PRELIMINARY PREPRINT VERSION – DO NOT CITE
This is not the version of record. The final version will appear in the Proceedings of IJCAI-ECAI 2026.

# NeuroDALEC: A Differentiable and Interpretable Mass-Conserving Framework for Terrestrial Ecosystem Carbon Cycle Dynamics

**Meng Wan**[1,2†], **Tiantian Liu**[1,3†], **Zhixin Xia**[4], **Ningming Nie**[1,3], **Jue Wang**[1,3*], **Rongqiang Cao**[1,3], **Honglin He**[3,5*], **Xiaoli Ren**[3,5], **Peng Shi**[2], **Yangang Wang**[1,3]
[1]Computer Network Information Center, Chinese Academy of Sciences, Beijing, China
[2]National Center for Materials Safety, University of Science and Technology Beijing, Beijing, China
[3]University of Chinese Academy of Sciences, Beijing, China
[4]North China Electric Power University, China
[5]Institute of Geographic Sciences and Natural Resources Research, Beijing, China
wanmengdamon@cnic.cn, liutiantian21a@mails.ucas.ac.cn, zhixinxia@ncepu.edu.cn,
{nienm,wangjue,caorq,wangyg}@sccas.cn,
{hehl,renxl}@igsnrr.ac.cn, pshi@ustb.edu.cn



## Abstract

Accurate simulation of terrestrial ecological carbon cycles is crucial for global climate change and ecosystem management. Process-based carbon models have high interpretability, but suffer from insufficient accuracy and slow computation due to fixed parameters. In contrast, deep-learning carbon models achieve high accuracy, but disregard physical principles, which prevents ecologists from explaining ecosystem dynamics. We propose NeuroDALEC, an interpretable framework that embeds the DALEC carbon-cycle model within a neural network, enabling differentiable computation of ecological processes. Key parameters and ensemble learning strategies are designed, and mass-conserving carbon pool state transition equations are introduced to ensure physical consistency. Experiments show NeuroDALEC outperforms existing models in both accuracy and efficiency. Moreover, it provides sufficient interpretability by predicting all components of the carbon cycle. Deployed in a real-time carbon assimilation system, NeuroDALEC supports daily carbon forecasting and decision-making. This work contributes to the United Nations' Sustainable Development Goals 13 (Climate Action) and 15 (Life on Land). The source code is available at: https://github.com/codesiena/NeuroDALEC.

## 1 Introduction

Driven by global climate mitigation and carbon neutrality targets, terrestrial ecosystems play a key role in regulating the global carbon balance by absorbing and storing atmospheric carbon dioxide [Jain and Mitra, 2025]. As important components of terrestrial ecosystems, forests, grasslands and other vegetation types together constitute the core process of terrestrial carbon sinks through vegetation growth, litter inputs, and soil carbon accumulation [Cao and Woodward, 1998; Hari and Tyagi, 2022]. Therefore, reliable modeling and prediction of the terrestrial ecosystem carbon cycle is a key scientific foundation for carbon peaking, carbon neutralization strategies, and ecosystem management [Wang *et al.*, 2021; Song *et al.*, 2025]. However, terrestrial ecosystem carbon dynamics involve multiple coupled carbon pools and complex nonlinear biophysical processes, whose evolution is jointly influenced by climate drivers, ecological processes, and external disturbances. Accurately forecasting net primary productivity (NPP), net ecosystem productivity (NEP), gross primary production (GPP) and carbon stock dynamics therefore remains a central challenge for sustainable ecosystem management and climate action [Ma *et al.*, 2025].

Although some progress has been made in ecosystem modeling, current approaches to the terrestrial carbon cycle face fundamental limitations [Liao *et al.*, 2023]. Firstly, process-based models utilize carbon pool-flux structure and mass conservation, which provides good ecological interpretability [Bernardini *et al.*, 2024]. However, key ecological parameters are fixed and estimated by Bayesian inversion, limiting the ability of time-varying states [Hui *et al.*, 2022]. Moreover, data assimilation and inference are computationally expensive, making real-time prediction and large-scale deployment difficult [Pierson *et al.*, 2022]. Secondly, data-driven deep learning time series models are relatively efficient and accurate, but the carbon turnover process is a black box [Perry *et al.*, 2022]. Without explicit physical and biological constraints, the predictions may violate carbon conservation and ecological principles, especially under extreme climatic conditions [Chen *et al.*, 2024]. The limited interpretability makes it difficult for ecological scientists to use them as reliable tools for ecosystem analysis. Besides, from an application perspective, carbon prediction models must maintain efficiency, stability, and transferability for the long-term run in online systems [Tian *et al.*, 2025]. The above challenges impede the transformation of carbon-cycle models into globally scaled, sustainably operated decision-making tools.

To solve these problems, we propose a differentiable and interpretable mass-conserving framework for terrestrial ecosystem carbon cycle dynamics. By transforming the carbon pool-flux equations of the DALEC model into differentiable operators and embedding them within a neural network, the model achieves adaptive responsiveness to climate-driven variations. To the best of our knowledge, this is one of the first end-to-end framework that embeds the DALEC pool–flux dynamics as a fully differentiable, mass-conserving operator within a neural forecasting model for multi-step prediction. The main contributions are summarised as follows:

- **Ecology-aware parameter learning with physical consistency constraints.** We group and optimise key ecological parameters like turnover rates and litterfall processes to ecological timescales to enhance model interpretability.
- **Differentiable DALEC-ACM Module.** We design the carbon pool and flux equations as a differentiable dynamic layer, embedding it into the time-series backbone network while ensuring carbon mass conservation. This significantly enhances model accuracy and efficiency.
- **Accurate full-component carbon cycle dynamics.** Compared with DALEC, NeuroDALEC improves prediction accuracy by 12% while reducing computational time from the scale of hours to minutes. It also presents the dynamics of all ecological variables.
- **Real-world deployment and social impact.** The model has been integrated into an operational ecosystem carbon assimilation and forecasting system that supports daily prediction at multiple ecological sites. It provides interpretable and auditable outputs for carbon sink assessment, ecosystem management, and climate-related decision support.

This work directly supports the United Nations Sustainable Development Goals SDG 13 (Climate Action) and SDG 15 (Life on Land). By enabling efficient, interpretable, and deployable forest carbon forecasting, the NeuroDALEC supports real-time carbon monitoring and near-term risk assessment. It demonstrates how physics-aware AI can be translated into trustworthy tools for ecological governance, aligning with the principles of AI for Social Good.

# 2 Related Work

## 2.1 Terrestrial Carbon Cycle Process Models

Terrestrial carbon cycle simulation and prediction depends on a model called Process-Based Models (PBMs) [McGuire *et al.*, 2001]. These models can describe the changes of carbon pool and the dynamics of carbon flux driven by weather [Yuan *et al.*, 2024]. Among them, the Data Assimilation Linked Ecosystem Carbon (DALEC) model and its various variants are an important and interpretable inference framework [Chuter *et al.*, 2015]. Generally, we combine DALEC with data assimilation technology, and use Bayesian inversion method like Markov Chain Monte Carlo Method (MCMC) to estimate ecological parameters and initial carbon state from past observation data, and then use it to make future simulation and prediction [Williams *et al.*, 2005; Delahaies *et al.*, 2017]. DALEC–ACM combines Aggregated Canopy Model (ACM), and estimates GPP and evapotranspiration by Leaf Area Index (LAI) and meteorological data, which is simpler than those complicated physiological models [Smallman and Williams, 2019]. It also extends the Ecological and Dynamic Constraints (EDCs) to limit the parameter range to a reasonable biological range [Bloom and Williams, 2015], and the CARDAMOM framework to estimate and verify the large-scale flux with DALEC [Famiglietti *et al.*, 2021; Yang *et al.*, 2022]. However, the standard data assimilation method assumes that the parameters will not change during the prediction period, which leads to the accumulation of errors under unsteady climate interference, and the high computational cost of iterative inversion limits the real-time applications.

## 2.2 Mechanism-Aware Deep Learning

In order to balance the expressive ability of data-driven models and physical consistency, recent research usually adds physical knowledge to deep learning in two directions: one is to use physics-informed neural networks (PINNs) and add the residual or constraint conditions of the control equation to the loss function to ensure consistency [Raissi *et al.*, 2019], which is summarized and applied in the physical information machine learning framework [Karniadakis *et al.*, 2021]. The other is to embed process modules or differentiable physical layers into the network in a structured way, and learn the relationship between parameters and mechanisms in training [Willard *et al.*, 2022; Shen *et al.*, 2023]. In the field of terrestrial carbon cycling, early applications mainly focused on data-driven flux scale expansion research. For example, FLUXCOM used machine learning to extrapolate observation station data to the global scale. But its essence is to establish a static mapping between meteorology and flux, lacking a dynamic mechanism [Jung *et al.*, 2020]. In order to alleviate the physical inconsistency, recent research attempts to introduce ecological constraints into deep learning. Fan et al. explicitly introduced the equilibrium constraint [Fan *et al.*, 2025] in the multi flux joint learning. The knowledge guided framework built by Liu et al. combines process structure priori with high-resolution observations, which significantly improves the quantification ability of the carbon cycle [Liu *et al.*, 2024]. However, most of the existing methods only focus on the numerical fitting at the flux level, ignoring the dynamic evolution of carbon storage state. The lack of constraints in this mechanism not only limits the interpretability of the model, but also leads to the drift phenomenon that violates the ecological law in the long-term rolling prediction.

# 3 Methods

## 3.1 Problem Formulation

Given a historical window of length $L$, the model takes meteorological driver sequences, including air temperature ($\mathrm{T_a}$), photosynthetically active radiation ($\mathrm{PAR}$), relative humidity ($\mathrm{RH}$), and soil water content ($\mathrm{SWC}$), and predicts ecosystem states and carbon-cycle variables over a future horizon of length $H$. Formally, the input is represented as $X \in$

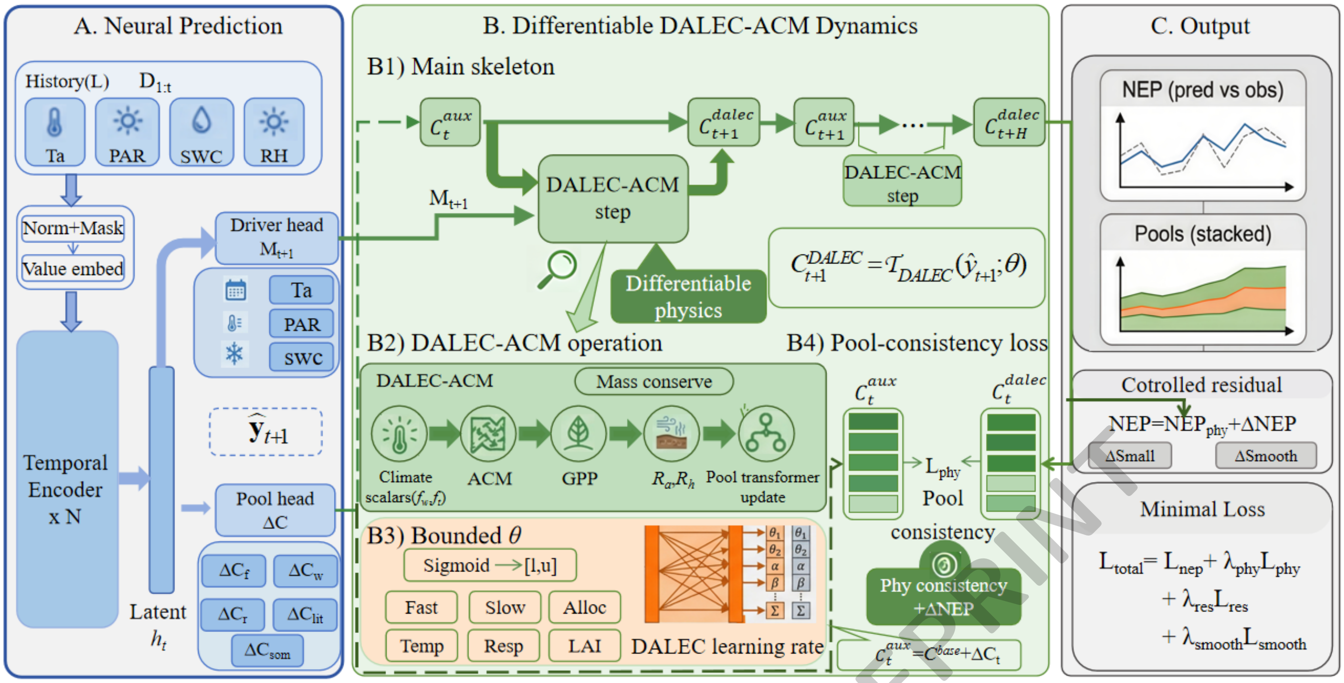


Figure 1: The overall structure of the proposed NeuroDALEC. Three key processes are shown in detail, including the standard flux chain GPP → NPP → NEP and mass-conserving carbon-pool updates.

$\mathbb{R}^{B\times L\times D}$, where $B$ denotes the batch size, $L$ denotes the input length, and $D$ denotes the number of input variables. A time series forecasting backbone maps the input sequence to a multi-channel auxiliary output $\hat{Y} \in \mathbb{R}^{B\times H\times C}$, whose channels are aligned with the subsequent physical layer.

The output channels follow a fixed ordering $[\mathrm{T_a}, \mathrm{PAR}, \mathrm{SWC}, C_f, C_w, C_r, C_{lit}, C_{som}]$, where the first three variables denote future meteorological drivers and the remaining variables correspond to ecosystem carbon pools. Based on these predicted states, a differentiable DALEC–ACM physical module computes carbon fluxes such as GPP, NPP, and NEP. Then carbon pools are updated via mass-conserving balance equations. During training, flux and carbon pool observations are used for supervision, while physical consistency constraints are imposed to ensure ecologically plausible predictions.

### 3.2 The Overall Framework

The overall framework of the proposed NeuroDALEC is illustrated in Figure 1. NeuroDALEC is an end-to-end differential physics guided prediction framework, which aims to embed the ecological mechanism into the neural network, and realize the integration of data-driven and physical processes. Our framework is mainly composed of the following three core processes:

(i) The first is the neural prediction process in Figure 1A. The module encodes the historical meteorological observations (History L) using the time-series backbone network to extract the latent representations $h_t$. Then, the decoder is used to predict the future meteorological driving factor $M_{t+1}$ and the instantaneous deviation $\Delta C$ of each carbon pool state relative to the baseline values. The design reconstructs the current state of the carbon pool, provides accurate initial conditions for the evolution of subsequent physical differential equations.

(ii) The second is the differentiable DALEC–ACM dynamics process in Figure 1B. The B1 main rollout skeleton establishes a time evolution based on differential equations. This module takes $C_t^{aux}$ as the starting point of physical evolution, and uses the differentiable DALEC–ACM operator for multi-step recursive deduction to generate the prediction sequence $C_{t+H}^{dalec}$. The DALEC-ACM operation (B2) encapsulates the complete ecological calculation process, receives driving data to calculate ACM photosynthesis, and strictly calculates respiration and carbon pool turnover flux based on the principle of conservation of mass. Specifically, this physical process follows the standard ecological relation chain GPP $\rightarrow$ NPP $\rightarrow$ NEP, where $\mathrm{NPP} = \mathrm{GPP} - R_a$ and $\mathrm{NEP} = \mathrm{GPP} - R_a - R_h$. The resulting fluxes are then used to update each carbon pool by inflow–outflow balance. At the same time, the B3 bounded module models key ecological parameters as learnable variables and groups them according to the fast/slow ecological process. Through Sigmoid transformation, these parameters are strictly constrained within the biological interval $[l, u]$, enabling the network to learn parameter with physical significance. The B4 pooling consistency loss module calculates the neural prediction carbon pools $C^{aux}$ and physical evolution carbon pools $C^{dalec}$. The difference between them forces neural networks to fit the real physics law.

(iii) The output process is in Figure 1C. The model combines the physical module output and neural residual term

($\Delta$NEP) to generate the final NEP prediction. At the same time, the whole framework is optimized end-to-end through the composite loss function to achieve a balance between the prediction accuracy, physical consistency and regularization constraints. Ultimately, the model not only outputs NEP prediction curves comparing with observed values, but also presents the stacked pools of the five major carbon pools over time, making the dynamics of the ecosystem transparent and visible.

### 3.3 Time Series Forecasting Backbone Module

The time series backbone encodes the history of meteorological drivers and extracts a latent representation of short-term environmental dynamics. At time step $t$, the model takes the driver sequence with standard temporal position encodings:

$$\mathbf{D}_{1:t} = [T_{a,1:t}, \mathrm{PAR}_{1:t}, \mathrm{SWC}_{1:t}, \mathrm{RH}_{1:t}], \tag{1}$$

where $\mathbf{D}_{1:t}$ denotes the concatenated multivariate time series at time $t$. Let $\mathrm{Enc}(\cdot)$ denote the encoding operation of the time-series backbone such as TimeXer [Wang *et al.*, 2024] or iTransformer [Liu *et al.*, 2023], which maps the input sequence to a latent representation:

$$h_t = \mathrm{Enc}(\mathbf{D}_{1:t}), \quad h_t \in \mathbb{R}^{d_h}. \tag{2}$$

The latent state $h_t$ is then mapped by a prediction mapping $\mathrm{Pred}(\cdot)$:

$$\hat{\mathbf{y}}_{t+1:t+H} = \mathrm{Pred}(h_t). \tag{3}$$

Specifically, the output vector at time step $t+1$ consists of predicted meteorological drivers at $t+1$ and auxiliary deviations of carbon pools at the current time step $t$, and is defined as:

$$\hat{\mathbf{y}}_{t+1} = \begin{bmatrix} T_{at+1},\ \mathrm{PAR}_{t+1},\ \mathrm{SWC}_{t+1},\ \Delta C_{f,t},\ \Delta C_{w,t}, \\ \Delta C_{r,t}, \Delta C_{\mathrm{lit},t}, \Delta C_{\mathrm{som},t} \end{bmatrix}. \tag{4}$$

Let $\mathbf{M}_{t+1} = [T_{a,t+1}, \mathrm{PAR}_{t+1}, \mathrm{SWC}_{t+1}]$ denote the predicted meteorological drivers at time $t+1$. Carbon pool $x \in \{f, w, r, \mathrm{lit}, \mathrm{som}\}$ indexes foliage, wood, root, litter, and soil organic matter pools, respectively. For each $x$, the backbone predicts a deviation $\Delta C_{x,t}$ from its long-term site-specific baseline, and reconstructs an auxiliary estimate $\hat{C}^{\mathrm{aux}}_{x,t}$ through the following equation.

$$\hat{C}^{\mathrm{aux}}_{x,t} = C^0_x + \Delta C_{x,t}, \tag{5}$$

where $C^0_x$ denotes the baseline value of pool $x$. This design is motivated by the fact that site-scale carbon pools are stable at year time scales, so we generate short-term deviations to fixed baseline values.

### 3.4 Ecology-Aware Parameter Learning

The DALEC–ACM physical module requires a set of ecological parameters:

$$\theta = \{p_0, \ldots, p_{10}\} \tag{6}$$

where each parameter $p_i$ controls a specific ecological process, such as carbon allocation fractions, turnover rates, or temperature sensitivity coefficients in Table 1.

| Group | Parameters |
|---|---|
| Slow decomposition and soil processes | $p_0,\ p_7,\ p_8$ |
| Autotrophic respiration | $p_1$ |
| Carbon allocation | $p_2,\ p_3$ |
| Fast turnover and transfer | $p_4,\ p_5,\ p_6$ |
| Temperature sensitivity | $p_9$ |
| Leaf area index(LAI)-related process | $p_{10}$ |

Table 1: Grouping of ecological parameters $p_0$–$p_{10}$ by dominant ecological process.

#### Bounded and Differentiable Parameterization

To preserve ecological interpretability and ensure stable training, these parameters of the DALEC–ACM model are designed as learnable variables under biological constraints. Instead of estimating by offline inversion, NeuroDALEC enables adaptive parameter learning within an end-to-end differentiable module.

Each ecological parameter is represented by an unconstrained variable $z_i$ and mapped to a biological range $[l_i, u_i]$ by a sigmoid transformation:

$$p_i = l_i + (u_i - l_i) \cdot \sigma(z_i), \quad i = 0, \ldots, 10, \tag{7}$$

where $l_i$ and $u_i$ denote the lower and upper bounds for parameter $p_i$, which prevents non-physical parameter fluctuation and maintains biological ranges during training.

#### Process-Aware Grouped Optimization

Ecological processes operate on different time scales and impose different sensitivity requirements on the parameters. Therefore, we group the ecological parameters by their dominant ecological functions, as summarized in Table 1.

During training, each parameter group is optimized using a specific learning rate $g$ and regularization strength, which is defined as:

$$\eta_g = \alpha_g\, \eta, \tag{8}$$

where $\eta$ denotes the base learning rate shared by all parameters and $\alpha_g$ is a fixed scaling factor, which controls the update magnitude for group $g$. Parameters associated with faster processes are assigned larger $\alpha_g$, while slower processes are updated more steadily to prevent unrealistic drift in carbon pools.

### 3.5 Differentiable DALEC-ACM Physical Module

Our framework implements the DALEC–ACM as a differentiable physical module. At each time step, this module accepts the predicted meteorological drivers and carbon pool states, denoted as $\hat{\mathbf{y}}_{t+1}$ (derived from Eq. 4). Formally, we define the differentiable transition operator as:

$$\mathbf{C}^{\mathrm{dalec}}_{t+1} = \mathcal{T}_{\mathrm{DALEC}}(\hat{\mathbf{y}}_{t+1}; \theta), \tag{9}$$

where $\theta$ represents ecological parameters. As detailed below, this operator $\mathcal{T}_{\mathrm{DALEC}}$ uses meteorological driver data and carbon pools to infer GPP and allocate fluxes for pool updates $\mathbf{C}^{\mathrm{dalec}}_{t+1}$.

A critical innovation of NeuroDALEC is the execution of these dynamics. Unlike the traditional DALEC–ACM that operates as a standalone iterative simulation, our module applies these physical laws to advance the time-series model's prediction by exactly one time step. By integrating this operator with Ecology-Aware Parameter Learning and loss functions, we enable gradients to flow through the previously rigid biophysical equations.

**Environmental Response Functions**

Environmental controls on photosynthesis and decomposition are implemented as bounded differentiable gates in DALEC–ACM. Given the predicted drivers, these gates output factors constrained to $[0, 1]$ to adjust carbon uptake and heterotrophic decomposition. We use bounding operators like $\max(\cdot)$ and $\mathrm{clip}(\cdot)$ to stabilize long-horizon rollouts, while keeping gate parameters learnable and interpretable. We design a temperature gate strategy for photosynthesis. Temperature limitation with a piecewise-linear response is bounded to $[0, 1]$:

$$f_{t,t+1} = \mathrm{clip}\left(1 - \frac{T_{a,t+1} - T_{\mathrm{opt}}}{T_{\mathrm{lim}} - T_{\mathrm{opt}}},\ 0,\ 1\right), \tag{10}$$

where $T_{a,t+1}$ is air temperature, and $T_{\mathrm{opt}}$ and $T_{\mathrm{lim}}$ denote an optimal and a limiting temperature, respectively.

Following the same bound, we define a soil-water limitation factor $f_{w,t+1} \in [f_{\min}, 1]$ driven by $\mathrm{SWC}_{t+1}$, and a strictly positive decomposition temperature factor $T_{\mathrm{rate},t+1}$ driven by $T_{a,t+1}$ with sensitivity $\gamma$ (where $\gamma \in \theta$). We omit their explicit functional forms for brevity, as they are standard constrained modulators.

**Photosynthesis and Carbon Fluxes**

The ecosystem carbon fluxes are computed through a sequential process in which GPP is first estimated from environmental drivers, followed by autotrophic respiration ($R_a$) and heterotrophic respiration ($R_h$), such that NEP emerges as the net balance of carbon uptake and release. Specifically, the GPP is computed using the ACM module, which links canopy photosynthesis to meteorological drivers and leaf area index (LAI). Specifically, LAI is derived from leaf carbon as

$$\mathrm{LAI}_t = \lambda\, C_{f,t}, \tag{11}$$

where $\lambda$ is a fixed conversion coefficient from foliage carbon to LAI. GPP is then obtained as

$$\mathrm{GPP}_{t+1} = f_{\mathrm{ACM}}(T_{a,t+1}, \mathrm{PAR}_{t+1}, \mathrm{LAI}_t) \cdot f_{t,t+1} \cdot f_{w,t+1}, \tag{12}$$

where $f_{\mathrm{ACM}}(\cdot)$ denotes the ACM photosynthesis function. Then $R_a$ is modeled as a fixed fraction of GPP,

$$R_{a,t+1} = \alpha_{\mathrm{resp}}\, \mathrm{GPP}_{t+1}, \tag{13}$$

where $\alpha_{\mathrm{resp}}$ is a scalar respiration fraction. This yields net primary productivity

$$\mathrm{NPP}_{t+1} = \mathrm{GPP}_{t+1} - R_{a,t+1}, \tag{14}$$

NPP is allocated to leaf, wood, and root pools following standard DALEC allocation schemes. Then $R_h$ is computed from litter and soil pools as

$$R_{h,t+1} = T_{\mathrm{rate},t+1}\left(\kappa_{\mathrm{lit}} C_{\mathrm{lit},t} + \kappa_{\mathrm{som}} C_{\mathrm{som},t}\right), \tag{15}$$

where $\kappa_{\mathrm{lit}}$ and $\kappa_{\mathrm{som}}$ are base decomposition rate coefficients for litter and soil organic matter, respectively. Ecosystem productivity NEP is then given by

$$\mathrm{NEP}_{\mathrm{phy},t+1} = \mathrm{GPP}_{t+1} - R_{a,t+1} - R_{h,t+1}. \tag{16}$$

**Mass-Conserving Carbon Pool Dynamics**

The daily evolution of the five carbon pools is driven by internal fluxes and strictly follows mass balance principles. In DALEC–ACM, the transition is formalized as a compact difference equation:

$$\mathbf{C}_{t+1}^{\mathrm{dalec}} = \underbrace{\mathbf{C}_t^{\mathrm{aux}} + \mathbf{A}_{t+1} - \mathbf{T}(\theta, \mathbf{M}_{t+1}) \odot \mathbf{C}_t^{\mathrm{aux}}}_{\mathcal{F}_{\mathrm{dalec}}(\mathbf{C}_t^{\mathrm{aux}}, \mathbf{M}_{t+1})}, \tag{17}$$

where $\mathbf{A}_{t+1}$ aggregates both NPP allocation and inter-pool transfer fluxes, and $\mathbf{T}(\theta, \mathbf{M}_{t+1})$ denotes environmentally modulated turnover and decomposition coefficients. Pool-wise, this reduces to the standard mass-conservation form

$$C_{i,t+1} = C_{i,t} + F_{i,t+1}^{\mathrm{in}} - F_{i,t+1}^{\mathrm{out}}, \tag{18}$$

where internal transfers between pools cancel in the system-wide sum, so total carbon storage changes only through atmospheric exchange. Accordingly, the net storage change satisfies

$$\begin{aligned} &\sum_i C_{i,t+1} - \sum_i C_{i,t} = \\ &\mathrm{GPP}_{t+1} - R_{a,t+1} - R_{h,t+1} = \mathrm{NEP}_{\mathrm{phy},t+1}. \end{aligned} \tag{19}$$

We refer to $\mathcal{F}_{\mathrm{dalec}}(\cdot)$ as the physical evolution operator, which maps the current auxiliary state $\mathbf{C}_t^{\mathrm{aux}}$ and meteorological drivers $\mathbf{M}_{t+1}$ to the physically consistent next state.

**Training Objective with Stepwise Physical Consistency**

Backbone networks are currently used for parallel multi-step prediction to improve computational efficiency. However, if the states of each step are predicted independently, the prediction may violate the internal dynamics of the ecosystem between neighboring steps. Therefore, we introduce a stepwise physical consistency loss that enforces agreement between (i) the backbone's direct prediction of the next auxiliary carbon state and (ii) the one-step transition produced by applying the differentiable DALEC operator $\mathcal{F}_{\mathrm{dalec}}$ to the current predicted state. Concretely, for each forecasting step and each carbon pool $x \in \mathcal{X}$, we compute a DALEC-updated next state $C_x^{\mathrm{dalec}}$ from the current predicted auxiliary state, and penalize its distance from the direct next-state prediction $\hat{C}_x^{\mathrm{aux}}$:

$$\mathcal{L}_{\mathrm{phy}} = \frac{1}{H \cdot |\mathcal{X}|} \sum \left(\frac{\hat{C}_x^{\mathrm{aux}} - C_x^{\mathrm{dalec}}}{\sigma_x}\right)^2, \tag{20}$$

where the summation is taken over all consecutive step pairs within the forecasting horizon of length $H$, $\mathcal{X} = \{f, w, r, \mathrm{lit}, \mathrm{som}\}$ denotes the set of carbon pools, and $\sigma_x$ normalizes scale differences across pools.

Despite mass conservation in the differentiable DALEC–ACM, structural model errors remain. We add a lightweight residual head to predict $\Delta\mathrm{NEP}$ from backbone features and correct the final estimate as

$$\widehat{\mathrm{NEP}} = \mathrm{NEP}_{\mathrm{phy}} + \Delta\mathrm{NEP}, \tag{21}$$

Finally, we optimize a composite objective:

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{nep}} + \lambda_{\mathrm{res}}\mathcal{L}_{\mathrm{res}} + \lambda_{\mathrm{smooth}}\mathcal{L}_{\mathrm{smooth}} + \lambda_{\mathrm{phy}}\mathcal{L}_{\mathrm{phy}}, \tag{22}$$

Here, $\mathcal{L}_{\mathrm{nep}} = \mathbb{E}[(\widehat{\mathrm{NEP}}_t - \mathrm{NEP}_t^{\mathrm{obs}})^2]$ represents the supervised loss against observations. The regularization terms $\mathcal{L}_{\mathrm{res}} = \mathbb{E}[(\Delta\mathrm{NEP}_t)^2]$ and $\mathcal{L}_{\mathrm{smooth}} = \mathbb{E}[(\Delta\mathrm{NEP}_t - \Delta\mathrm{NEP}_{t-1})^2]$ constrain the magnitude and temporal variation of the correction, respectively. The hyperparameters $\lambda$ control the trade-off between data fitting, physical consistency $\mathcal{L}_{\mathrm{phy}}$, and regularization.

# 4 Experiments

## 4.1 Experimental Settings

**Datasets.** Experiments are conducted on the DingHuShan (DHF) national ecological station dataset, which consists of long-term observations from 2005-2025. Input variables include meteorological drivers including air temperature ($T_a$), photosynthetically active radiation (PAR), and soil water content (SWC), together with available remote-sensing indicators. The primary prediction target is NEP, while carbon pool variables ($C_f$, $C_w$, $C_r$, $C_{lit}$, $C_{som}$) are used to enforce physical consistency and support interpretability analysis. All variables are temporally aligned. The dataset is split into training, validation, and test sets with a ratio of 7:2:1.

**Baselines.** We compare the proposed method with both process-based and data-driven baselines. Process-based baselines include the DALEC with default parameters and DALEC with Bayesian parameter inversion using MCMC (Metropolis). Data-driven baselines include representative time-series forecasting models including TimeXer[Wang *et al.*, 2024], iTransformer[Liu *et al.*, 2023] and PatchTST[Nie, 2022].

**Evaluation Metrics.** Prediction performance is evaluated using root mean squared error (RMSE) and mean absolute error (MAE) for NEP forecasts.

**Implementation Details.** We evaluate the proposed NeuroDALEC and data-driven baselines under multiple forecasting settings covering short-term to long-term horizons. Specifically, for the data-driven models, we consider varying input lengths $L$ and prediction length $H$, including $(L, H) \in \{(14, 7), (28, 14), (56, 28), (56, 56), (90, 90)\}$, covering weekly to seasonal time scales. For the DALEC baseline, $L$ corresponds to the full length of the available historical sequence used for model parameter inversion to forecast the subsequent $H$ days. We employ either a TimeXer-based, an iTransformer-based and a PatchTST-based architecture as the time series forecasting backbone. The model is trained using AdamW with a learning rate of $1 \times 10^{-4}$ for neural network parameters. Ecological parameters are optimized jointly using a smaller, fixed learning rate through the proposed process-aware grouped optimization strategy. Loss weights are kept fixed across all experiments, and early stopping is applied based on validation performance. All experiments are conducted on a single workstation equipped with a NVIDIA P100 GPU.

## 4.2 Main Results

Table 2 reports the performance of different methods on the NEP forecasting task. NeuroDALEC achieves the lowest RMSE and MAE errors across all baselines. We observe that data-driven models exhibit higher RMSE values compared to DALEC methods which indicates that general time-series models are insufficient to capture the complex non-linear relationships within the carbon cycle system. NeuroDALEC reduces RMSE by about 20% and MAE by nearly 50% compared with data-driven models and also reduces RMSE by about 12% compared with DALEC (calibration) while maintaining comparable MAE levels. These results demonstrate the differentiable DALEC-ACM physical layer can not only introduce essential physical inductive bias to prevent generating intermediate variables that violate conservation laws but also optimize the physical parameters via backpropagation and enhance the fitting capability of the model. Regarding the effects of input and prediction window lengths, data-driven models exhibit performance degradation when the prediction lengths extends from 14 to 90 (RMSE increase from 3.6 to 4.1). In contrast, NeuroDALEC maintains stability with an RMSE around 2.7 across all prediction length and achieves even better performance when window lengths increase (NeuroDALEC+TimeXer). This trend may suggest that fitting longer window data becomes increasingly difficult for data-driven models lacking physical formula priors. By contrast, NeuroDALEC utilizes the mass conservation relationship between time steps, which enables the model to learn from longer data sequences for precise forecasting.

## 4.3 Ablation Study

We perform ablation experiments to quantify the contribution of each component in NeuroDALEC. As shown in Table 3, **W/O NEP residual correction** results in the largest performance degradation.This may suggest that after DALEC provides predictions within physical constraints, a lightweight neural network for correction can enhance the representation capacity of the model. Strict regularization in the correction can focus on correcting high-frequency residuals without overriding the physical backbone. **W/O physical consistency loss** leads to a relatively small rise in RMSE and MAE. This is because this the primary role of $\mathcal{L}_{\mathrm{phy}}$ is not to directly minimize point-wise prediction errors, but to constrain carbon-pool transitions across consecutive forecasting steps under mass conservation, thereby reducing physical drift, error accumulation, and implausible intermediate states. Hence, this term mainly improves physical plausibility and long-horizon stability, which are essential for reliable multi-step forecasting. **W/O learning ecological parameters** and **W/O grouped optimization** also degrades performance, indicating that Process-Aware Grouped Optimization can coordinates ecological groups with different physical sensitivities to improve accuracy.

## 4.4 Efficiency Analysis

We compare computational efficiency with DALEC models using Bayesian parameter inversion. As shown in Table 4, the differentiable physical layer enables end-to-end gradient-based optimization and yields orders-of-magnitude reductions in training and inference time.

# 5 Interpretability Analysis

To validate the interpretability of NeuroDALEC, we conduct a component-level consistency analysis on the model outputs.

| Settings | | Ours (+TimeXer) | | TimeXer | | Ours (+iTrans) | | iTransformer | | Ours (+PatchTST) | | PatchTST | | DALEC (calibration) | | DALEC (default) | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Pred | Seq | RMSE | MAE | RMSE | MAE | RMSE | MAE | RMSE | MAE | RMSE | MAE | RMSE | MAE | RMSE | MAE | RMSE | MAE |
| 7 | 14 | 2.7713 | 1.9927 | 3.6646 | 2.9329 | 2.7564 | 1.9268 | 3.6796 | 2.9290 | **2.7529** | 1.9299 | 3.6691 | 2.9415 | 3.1457 | **1.9095** | 3.5578 | 1.9829 |
| 14 | 28 | 2.7904 | 1.9231 | 3.7970 | 3.0826 | 2.7783 | 1.9221 | 3.8629 | 3.1448 | **2.7632** | 1.9467 | 3.7577 | 3.0414 | 3.1488 | **1.9124** | 3.3577 | 1.9875 |
| 28 | 56 | **2.7367** | 1.9265 | 3.8163 | 3.1661 | 2.7546 | **1.9154** | 4.1316 | 3.4827 | 2.7696 | 1.9389 | 4.0328 | 3.3725 | 3.1514 | 1.9165 | 3.3580 | 1.9932 |
| 42 | 90 | 2.8754 | 1.9064 | 4.0323 | 3.3094 | 2.7700 | **1.8644** | 4.1832 | 3.5712 | **2.7681** | 1.9528 | 3.9315 | 3.2102 | 3.1516 | 1.9080 | 3.3582 | 1.9951 |
| 56 | 56 | 2.8341 | **1.8661** | 4.0357 | 3.3674 | 2.7824 | 1.8968 | 4.2515 | 3.5662 | **2.7793** | 1.8874 | 4.1258 | 3.4279 | 3.1515 | 1.9077 | 3.3585 | 1.9980 |
| 90 | 90 | **2.7593** | 1.8661 | 4.1196 | 3.3294 | 2.7691 | 1.8538 | 4.2118 | 3.4706 | 2.7705 | **1.8009** | 4.0017 | 3.1010 | 3.1521 | 1.9121 | 3.3590 | 2.0002 |

Table 2: Results on DHF Station. The best and second-best results for each metric are shown in bold and underlined, respectively.

| **Models** | **Metrics** | **Ours (+TimeXer)** |
|---|---|---|
| W/O physical consistency loss | RMSE<br>MAE | 2.7652<br>1.8877 |
| W/O learning ecological parameters | RMSE<br>MAE | 2.9277<br>1.9500 |
| W/O grouped optimization | RMSE<br>MAE | 2.8363<br>1.9288 |
| W/O NEP residual correction | RMSE<br>MAE | 3.0008<br>1.9462 |

Table 3: Ablation study on different modules and components (look-back window length=56, prediction window=28).

| **Method** | **Stage I** (Inversion / Training) | **Stage II** (Prediction) | **Total Time** |
|---|---|---|---|
| DALEC | 5.8 h | 1.2 h | $\sim$ 7 h |
| **Ours** | **5 min** | **30 s** | $\sim$ **0.1h** |

Table 4: Comparison Results on Computational Efficiency.

We verify that the predicted carbon pools and fluxes satisfy mass conservation. The changes in carbon storage are fully explained by the balance between carbon inputs and outputs at each prediction step.

Figure 2 presents a daily ecosystem carbon state predicted by NeuroDALEC. In addition to NEP, the model estimates gross and net primary productivity (GPP, NPP), respiration (Ra, Rh), carbon allocation fluxes to leaves, wood, and roots (Af, Aw, Ar), as well as vegetation and soil carbon pool sizes (Cf, Cw, Cr, Csom). These quantities form a closed and physically consistent carbon balance, where NEP emerges as the net outcome of carbon uptake and respiratory losses rather than an independently regressed variable. The availability of carbon stocks and fluxes further enables the derivation of ecosystem-level indicators. For instance, vegetation and soil carbon turnover times ($\tau_{veg}$ and $\tau_{soil}$) can be computed from the ratios between carbon pools and outgoing fluxes, providing insights into global carbon residence time, ecosystem stability, and long-term carbon storage capacity.

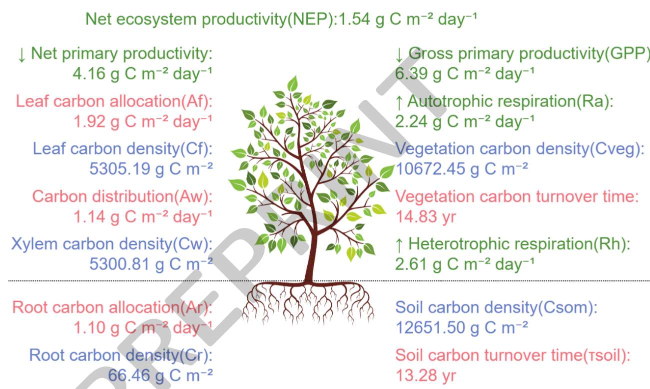


Figure 2: Interpretable carbon-cycle state results, showing the flux chain GPP $\rightarrow$ NPP $\rightarrow$ NEP, mass-conserving pool updates, and ecosystem-level turnover times.

## 6 Social Impact Analysis

The proposed NeuroDALEC model has been adopted in an operational, nation-wide forecasting workflow for forest carbon dynamics across China. The daily forecasting is compatible with Hygon GPU resources. As shown in Figure 3, the system produces routine, daily estimates of carbon fluxes (e.g., NEP, NPP, and GPP) and carbon pool states for national ecological research sites, in a manner analogous to operational weather forecasting. This system provides automated and low-latency monitoring of ecosystem carbon dynamics, which reduces the need for frequent field campaigns and manual sampling by scientists and monitoring staff. The system supports practical tasks such as carbon accounting, ecosystem assessment, and the identification of anomalous carbon responses under climate extremes, contributing to climate governance and sustainable decision-making management. Besides, like weather forecasts, the system's daily terrestrial carbon-cycle predictions can be visualized to intuitively illustrate global carbon turnover for public outreach and K–12 education.

This work is carried out in close collaboration with ecosystem scientists and operational monitoring teams at national ecological research stations, who contributed domain knowledge, data validation, and feedback on model interpretability and usability. Our aim is to help domain experts to diagnose whether variations in NEP are driven by photosynthetic uptake, carbon allocation strategies, or respiratory losses.


Figure 3: Operational interface of the deployed NeuroDALEC-based ecosystem carbon forecasting system.

## 7 Conclusion

We proposed NeuroDALEC, coupling a differentiable DALEC–ACM operator with sequence modeling for long-horizon carbon forecasting under mass-conserving, physically consistent rollouts. Experiments at the DHF station demonstrate improved accuracy and stability with interpretable ecological parameters, and the model has been deployed across multiple Chinese national ecological stations in an operational forecasting system. Future work will extend the model globally to diverse biomes and climatic regions to validate cross-regional generalization, and will explore parameter transfer under ecosystem shifts and disturbances. We also plan to integrate additional observations and quantify predictive uncertainty for carbon monitoring and visualization.

## Ethical Statement

There are no ethical issues.

## Acknowledgements

This research is supported by the National Key Research and Development Program of China (Grant No.2024YFB4506000). We also thank the VenusAI platform (http://data.aicnic.cn) and technical support teams for providing NVIDIA and Hygon GPU resources for model evaluation.

## Contribution Statement

† Equal contribution, ∗ Corresponding author.

## References

[Bernardini *et al.*, 2024] Luca G Bernardini, Christoph Rosinger, Gernot Bodner, Katharina M Keiblinger, Emma Izquierdo-Verdiguier, Heide Spiegel, Carl O Retzlaff, and Andreas Holzinger. Learning vs. understanding: When does artificial intelligence outperform process-based modeling in soil organic carbon prediction? *New Biotechnology*, 81:20–31, 2024.

[Bloom and Williams, 2015] AA Bloom and M Williams. Constraining ecosystem carbon dynamics in a data-limited world: integrating ecological" common sense" in a model–data fusion framework. *Biogeosciences*, 12(5):1299–1315, 2015.

[Cao and Woodward, 1998] Mingkui Cao and F Ian Woodward. Dynamic responses of terrestrial ecosystem carbon cycling to global climate change. *Nature*, 393(6682):249–252, 1998.

[Chen *et al.*, 2024] Zeqiang Chen, Lei Wu, Nengcheng Chen, and Ke Wan. Modeling terrestrial net ecosystem exchange based on deep learning in china. *Remote Sensing*, 17(1):92, 2024.

[Chuter *et al.*, 2015] Anna M Chuter, Philip J Aston, Anne C Skeldon, and Ian Roulstone. A dynamical systems analysis of the data assimilation linked ecosystem carbon (dalec) models. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(3), 2015.

[Delahaies *et al.*, 2017] Sylvain Delahaies, Ian Roulstone, and Nancy Nichols. Constraining dalecv2 using multiple data streams and ecological constraints: analysis and application. *Geoscientific Model Development*, 10(7):2635–2650, 2017.

[Famiglietti *et al.*, 2021] Caroline A Famiglietti, T Luke Smallman, Paul A Levine, Sophie Flack-Prain, Gregory R Quetin, Victoria Meyer, Nicholas C Parazoo, Stephanie G Stettz, Yan Yang, Damien Bonal, et al. Optimal model complexity for terrestrial carbon cycle prediction. *Biogeosciences*, 18(8):2727–2754, 2021.

[Fan *et al.*, 2025] Bin Fan, Hankui K Zhang, Zhongbin B Li, Jingfeng Xiao, Xianghong Che, Zhihua Liu, Gustau Camps-Valls, and Jing M Chen. Estimating carbon fluxes over north america using a physics-constrained deep learning model. *ISPRS Journal of Photogrammetry and Remote Sensing*, 227:551–569, 2025.

[Hari and Tyagi, 2022] Manoj Hari and Bhishma Tyagi. Terrestrial carbon cycle: tipping edge of climate change between the atmosphere and biosphere ecosystems. *Environmental Science: Atmospheres*, 2(5):867–890, 2022.

[Hui *et al.*, 2022] Edwin Hui, Richard Stafford, Iain M Matthews, and V Anne Smith. Bayesian networks as a novel tool to enhance interpretability and predictive power of ecological models. *Ecological Informatics*, 68:101539, 2022.

[Jain and Mitra, 2025] Vishal Jain and Archan Mitra. Toward a balanced earth climate neutrality and the fight against global warming. In *Adaptive Strategies for Green Economy and Sustainability Policies*, pages 391–404. IGI Global Scientific Publishing, 2025.

[Jung *et al.*, 2020] Martin Jung, Christopher Schwalm, Mirco Migliavacca, Sophia Walther, Gustau Camps-Valls, Sujan Koirala, Peter Anthoni, Simon Besnard, Paul Bodesheim, Nuno Carvalhais, et al. Scaling carbon fluxes from eddy covariance sites to globe: synthesis and evaluation of the fluxcom approach. *Biogeosciences*, 17(5):1343–1365, 2020.

[Karniadakis *et al.*, 2021] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

[Liao *et al.*, 2023] Zhangze Liao, Binghuang Zhou, Jingyu Zhu, Hongyu Jia, and Xuehai Fei. A critical review of methods, principles and progress for estimating the gross primary productivity of terrestrial ecosystems. *Frontiers in Environmental Science*, 11:1093095, 2023.

[Liu *et al.*, 2023] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. *arXiv preprint arXiv:2310.06625*, 2023.

[Liu *et al.*, 2024] Licheng Liu, Wang Zhou, Kaiyu Guan, Bin Peng, Shaoming Xu, Jinyun Tang, Qing Zhu, Jessica Till, Xiaowei Jia, Chongya Jiang, et al. Knowledge-guided machine learning can improve carbon cycle quantification in agroecosystems. *Nature communications*, 15(1):357, 2024.

[Ma *et al.*, 2025] YQ Ma, JH Li, W Cao, and L Huang. Grain for green program to grassland might lead to carbon sink leakage in the loess plateau. *Earth's Future*, 13(4):e2024EF005261, 2025.

[McGuire *et al.*, 2001] AD McGuire, S Sitch, Joy S Clein, R Dargaville, G Esser, J Foley, Martin Heimann, F Joos, J Kaplan, DW Kicklighter, et al. Carbon balance of the terrestrial biosphere in the twentieth century: Analyses of co2, climate and land use effects with four process-based ecosystem models. *Global biogeochemical cycles*, 15(1):183–206, 2001.

[Nie, 2022] Y Nie. A time series is worth 64words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*, 2022.

[Perry *et al.*, 2022] George LW Perry, Rupert Seidl, André M Bellvé, and Werner Rammer. An outlook for deep learning in ecosystem science. *Ecosystems*, 25(8):1700–1718, 2022.

[Pierson *et al.*, 2022] Derek Pierson, Kathleen A Lohse, William R Wieder, Nicholas R Patton, Jeremy Facer, Marie-Anne de Graaff, Katerina Georgiou, Mark S Seyfried, Gerald Flerchinger, and Ryan Will. Optimizing process-based models to predict current and future soil organic carbon stocks at high-resolution. *Scientific Reports*, 12(1):10824, 2022.

[Raissi *et al.*, 2019] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

[Shen *et al.*, 2023] Chaopeng Shen, Alison P Appling, Pierre Gentine, Toshiyuki Bandai, Hoshin Gupta, Alexandre Tartakovsky, Marco Baity-Jesi, Fabrizio Fenicia, Daniel Kifer, Li Li, et al. Differentiable modelling to unify machine learning and physical models for geosciences. *Nature Reviews Earth & Environment*, 4(8):552–567, 2023.

[Smallman and Williams, 2019] Thomas Luke Smallman and Mathew Williams. Description and validation of an intermediate complexity model for ecosystem photosynthesis and evapotranspiration: Acm-gpp-etv1. *Geoscientific Model Development*, 12(6):2227–2253, 2019.

[Song *et al.*, 2025] Jian Song, Shiqiang Wan, Shilong Piao, Jianyang Xia, Yan Ning, Mengmei Zheng, Dafeng Hui, Jingyi Ru, Juanjuan Han, Jiayin Feng, et al. Global change and china's terrestrial carbon sink: A quantitative review of 30 years' ecosystem manipulative experiments. *Ecological Monographs*, 95(1):e70005, 2025.

[Tian *et al.*, 2025] Yaxin Tian, Xiang Ren, Keke Li, and Xiangqian Li. Carbon dioxide emission forecast: A review of existing models and future challenges. *Sustainability*, 17(4):1471, 2025.

[Wang *et al.*, 2021] Fang Wang, Jean Damascene Harindintwali, Zhizhang Yuan, Min Wang, Faming Wang, Sheng Li, Zhigang Yin, Lei Huang, Yuhao Fu, Lei Li, et al. Technologies and perspectives for achieving carbon neutrality. *The innovation*, 2(4), 2021.

[Wang *et al.*, 2024] Yuxuan Wang, Haixu Wu, Jiaxiang Dong, Guo Qin, Haoran Zhang, Yong Liu, Yunzhong Qiu, Jianmin Wang, and Mingsheng Long. Timexer: Empowering transformers for time series forecasting with exogenous variables. *Advances in Neural Information Processing Systems*, 37:469–498, 2024.

[Willard *et al.*, 2022] Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating scientific knowledge with machine learning for engineering and environmental systems. *ACM Computing Surveys*, 55(4):1–37, 2022.

[Williams *et al.*, 2005] Mathew Williams, Paul A Schwarz, Beverly E Law, James Irvine, and Meredith R Kurpius. An improved analysis of forest carbon dynamics using data assimilation. *Global change biology*, 11(1):89–105, 2005.

[Yang *et al.*, 2022] Yan Yang, A Anthony Bloom, Shuang Ma, Paul Levine, Alexander Norton, Nicholas C Parazoo, John T Reager, John Worden, Gregory R Quetin, T Luke Smallman, et al. Cardamom-fluxval version 1.0: a fluxnet-based validation system for cardamom carbon and water flux estimates. *Geoscientific Model Development*, 15(4):1789–1802, 2022.

[Yuan *et al.*, 2024] Wenping Yuan, Jiangzhou Xia, Chaoqing Song, and Ying-Ping Wang. Simulating the land carbon sink: Progresses and challenges of terrestrial ecosystem models. *Agricultural and Forest Meteorology*, 358:110264, 2024.