, S., Antonelli, L., Başgeçmez du Pree, S., Citrin, J., Fajardo, M., Knapp, P., Kettle, B., Kustowski, B., MacDonald, M.J., Mariscal, D., Martin, M.E., Nagayama, T., Palmer, C.A.J., Peterson, J.L., Rose, S., Ruby, J.J., Shneider, C., Streeter, M.J.V., Trickey, W., Williams, B., 2021. The data-driven future of high-energy-density physics. *Nature* 593, 351–361. <https://doi.org/10.1038/s41586-021-03382-w>.
- Helsen, K., Bassi, L., Feilhauer, H., Kattenborn, T., Matsushima, H., Van Cleemput, E., Somers, B., Honnay, O., 2021. Evaluating different methods for retrieving intraspecific leaf trait variation from hyperspectral leaf reflectance. *Ecol. Indic.* 130, 108111. <https://doi.org/10.1016/j.ecolind.2021.108111>.
- Huang, C., Zhang, L., Fang, J., Tong, Q., 2013. A radiometric calibration model for the field imaging spectrometer system. *IEEE Trans. Geosci. Remote Sens.* 51, 2465–2475. <https://doi.org/10.1109/TGRS.2012.2211026>.
- Hughes, N.M., Morley, C.B., Smith, W.K., 2007. Coordination of anthocyanin decline and photosynthetic maturation in juvenile leaves of three deciduous tree species. *New Phytol.* 175, 675–685. <https://doi.org/10.1111/j.1469-8137.2007.02133.x>.
- Jacquemoud, S., Baret, F., 1990. PROSPECT: a model of leaf optical properties spectra. *Remote Sens. Environ.* 34, 75–91. [https://doi.org/10.1016/0034-4257\(90\)90100-Z](https://doi.org/10.1016/0034-4257(90)90100-Z).
- Jacquemoud, S., Ustin, S., 2019. Leaf Optical Properties. Leaf Optical Properties, Cambridge. <https://doi.org/10.1017/9781108686457>.
- Jin, H., Chollet, F., Song, Q., Hu, X., 2023. AutoKeras: An AutoML Library for Deep Learning. *J. Mach. Learn. Res.* 24, 1–6. <https://doi.org/10.5555/3648699.3648705>.
- Lam, F., Peng, X., Liang, Z.P., 2023. High-dimensional MR Spatiotemporal imaging by integrating physics-based modeling and data-driven machine learning: current progress and future directions. *IEEE Signal Process. Mag.* 40, 101–115. <https://doi.org/10.1109/MSP.2022.3203867>.
- Liu, S., Yan, Z., Wang, Z., Serbin, S., Visser, M., Zeng, Y., Ryu, Y., Su, Y., Guo, Z., Song, G., Wu, Q., Zhang, H., Cheng, K.H., Dong, J., Hau, B.C.H., Zhao, P., Yang, X., Liu, L., Rogers, A., Wu, J., 2023. Mapping foliar photosynthetic capacity in sub-tropical and tropical forests with UAS-based imaging spectroscopy: scaling from leaf to canopy. *Remote Sens. Environ.* 293, 113612. <https://doi.org/10.1016/j.rse.2023.113612>.
- Luebke, D., 2008. CUDA: scalable parallel programming for high-performance scientific computing, in: 2008 5th IEEE international symposium on biomedical imaging: from Nano to macro, proceedings. ISBI 836–838. <https://doi.org/10.1109/ISBI.2008.4541126>.
- Martin, J., Schaub, H., 2022. Physics-informed neural networks for gravity field modeling of the earth and moon. *Celest. Mech. Dyn. Astron.* 134, 13. <https://doi.org/10.1007/s10569-022-10069-5>.
- Martínez-Ferrer, L., Moreno-Martínez, Á., Campos-Taberner, M., García-Haro, F.J., Muñoz-Mari, J., Running, S.W., Kimball, J., Clinton, N., Camps-Valls, G., 2022. Quantifying uncertainty in high resolution biophysical variable retrieval with machine learning. *Remote Sens. Environ.* 280, 113199. <https://doi.org/10.1016/j.rse.2022.113199>.
- McGreivy, N., Hakim, A., 2024. Weak baselines and reporting biases lead to overoptimism in machine learning for fluid-related partial differential equations. *Nat. Mach. Intell.* 6, 1256–1269. <https://doi.org/10.1038/s42256-024-00897-5>.
- Medvigy, D., Wofsy, S.C., Munger, J.W., Hollinger, D.Y., Moorcroft, P.R., 2009. Mechanistic scaling of ecosystem function and dynamics in space and time: ecosystem demography model version 2. *J. Geophys. Res. Biogeosci.* 114. <https://doi.org/10.1029/2008JG000812>.
- Meireles, J.E., Cavender-Bares, J., Townsend, P.A., Ustin, S., Gamon, J.A., Schweiger, A. K., Schaepman, M.E., Asner, G.P., Martin, R.E., Singh, A., Schrodt, F., Chlus, A., O'Meara, B.C., 2020. Leaf reflectance spectra capture the evolutionary history of seed plants. *New Phytol.* 228, 485–493. <https://doi.org/10.1111/nph.16771>.
- Meng, Q., Catchpole, D., Skillicorn, D., Kennedy, P.J., 2017. Relational autoencoder for feature extraction. In: Proceedings of the International Joint Conference on Neural Netw, pp. 364–371. <https://doi.org/10.1109/IJCNN.2017.7965877>.
- Meyer, H., Pebesma, E., 2022. Machine learning-based global maps of ecological variables and the challenge of assessing them. *Nat. Commun.* 13, 2208. <https://doi.org/10.1038/s41467-022-29838-9>.
- Meyer, H., Reudenbach, C., Hengl, T., Katurji, M., Nauss, T., 2018. Improving performance of spatio-temporal machine learning models using forward feature selection and target-oriented validation. *Environ. Model Softw.* 101, 1–9. <https://doi.org/10.1016/j.envsoft.2017.12.001>.
- Miralles, P., Thangavel, K., Fulvio Scannapieco, A., Jagadam, N., Baranwal, P., Faldu, B., Abhang, R., Bhatia, S., Bonnart, S., Bhatnagar, I., Batul, B., Prasad, P., Ortega-González, H., Joseph, H., More, H., Morchedi, S., Kumar Panda, A., Zaccaria Di Fraia, M., Wischert, D., Stepanova, D., 2023. A critical review on the state-of-the-art and future prospects of machine learning for earth observation operations. *Adv. Sp. Res.* 71, 4959–4986. <https://doi.org/10.1016/j.asr.2023.02.025>.
- Onoda, Y., Westoby, M., Adler, P.B., Choong, A.M.F., Clissold, F.J., Cornelissen, J.H.C., Díaz, S., Dominy, N.J., Elgart, A., Enrico, L., Fine, P.V.A., Howard, J.J., Jalili, A., Kitajima, K., Kurokawa, H., McArthur, C., Lucas, P.W., Markesteijn, L., Pérez-Harguindeguy, N., Poorter, L., Richards, L., Santiago, L.S., Sosinski, E.E., Van Bael, S. A., Warton, D.I., Wright, I.J., Joseph Wright, S., Yamashita, N., 2011. Global patterns of leaf mechanical properties. *Ecol. Lett.* 14, 301–312. <https://doi.org/10.1111/j.1461-0248.2010.01582.x>.
- Pérez-Cutillas, P., Pérez-Navarro, A., Conesa-García, C., Zema, D.A., Amado-Álvarez, J. P., 2023. What is going on within google earth engine? A systematic review and meta-analysis. *Remote Sens. Appl. Soc. Environ.* 29, 100907. <https://doi.org/10.1016/j.rsase.2022.100907>.
- Persello, C., Wegner, J.D., Hansch, R., Tuia, D., Ghamisi, P., Koeva, M., Camps-Valls, G., 2022. Deep learning and earth observation to support the sustainable development goals: current approaches, open challenges, and future opportunities. *IEEE Geosci. Remote Sens. Mag.* 10, 172–200. <https://doi.org/10.1109/MGRS.2021.3136100>.
- Raissi, M., Perdikaris, P., Karniadakis, G.E., 2019. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- Rao, K., Williams, A.P., Flefil, J.F., Konings, A.G., 2020. SAR-enhanced mapping of live fuel moisture content. *Remote Sens. Environ.* 245, 111797. <https://doi.org/10.1016/j.rse.2020.111797>.
- Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., Prabhat, 2019. Deep learning and process understanding for data-driven earth system science. *Nature* 566, 195–204. <https://doi.org/10.1038/s41586-019-0912-1>.
- Schweiger, A.K., Cavender-Bares, J., Townsend, P.A., Hobbie, S.E., Madritch, M.D., Wang, R., Tilman, D., Gamon, J.A., 2018. Plant spectral diversity integrates functional and phylogenetic components of biodiversity and predicts ecosystem function. *Nat. Ecol. Evol.* 2, 976–982. <https://doi.org/10.1038/s41559-018-0551-1>.
- Serbin, S.P., Wu, J., Ely, K.S., Kruger, E.L., Townsend, P.A., Meng, R., Wolfe, B.T., Chlus, A., Wang, Z., Rogers, A., 2019. From the Arctic to the tropics: multibiome prediction of leaf mass per area using leaf reflectance. *New Phytol.* 224, 1557–1568. <https://doi.org/10.1111/nph.16123>.
- Shiklomanov, A.N., Dietze, M.C., Viskari, T., Townsend, P.A., Serbin, S.P., 2016. Quantifying the influences of spectral resolution on uncertainty in leaf trait estimates through a Bayesian approach to RTM inversion. *Remote Sens. Environ.* 183, 226–238. <https://doi.org/10.1016/j.rse.2016.05.023>.
- Stokes, G.G., 1862. IV. On the intensity of the light reflected from or transmitted through a pile of plates. *Proc. R. Soc. Lond.* 11, 545–556. <https://doi.org/10.1098/rspl.1860.0119>.
- Timmermans, J., Daniel Kissling, W., 2023. Advancing terrestrial biodiversity monitoring with satellite remote sensing in the context of the Kunming-Montreal global biodiversity framework. *Ecol. Indic.* 154, 110773. <https://doi.org/10.1016/j.ecolind.2023.110773>.
- Todman, L.C., Bush, A., Hood, A.S.C., 2023. 'Small data' for big insights in ecology. *Trends Ecol. Evol.* 38, 615–622. <https://doi.org/10.1016/j.tree.2023.01.015>.
- Verrelst, J., Camps-Valls, G., Muñoz-Mari, J., Rivera, J.P., Veroustraete, F., Clevers, J.G. P.W., Moreno, J., 2015. Optical remote sensing and the retrieval of terrestrial vegetation bio-geophysical properties - a review. *ISPRS J. Photogramm. Remote Sens.* 108, 273–290. <https://doi.org/10.1016/j.isprsjprs.2015.05.005>.
- Vilfan, N., van der Tol, C., Muller, O., Rascher, U., Verhoef, W., 2016. Fluspect-B: a model for leaf fluorescence, reflectance and transmittance spectra. *Remote Sens. Environ.* 186, 596–615. <https://doi.org/10.1016/j.rse.2016.09.017>.
- Visser, M.D., Detto, M., Meunier, F., Wu, J., Foster, J.R., Marvin, D.C., van Bodegom, P. M., Bongalov, B., Nunes, M.H., Coomes, D., Verbeeck, H., Guzmán, Q., Sanchez-Azofeifa, A., Chandler, C.J., van der Heijden, G.M.F., Boyd, D.S., Foody, G.M.,

- Cutler, M.E.J., Broadbent, E.N., Serbin, S.P., Schnitzer, S., Rodríguez-Ronderos, M. E., Sterck, F., Medina-Vega, J.A., Pacala, S.W., 2025. When can we detect lianas from space? Toward a mechanistic understanding of liana-infested forest optics. *Ecology* 106, e70082. <https://doi.org/10.1002/ecy.70082>.
- Wang, S., Wang, H., Perdikaris, P., 2021. Learning the solution operator of parametric partial differential equations with physics-informed DeepONets. *Sci. Adv.* 7, eabi8605. <https://doi.org/10.1126/sciadv.abi8605>.
- Willard, J., Jia, X., Xu, S., Steinbach, M., Kumar, V., 2022. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Comput. Surv.* 55, 1–37. <https://doi.org/10.1145/3514228>.
- Yan, Z., Guo, Z., Serbin, S.P., Song, G., Zhao, Y., Chen, Y., Wu, S., Wang, J., Wang, X., Li, J., Wang, B., Wu, Y., Su, Y., Wang, H., Rogers, A., Liu, L., Wu, J., 2021. Spectroscopy outperforms leaf trait relationships for predicting photosynthetic capacity across different forest types. *New Phytol.* 232, 134–147. <https://doi.org/10.1111/nph.17579>.
- Yang, P., van der Tol, C., Yin, T., Verhoef, W., 2020. The SPART model: a soil-plant-atmosphere radiative transfer model for satellite measurements in the solar spectrum. *Remote Sens. Environ.* 247, 111870. <https://doi.org/10.1016/j.rse.2020.111870>.
- Zérah, Y., Valero, S., Inglada, J., 2024. Physics-constrained deep learning for biophysical parameter retrieval from Sentinel-2 images: inversion of the PROSAIL model. *Remote Sens. Environ.* 312, 114309. <https://doi.org/10.1016/j.rse.2024.114309>.