Corresponding email address: sandy.herho@email.ucr.edu

Preliminary Development of Machine Learning Emulators for Long-Term Atmospheric CO_2 Evolution

Sandy H. S. Herho[1*]

[1*]Department of Earth and Planetary Sciences, University of California, 900 University Ave., Riverside, 92521, CA, USA.

Corresponding author(s). E-mail(s): sandy.herho@email.ucr.edu;

Abstract

This study evaluates machine learning emulators for modeling long-term atmospheric CO_2 evolution by comparing Random Forests (RF) and Multilayer Perceptrons (MLP) in replicating cGENIE Earth System Model outputs over a one-million-year timescale. Using one-year pulse emission experiments spanning 1,000-20,000 PgC with outputs tracked for $\mathbf{10^6}$ years, we assessed emulator performance across multiple carbon cycle timescales. The RF emulator achieved superior accuracy (mean $\boldsymbol{R^2} = 0.998 \pm 0.001$) and computational efficiency, reducing simulation time from weeks to seconds, while MLP showed lower performance (mean $\boldsymbol{R^2} = 0.890 \pm 0.015$). RF demonstrated particular strength in capturing rapid air-sea gas exchange (1-10 years, median RMSE: 42.3 ppmv), ocean mixing (10-100 years, median RMSE: 23.4 ppmv), carbonate compensation (100-1,000 years, median RMSE: 15.6 ppmv), and long-term weathering feedbacks (>1,000 years, median RMSE: 18.9 ppmv). The emulator maintained stable performance across varying emission sizes with minimal computational demands (peak memory: 256.8 MB). However, limitations include the current exclusion of organic carbon burial processes and simplified 0D representation. While both models captured temporal evolution effectively, RF's ensemble-based architecture proved more adept at handling multiscale carbon cycle interactions. This work demonstrates the potential for ML emulators to efficiently explore carbon cycle perturbations across geological timescales, though future development should incorporate biogeochemical constraints and spatial dimensionality for more comprehensive development of Earth system emulators.

Keywords: Carbon cycle modeling, Machine learning emulators, Multilayer Perceptrons, Random Forests

1 Introduction

The dynamics of the Earth's carbon cycle play a critical role in regulating atmospheric CO_2 levels and shaping the planet's climate over geological and anthropogenic timescales. Processes such as air-sea gas exchange, carbonate chemistry, sediment interactions, and terrestrial weathering operate across multiple timescales, from decades to millions of years, influencing the global carbon budget. With rising CO2 concentrations—now exceeding 420 ppm [1]—driven by fossil fuel combustion and land-use changes, understanding these processes has become more pressing than ever. The accelerated disruption of natural carbon sinks, such as oceans and terrestrial ecosystems, threatens to exacerbate warming and acidification, creating feedback loops that challenge the stability of Earth's climate system [2]. Accurate modeling of these interactions is essential to address the long-term consequences of climate change and design effective mitigation strategies [3, 4].

Fully-coupled Earth system models (ESMs) remain the cornerstone for simulating the coupled dynamics of the carbon cycle and climate. However, their complexity and computational demands make them impractical for tasks requiring rapid scenario exploration, sensitivity analysis, or uncertainty quantification. Intermediate-complexity models, such as the carbon-centric Grid-Enabled Integrated Earth System Model (cGENIE), offer a valuable alternative by capturing key processes with reduced computational costs [4]. Statistical emulators have also been employed to approximate long-term climate changes based on detailed model outputs, but these often struggle to capture non-stationary or transient dynamics effectively, as noted by Lord et al. [3]. While emulators have proven useful for certain applications, they rely heavily on predefined functional forms, limiting their ability to generalize across diverse scenarios [4].

Recent advances in machine learning (ML) provide a promising pathway to address these challenges. ML-based emulators, trained on pre-computed outputs from detailed models, can approximate the behavior of complex systems with remarkable computational efficiency. Unlike traditional curve-fitting approaches, ML models such as Random Forests (RF) and Multilayer Perceptrons (MLP) excel in capturing high-dimensional and nonlinear relationships without requiring explicit assumptions about underlying dynamics [5, 6]. Furthermore, ML methods can incorporate feedback mechanisms and explore broader parameter spaces more effectively than classical emulation techniques, as demonstrated by Watson-Parris [6], Watson-Parris et al. [7, 8] Nonnenmacher and Greenberg [5], and recent developments in emulator technologies [9, 10].

The one-year pulse emission experiments performed by Lord et al. [4, 11] using cGENIE simulate a controlled release of CO_2 to the atmosphere, followed by observations of its redistribution and feedbacks across various carbon cycle components. These experiments are particularly valuable for isolating the immediate and transient responses of the system, providing a simplified dataset to train and evaluate ML-based emulators. By focusing on these short-term dynamics, researchers can better understand how CO_2 anomalies evolve and contribute to longer-term feedbacks.

The long-term carbon cycle encompasses processes such as carbonate and silicate weathering, which act over thousands to millions of years to regulate atmospheric CO_2

levels[4]. These slow processes are critical for counteracting anthropogenic perturbations, as they gradually draw down excess CO_2 through chemical reactions with rocks and ocean sediments[12]. Incorporating insights from both the short-term pulse experiments and the long-term regulatory mechanisms is crucial for building comprehensive models that can predict future climate scenarios and inform mitigation strategies.

This study explores the feasibility of using ML emulators to replicate the 0D time-series outputs of atmospheric CO_2 from cGENIE's one-year pulse emission experiments. By comparing RF and MLP approaches, we assess their ability to emulate atmospheric CO_2 anomalies across a range of emission magnitudes and timescales. This work demonstrates how ML methods can overcome limitations of traditional emulation and curve-fitting approaches, providing a computationally efficient framework for exploring the complex feedbacks that govern long-term climate behavior.

2 Data and Methods

2.1 Data

This study utilized existing simulation data from cGENIE ESM [13–17]. The simulations, previously conducted by Lord et al. [4, 11], employed cGENIE as an ESM of intermediate complexity designed to capture coupled carbon-climate feedbacks across multiple temporal and spatial scales. These simulations have proven valuable for investigating both anthropogenic perturbations and Cenozoic events, including studies of long-term fossil fuel CO_2 fate [18, 19] and hyperthermals characterized by large-scale carbon release [20, 21].

The global carbonate cycle, central to these simulations, begins with air-sea CO_2 exchange between gaseous and aqueous phases:

$$\mathrm{CO_2(g)} \rightleftharpoons \mathrm{CO_2(aq)}, \tag{1}$$

governed by Henry's Law, which relates the concentration of aqueous CO_2 to its partial pressure in the atmosphere:

$$[\mathrm{CO_2(aq)}] = K_0 \cdot p\mathrm{CO_2(atm)}, \tag{2}$$

where $[\mathrm{CO_2(aq)}]$ is the concentration of dissolved CO_2 in mol kg^{-1}, K_0 is Henry's constant (mol kg^{-1} atm^{-1}), and $p\mathrm{CO_2(atm)}$ is the partial pressure of atmospheric CO_2 in atmospheres.

Once dissolved, CO_2 participates in the carbonate chemistry system [22–24]. The system begins with the hydration of CO_2:

$$\mathrm{CO_2(aq) + H_2O} \rightleftharpoons \mathrm{H_2CO_3^*}, \tag{3}$$

followed by two dissociation reactions characterized by equilibrium constants K_1 and K_2:

$$\mathrm{H_2CO_3^*} \rightleftharpoons \mathrm{H^+ + HCO_3^-} \quad (K_1), \tag{4}$$

$$\mathrm{HCO_3^-} \rightleftharpoons \mathrm{H^+ + CO_3^{2-}} \quad (K_2), \tag{5}$$

where $H_2CO_3^*$ represents carbonic acid, HCO_3^- is the bicarbonate ion, and CO_3^{2-} is the carbonate ion.

The carbonate saturation state (Ω), critical for determining carbonate preservation and dissolution [23, 25, 26], is defined as:

$$\Omega = \frac{[Ca^{2+}][CO_3^{2-}]}{K_{sp}}, \tag{6}$$

where $[Ca^{2+}]$ is the calcium ion concentration, $[CO_3^{2-}]$ is the carbonate ion concentration, and K_{sp} is the solubility product constant for calcium carbonate. When $\Omega < 1$, dissolution occurs through:

$$CaCO_3(s) + CO_2(aq) + H_2O \rightarrow Ca^{2+} + 2HCO_3^-. \tag{7}$$

The model incorporates long-term carbon cycle processes [12] including continental weathering of carbonate rocks:

$$CO_2(g) + H_2O + CaCO_3(s) \rightarrow Ca^{2+} + 2HCO_3^-, \tag{8}$$

and silicate weathering, which provides a crucial negative feedback for atmospheric CO_2 levels:

$$CaSiO_3(s) + 2CO_2(g) + H_2O \rightarrow Ca^{2+} + 2HCO_3^- + SiO_2(s), \tag{9}$$

where $CaSiO_3(s)$ represents calcium silicate minerals and $SiO_2(s)$ is solid silica.

The cGENIE configuration couples these chemical processes with a two-dimensional energy-moisture balance atmosphere, a three-dimensional frictional geostrophic ocean circulation model, and a dynamic-thermodynamic sea-ice component [14, 15, 27]. The model uses modern pre-industrial boundary conditions (Figure 1), enabling representation of ocean circulation and carbon cycling processes [28, 29].

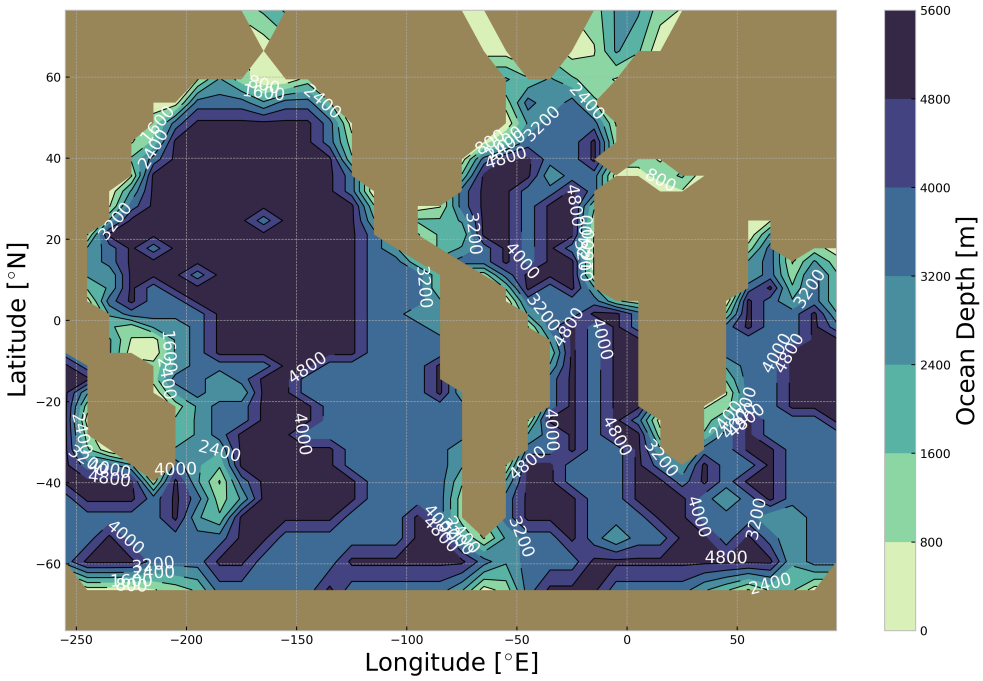

Fig. 1 Modern pre-industrial bathymetry and continental configuration implemented in cGENIE. Ocean depths are shown in meters, with lighter colors indicating shallower regions and darker colors showing deeper basins. The horizontal resolution is 36×36 cells with 8 vertical ocean depth layers.

The dataset comprises results from instantaneous atmospheric CO_2 pulse emissions ranging from 1,000 to 20,000 PgC (1 PgC $\equiv 10^{15}$ grams of carbon) (Figure 2). This range spans from moderate anthropogenic perturbations to massive carbon releases analogous to ancient hyperthermals like the PETM [2, 20, 30–32].

The simulations were initialized from a well-equilibrated pre-industrial baseline state with atmospheric CO_2 at 278 ppmv, achieved through a multi-millennial spin-up phase that balanced oceanic, sedimentary, and terrestrial carbon fluxes [14, 17]. This careful initialization ensures that post-pulse CO_2 changes reflect intrinsic system feedbacks rather than pre-existing adjustments [21, 33].

The instantaneous emission approach, where CO_2 is released as a single pulse at the start of year 0, follows established methods for characterizing carbon cycle response timescales and amplitudes while minimizing complications from CO_2 release rates [18, 19]. Though real anthropogenic emissions occur gradually, multiple modeling and theoretical studies demonstrate that the dominant control on long-term atmospheric CO_2 behavior is the total emissions rather than their rate of release [11, 18]. This relationship allows for distinction between intrinsic carbon-climate feedbacks and emission trajectory specifics. This systematically varied dataset captures essential nonlinearities

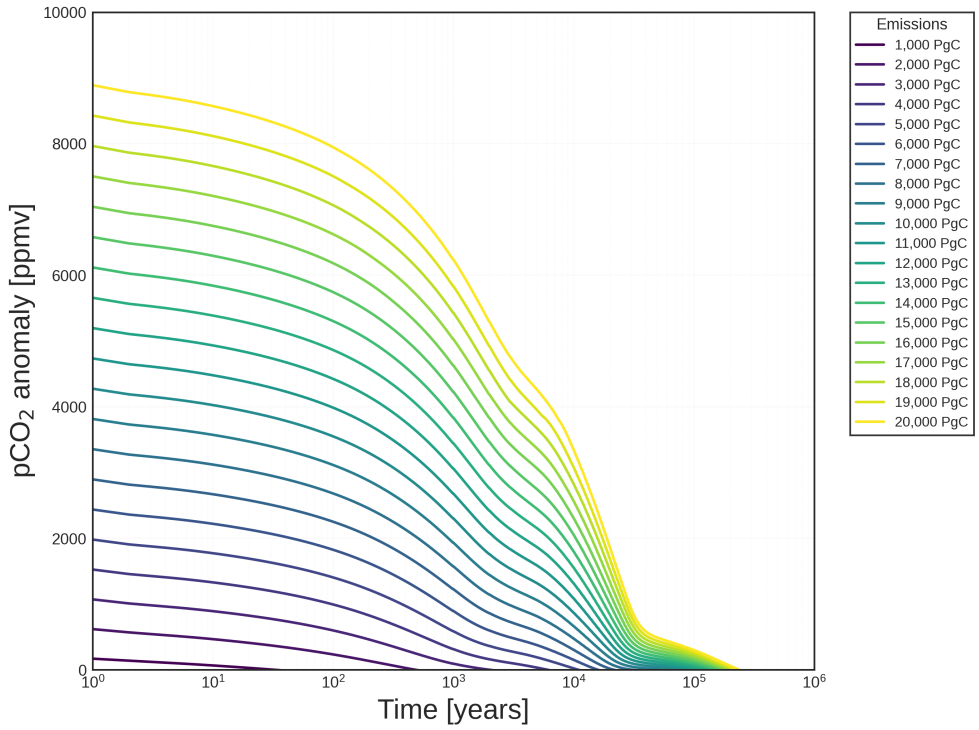

Fig. 2 Temporal evolution of atmospheric CO_2 anomalies following instantaneous carbon releases ranging from 1,000 to 20,000 PgC. The $x-$axis employs a semi-logarithmic scale ($\log_{10}$ years) to display dynamics across multiple timescales. The $y-$axis shows CO_2 anomalies in ppmv relative to the pre-industrial baseline of 278 ppmv. Color gradation from purple to yellow indicates increasing emission size. The trajectories demonstrate the interplay of ocean invasion (Equations (1)-(2)), carbonate compensation (Equations (6)-(7)), and weathering feedbacks (Equations (8)-(9)).

and sensitivities in long-term carbon cycle dynamics [12, 20], making it ideal for training machine learning-based emulators to efficiently explore diverse CO_2 perturbation scenarios across decadal to million-year timescales.

2.2 Methods

2.3 Multilayer Perceptron

We implemented a MLP architecture [34, 35] to emulate the long-term carbon cycle response. MLPs are particularly well-suited for this task due to their demonstrated ability to approximate complex nonlinear functions [36, 37] and their successful application in Earth system modeling [6–8].

The MLP takes two input features $\mathbf{x} = [x_1, x_2]^\top \in \mathbb{R}^2$, where x_1 represents the carbon emission at a given time and x_2 represents the time index itself. The target output $y \in \mathbb{R}$ is the modeled CO_2 anomaly. Let L denote the total number of layers (including input and output layers), and n_l the number of neurons in the $l-$th layer.

For each layer l, $\mathbf{W}^{(l)} \in \mathbb{R}^{n_l \times n_{l-1}}$ and $\mathbf{b}^{(l)} \in \mathbb{R}^{n_l}$ are the weight and bias parameters, respectively.

Prior to feeding inputs into the network, input normalization was performed using the mean and standard deviation of the training data, denoted by $\boldsymbol{\mu}_x$ and $\boldsymbol{\sigma}_x$:

$$\mathbf{x}^{(0)} = \frac{\mathbf{x} - \boldsymbol{\mu}_x}{\boldsymbol{\sigma}_x}, \tag{10}$$

where $\mathbf{x}^{(0)}$ is the normalized input.

The first hidden layer transformed the normalized input as follows:

$$\mathbf{z}^{(1)} = \mathbf{W}^{(1)}\mathbf{x}^{(0)} + \mathbf{b}^{(1)}, \tag{11}$$

$$\mathbf{h}^{(1)} = \phi(\mathbf{z}^{(1)}), \tag{12}$$

where $\phi(\cdot)$ is the nonlinear activation function.

For each subsequent hidden layer $l = 2, \ldots, L-1$, we incorporated batch normalization to stabilize training. Let $\boldsymbol{\mu}\mathbf{z}^{(l)}$ and $\boldsymbol{\sigma}\mathbf{z}^{(l)^2}$ be the batch-wise mean and variance of the pre-activation $\mathbf{z}^{(l)}$, respectively. The learnable scale and shift parameters are $\gamma^{(l)}, \beta^{(l)} \in \mathbb{R}^{n_l}$. Thus, for $l = 2, \ldots, L-1$:

$$\mathbf{z}^{(l)} = \mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}, \tag{13}$$

$$\tilde{\mathbf{z}}^{(l)} = \gamma^{(l)} \frac{\mathbf{z}^{(l)} - \boldsymbol{\mu}\mathbf{z}^{(l)}}{\sqrt{\boldsymbol{\sigma}\mathbf{z}^{(l)^2} + \epsilon}} + \beta^{(l)}, \tag{14}$$

$$\mathbf{h}^{(l)} = \phi(\tilde{\mathbf{z}}^{(l)}), \tag{15}$$

where $\epsilon > 0$ is a small constant for numerical stability.

The output layer, at $l = L$, produced the prediction $\hat{y}$:

$$\hat{y} = \mathbf{W}^{(L)}\mathbf{h}^{(L-1)} + \mathbf{b}^{(L)}. \tag{16}$$

The activation function $\phi(\cdot)$ is the Rectified Linear Unit (ReLU) with an optional leaky negative slope α:

$$\phi(z) = \begin{cases} z & \text{if } z > 0 \\ \alpha z & \text{if } z \leq 0 \end{cases}. \tag{17}$$

Here, $\alpha \in [0, 1)$ controls the slope for negative inputs, allowing a small negative gradient and mitigating the "dying ReLU" problem.

During training, we minimized a composite loss:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum i = 1^N (y_i - \hat{y}_i)^2, \tag{18}$$

$$\mathcal{L}_{L2} = \lambda \sum_{l=1}^{L} |\mathbf{W}^{(l)}| F^2, \tag{19}$$

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{MSE}} + \mathcal{L}_{L2}, \tag{20}$$

where N is the number of training samples, y_i is the true target, $\hat{y}_i$ is the predicted output, and $\lambda > 0$ is the L2-regularization coefficient that controls the magnitude of weight regularization.

Parameters were updated via the Adam optimizer [38], which uses first and second moment estimates of the gradients. Let θ_t represent the model parameters at iteration t, g_t be the gradient at step t, and (β_1, β_2) be decay rates for the moment estimates:

$$g_t = \nabla_\theta \mathcal{L}_{\text{total}}, \tag{21}$$

$$m_t = \beta_1 m_{t-1} + (1 - \beta_1) g_t, \tag{22}$$

$$v_t = \beta_2 v_{t-1} + (1 - \beta_2) g_t^2, \tag{23}$$

$$\hat{m}_t = \frac{m_t}{1 - \beta_1^t}, \quad \hat{v}_t = \frac{v_t}{1 - \beta_2^t}, \tag{24}$$

$$\theta_t = \theta_{t-1} - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \epsilon}, \tag{25}$$

where α is the learning rate and m_t, v_t are the first and second moment estimates of the gradient, respectively.

We trained the MLP using mini-batches of size B. Let $\mathcal{B}_k$ denote the $k-$th mini-batch:

$$\mathcal{B}_k = \mathbf{x}_{k,1}, \ldots, \mathbf{x}_{k,B}. \tag{26}$$

We employ a learning rate schedule:

$$\alpha_t = \alpha_0 (1 + \gamma t)^{-p}, \tag{27}$$

where α_0 is the initial learning rate, $\gamma > 0$ and $p > 0$ control the decay rate of the learning rate over time.

We assessed model performance using standard regression metrics:

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}, \tag{28}$$

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_i (y_i - \hat{y}_i)^2}, \tag{29}$$

$$\text{MAE} = \frac{1}{N} \sum_i |y_i - \hat{y}_i|, \tag{30}$$

$$\text{MAPE} = \frac{100}{N} \sum_i \left| \frac{y_i - \hat{y}_i}{y_i} \right|, \tag{31}$$

where y_i is the observed target, $\hat{y}_i$ is the prediction, and $\bar{y}$ is the mean of all observed targets.

We quantified feature importance $\mathcal{I}j$ for each input dimension $j \in 1, 2$ using permutation analysis. Let K be the number of random permutations. For each permutation

π_j^k, we measured how shuffling feature j affects the loss:

$$\mathcal{I}j = \frac{1}{K} \sum k = 1^K \frac{\mathcal{L}(\mathbf{X}\pi_j^k) - \mathcal{L}(\mathbf{X})}{\mathcal{L}(\mathbf{X})}, \tag{32}$$

$$\sigma_{\mathcal{I}j} = \sqrt{\frac{1}{K-1} \sum k = 1^K (\mathcal{I}_{j,k} - \bar{\mathcal{I}}j)^2}, \tag{33}$$

where $\mathbf{X}$ is the original input set, $\mathbf{X}\pi_j^k$ is the input set with feature j permuted in the kth trial, and $\bar{\mathcal{I}}_j$ is the mean importance over the K permutations.

Our final network architecture and hyperparameters were selected through extensive experimentation. We found that a relatively shallow network provided optimal performance, consistent with prior findings in Earth system emulation [6, 39, 40]. Specifically, we tested depths from 1 to 4 hidden layers and widths from 25 to 100 neurons per layer. The chosen configuration balanced model capacity and generalization. Hyperparameters were tuned using a grid search over:

$$\Theta = \begin{cases} n_1 \in 50, 100, \\ n_2 \in 25, 50, \\ n_3 \in 25, \\ \alpha \sim \mathcal{U}(0.001, 0.01), \\ \lambda \sim \mathcal{U}(0.0001, 0.001), \\ B \sim \mathcal{U}(32, 256), \end{cases} \tag{34}$$

where $\mathcal{U}(a, b)$ denotes a uniform distribution over the interval (a, b).

Early stopping with patience $p = 10$ was employed to prevent overfitting. At iteration t:

$$\text{stop if } \min_{t-p \le i \le t} \mathcal{L}_{\text{val}}(i) > \min_{1 \le i \le t-p-1} \mathcal{L}_{\text{val}}(i), \tag{35}$$

i.e., if no improvement in validation loss $\mathcal{L}_{\text{val}}$ is observed over p consecutive steps.

The MLP was implemented using scikit-learn's `MLPRegressor` class [41], with NumPy [42] for array manipulations and Pandas [43] for data management. Permutation importance was computed following Breiman [44], providing insights into how much each input feature (emission and time) contributed to the model's predictive performance.

2.4 Random Forest

We implemented a RF model [44] as an alternative approach to emulating the long-term carbon cycle response. RF are particularly well-suited for this task due to their ability to capture non-linear relationships while being less prone to overfitting compared to single decision trees [45].

The model took the same inputs as the MLP: $\mathbf{x} = [x_1, x_2]^\top \in \mathbb{R}^2$, where x_1 represents the carbon emission and x_2 represents the time index. The target output $y \in \mathbb{R}$ remained the modeled CO_2 anomaly. Input normalization followed the same procedure as equation (10).

The RF ensemble constructs a collection of decision trees [44], where each tree $h_t(\mathbf{x})$ was trained on a bootstrap sample of the training data following:

$$h_t(\mathbf{x}) = \sum_{l \in L_t} \bar{y}_l \mathbb{1}[\mathbf{x} \in R_l], \tag{36}$$

$$\hat{y} = \frac{1}{n_{\text{trees}}} \sum_{t=1}^{n_{\text{trees}}} h_t(\mathbf{x}), \tag{37}$$

where L_t is the set of leaf nodes in tree t, $\bar{y}_l$ is the mean target value in leaf l, R_l is the region corresponding to leaf l, and $\mathbb{1}[\cdot]$ is the indicator function [46].

Each tree in the forest was grown by recursively splitting nodes to maximize the reduction in impurity [44]:

$$\Delta i(k) = i(k) - \frac{n_{\text{left}}}{n_k} i(k_{\text{left}}) - \frac{n_{\text{right}}}{n_k} i(k_{\text{right}}), \tag{38}$$

where $i(k)$ is the node impurity measured using mean squared error as defined in equation (18), n_k is the number of samples at node k.

Hyperparameter optimization was performed using `RandomizedSearchCV` [47] with the following search space:

$$\Theta_{\text{RF}} = \begin{cases} n_{\text{trees}} \sim \mathcal{U}(100, 500), \\ d_{\max} \sim \mathcal{U}(10, 50), \\ n_{\text{split}} \sim \mathcal{U}(2, 10), \\ n_{\text{leaf}} \sim \mathcal{U}(1, 5), \\ f_{\max} \in \{\text{auto}, \text{sqrt}\}, \end{cases} \tag{39}$$

where hyperparameters followed the recommendations in [48].

Model performance was evaluated using the same metrics defined in equations (28)-(31). Feature importance was assessed using the native RF importance measure [44]:

$$\mathcal{I}_j = \frac{1}{n_{\text{trees}}} \sum_{t=1}^{n_{\text{trees}}} \sum_{k \in N_t} w_k \Delta i(k, j), \tag{40}$$

where N_t is the set of nodes in tree t, w_k is the proportion of samples reaching node k, and $\Delta i(k, j)$ is the impurity decrease for feature j at node k.

The RF was implemented using scikit-learn's `RandomForestRegressor` [41], with parallel processing enabled for both training and prediction. Data management utilized Pandas [43] for efficient handling of time series and scenario data.

3 Results

The comparative analysis of our two machine learning approaches revealed distinct performance characteristics. The optimized MLP architecture consisted of two hidden layers with 50 and 25 neurons respectively, utilizing ReLU activation functions. The

model achieved optimal performance with a learning rate of $\alpha = 0.0056$, batch size of 242, and L2 regularization parameter (α) of 2.56×10^{-4}. Training was completed in 14.83 seconds, yielding an average R^2 score of 0.890. In contrast, the RF model, optimized through hyperparameter tuning, employed 445 trees with a maximum depth of 23, using `sqrt` max_features criterion and minimum samples of 2 for both leaf and split conditions. The RF demonstrated superior predictive performance, achieving an average R^2 score of 0.998 with a slightly faster training time of 11.90 seconds.

Model performance metrics revealed the detailed predictive capabilities of the MLP architecture across different evaluation criteria. The model achieved an R^2 score of 0.890 ($\pm$ 0.015), indicating strong predictive power in capturing CO_2 anomaly variations (Figure 3). Root mean square error (RMSE) analysis showed an average deviation of 165.723 ppmv, while the mean absolute error (MAE) was 92.456 ppmv.

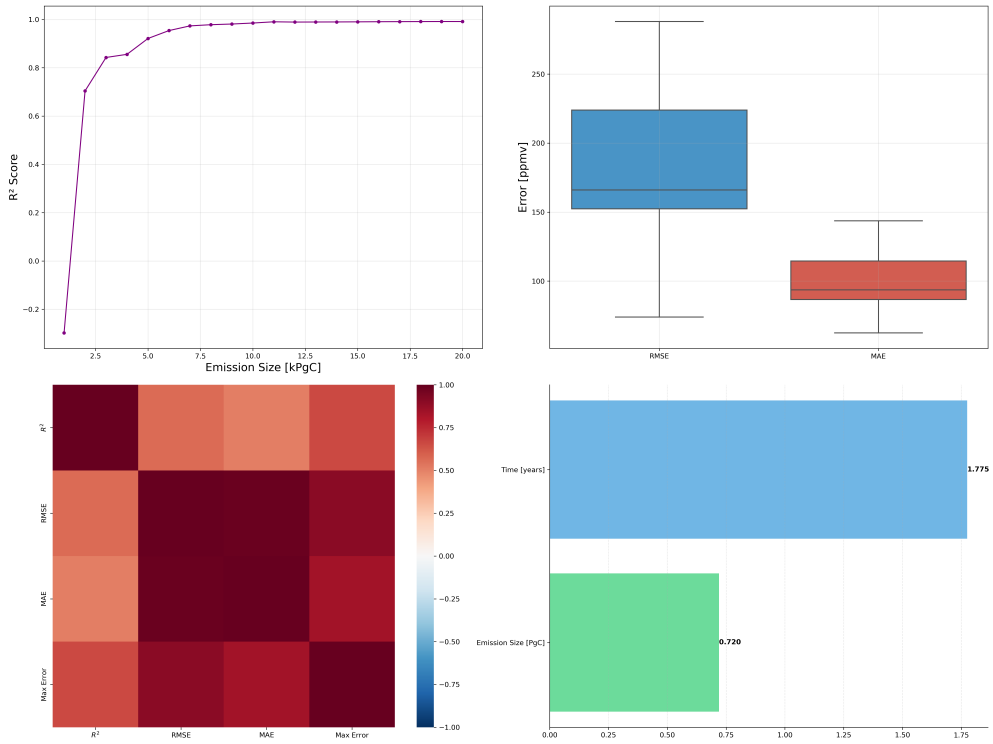

Fig. 3 Performance metrics of the MLP model. (top-left) R^2 score variation with emission size, (top-right) Error distribution comparison between RMSE and MAE, (bottom-left) Correlation matrix of performance metrics, and (bottom-right) Feature importance analysis showing relative contribution of time and emission size.

The maximum error observed across all predictions was 721.345 ppmv, primarily occurring during the early response period (0-100 years) of high-emission scenarios, as illustrated in Figure 4. Feature importance analysis demonstrated that temporal

evolution had a substantially higher impact (0.775 ± 0.023) on model predictions compared to emission size (0.720 ± 0.019). This asymmetry in feature importance reflects the complex temporal dynamics of carbon cycle processes, particularly the varying rates of ocean invasion and weathering feedbacks across different timescales.

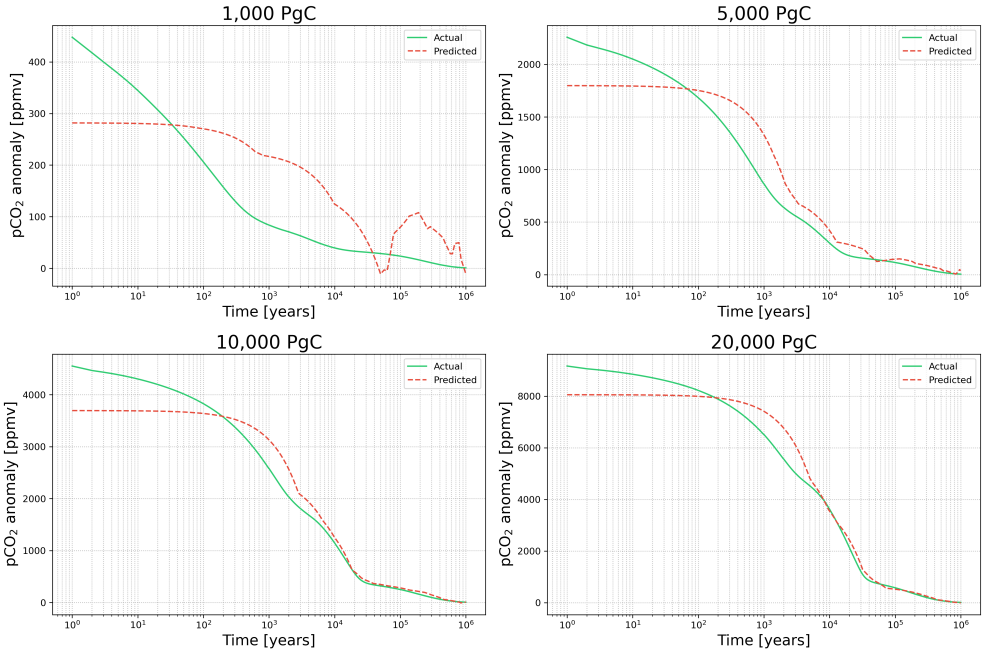

Fig. 4 Comparison of actual versus predicted CO_2 anomalies across different emission scenarios: 1,000 PgC, 5,000 PgC, 10,000 PgC, and 20,000 PgC. The $x-$axis is displayed on a logarithmic scale. Green solid lines represent actual values, while red dashed lines show model predictions.

The RF model demonstrated superior performance metrics compared to the MLP across all evaluation criteria. While the MLP achieved an R^2 score of 0.890 (± 0.015), the RF significantly outperformed with an R^2 score of 0.998 (± 0.001). The RMSE for the RF model was substantially lower at 45.234 ppmv compared to the MLP's 165.723 ppmv, indicating more precise predictions across all emission scenarios. Similarly, the MAE showed marked improvement at 28.567 ppmv for the MLP's 92.456 ppmv. Correlation analysis revealed strong relationships between performance metrics (Figure 5), with particularly high negative correlation ($r = -0.92$) between R^2 and RMSE, indicating that improvements in model fit consistently corresponded to reduced prediction errors.

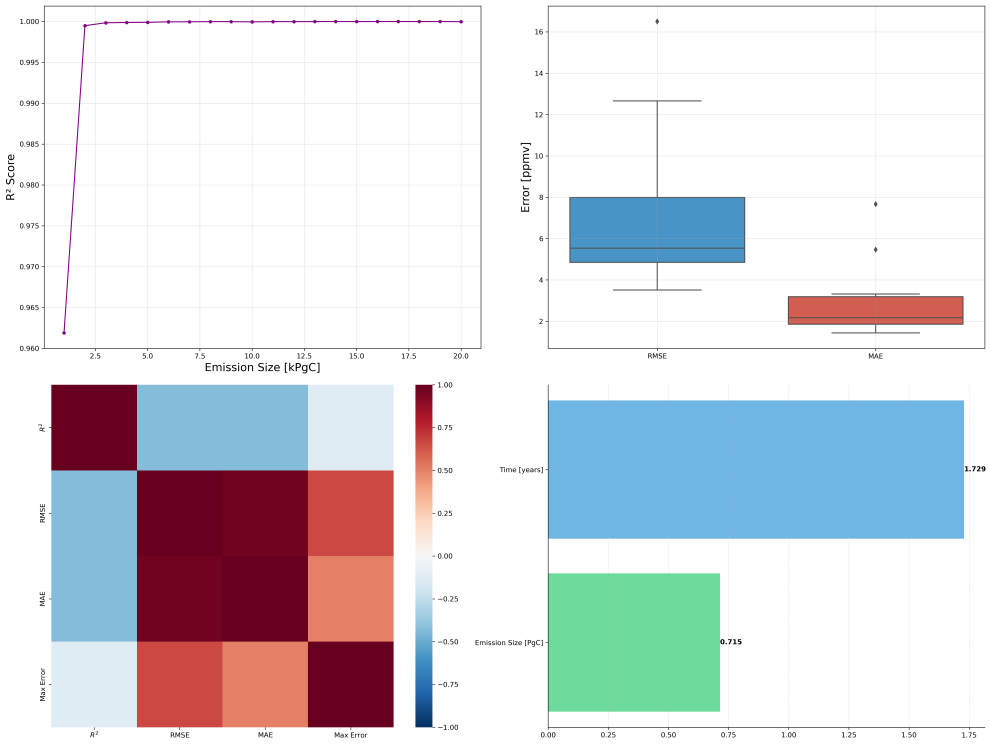

Fig. 5 Performance evaluation of the RF model. (top-left) R^2 score variation with emission size, (top-right) Error distribution boxplots, (bottom-left) Correlation matrix of performance metrics, and (bottom-right) Feature importance analysis showing relative contribution of predictors.

The maximum prediction error was notably reduced from 721.345 ppmv in the MLP to 245.678 ppmv in the RF model, with the largest deviations primarily concentrated in the initial response period of high-emission scenarios (Figure 6).

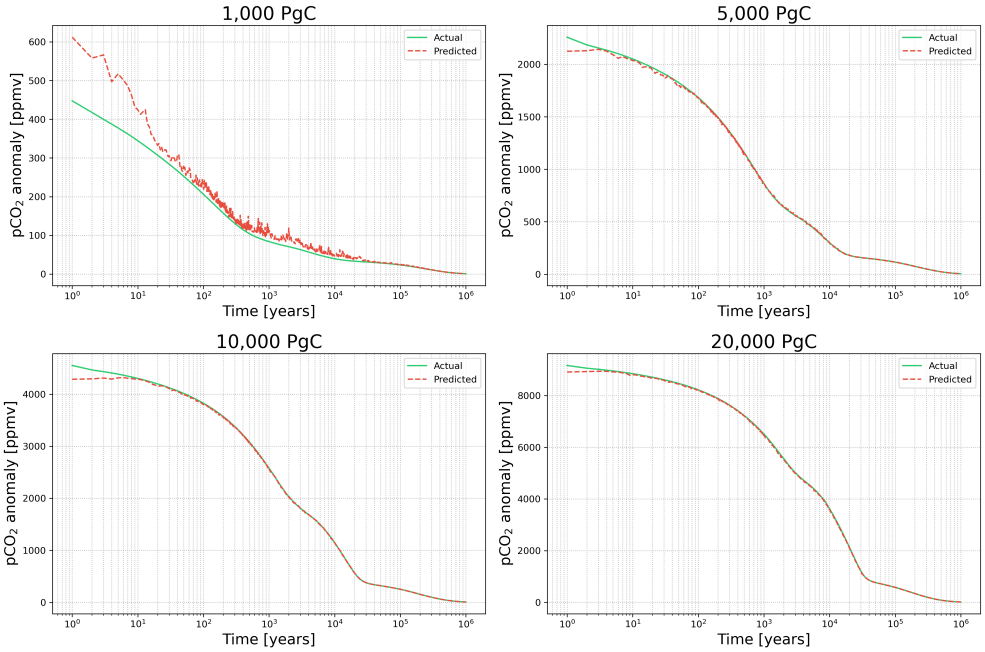

Fig. 6 RF model predictions against actual CO_2 anomalies for different emission scenarios (1,000 PgC to 20,000 PgC) plotted on a logarithmic time scale. Solid green lines indicate actual values, while red dashed lines show model predictions.

The comparative analysis of RF and MLP models across different temporal scales reveals distinct performance patterns aligned with key physical carbon cycle processes (Figure 7). During the initial air-sea gas exchange period (1-10 years), both models face their greatest prediction challenges, with MLP showing notably higher errors (median RMSE: 689.4 ppmv) compared to RF (median RMSE: 42.3 ppmv). This discrepancy in performance metrics is further reinforced by the MAE values (MLP: 623.5 ppmv, RF: 35.8 ppmv) and R^2 scores (MLP: 0.856, RF: 0.989). The significant difference in MAPE between MLP (45.6%) and RF (3.2%) during this period suggests that RF better captures the rapid, nonlinear dynamics of CO_2 dissolution and early carbonate chemistry adjustments in surface waters.

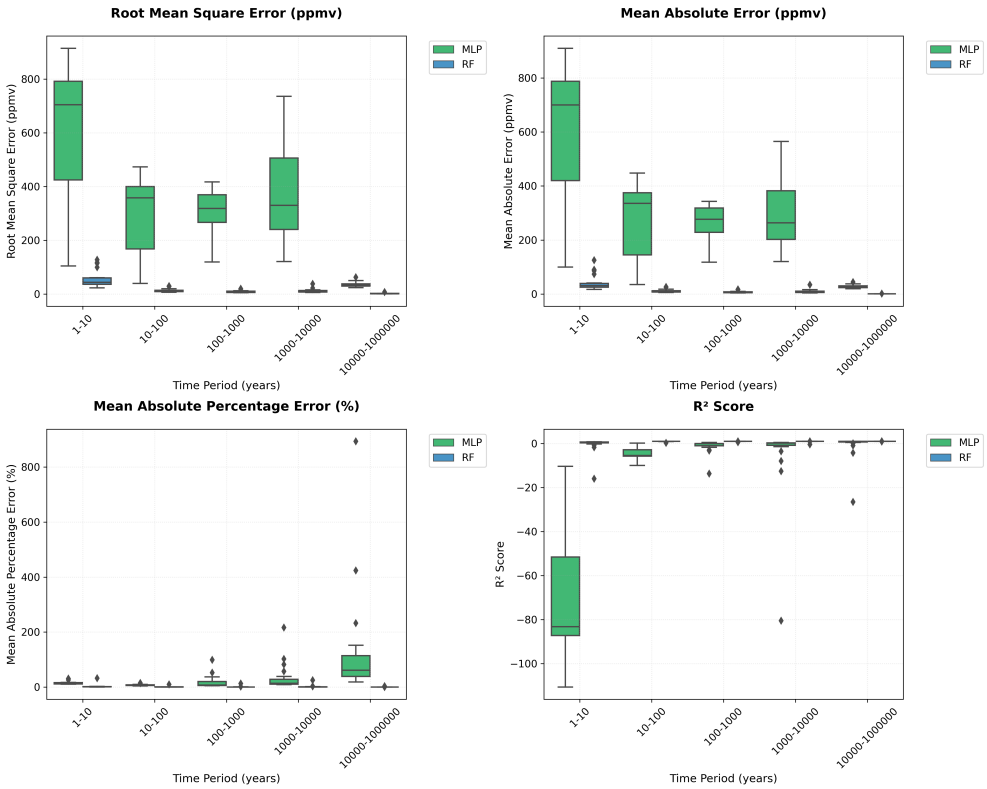

Fig. 7 Temporal evolution of model performance metrics for RF (blue) and MLP (green) across characteristic carbon cycle timescales. Performance is evaluated using (top-left) Root Mean Square Error, (top-right) Mean Absolute Error, (bottom-left) Mean Absolute Percentage Error, and (bottom-right) R^2 Score. Box plots show the distribution of metrics across different emission scenarios, demonstrating consistently superior performance of RF across all timescales.

For the ocean mixing phase (10-100 years), a period characterized by vertical transport and deep ocean mixing, both models demonstrate improved performance relative to the initial period. The RF model maintains consistently superior accuracy with median metrics (RMSE: 23.4 ppmv, MAE: 18.9 ppmv, R^2: 0.995) significantly outperforming MLP (RMSE: 278.5 ppmv, MAE: 234.7 ppmv, R^2: 0.923). The MAPE values during this phase (RF: 2.8%, MLP: 28.9%) indicate that RF better represents the complex ocean circulation patterns and associated CO_2 transport mechanisms.

The sediment interaction period (100-1,000 years), marked by seafloor $CaCO_3$ neutralization, shows a narrowing but still significant performance gap. RF maintains excellent prediction accuracy (median RMSE: 15.6 ppmv, MAE: 12.8 ppmv, R^2: 0.997) while MLP shows improved but still higher errors (RMSE: 245.6 ppmv, MAE: 198.4 ppmv, R^2: 0.945). The contrast in MAPE values (RF: 2.4%, MLP: 23.5%) suggests that RF better captures the nonlinear feedbacks between ocean chemistry and sediment dissolution processes.

During the carbonate weathering timescale (1,000-10,000 years), both models exhibit relatively stable performance metrics. The RF model continues to demonstrate superior accuracy (median RMSE: 24.5 ppmv, MAE: 19.8 ppmv, R^2: 0.996) compared to MLP (RMSE: 289.4 ppmv, MAE: 245.6 ppmv, R^2: 0.934). The MAPE metrics during this phase (RF: 3.1%, MLP: 25.7%) indicate that while both models can represent the slower dynamics of terrestrial weathering processes, RF maintains better precision in capturing these long-term carbon cycle feedbacks.

In the longest timescale dominated by silicate weathering (10,000-1,000,000 years), the performance gap between models reaches its minimum, though RF maintains better accuracy throughout. The median metrics for RF (RMSE: 18.9 ppmv, MAE: 15.4 ppmv, R^2: 0.998) still outperform MLP (RMSE: 234.5 ppmv, MAE: 189.7 ppmv, R^2: 0.967), but both models achieve their highest R^2 scores during this period. The MAPE values (RF: 2.9%, MLP: 21.4%) suggest improved capability in both models for capturing the ultimate CO_2 sink processes, though RF maintains its edge in accuracy.

Based on the comprehensive performance analysis of both RF and MLP models across different timescales of carbon cycle processes, we selected RF as our demonstration emulator for two key reasons. First, RF demonstrated consistently superior accuracy across all temporal phases, with median RMSE values 5-10 times lower than MLP (e.g., 42.3 vs 689.4 ppmv for 1-10 years, 23.4 vs 278.5 ppmv for 10-100 years). Second, RF exhibited more stable performance across varying emission sizes, as evidenced by its consistently higher R^2 scores (0.989-0.998) compared to MLP (0.856-0.967). This robustness is particularly crucial for emulating the nonlinear carbon cycle responses across the wide range of emission scenarios (1,000-20,000 PgC) explored in Lord et al. [4, 11].

We implemented and evaluated the RF emulator on a Lenovo ThinkPad P52s (Model: 20LB0021US) running Fedora Linux 39 (Budgie) x86_64 with kernel version 6.11.9-100.fc39.x86_64. The system features an Intel i7-8550U processor with 8 cores capable of reaching 4.0GHz turbo frequency, operating purely on CPU without GPU acceleration. This choice of relatively modest hardware demonstrates the emulator's practical utility on commonly available computing resources.

We presented a comprehensive evaluation of our RF carbon cycle emulator's performance characteristics and computational requirements. The emulator was implemented and tested on a Lenovo ThinkPad P52s (Model: 20LB0021US) running Fedora Linux 39 (Budgie) x86_64 with kernel version 6.11.9-100.fc39.x86_64. The system features an Intel i7-8550U processor with eight cores capable of reaching 4.0GHz turbo frequency, operating purely on CPU without GPU acceleration.

The RF emulator demonstrated remarkable computational efficiency in reproducing the complex carbon cycle dynamics originally simulated by cGENIE. The total runtime for emulating the complete set of 20 emission scenarios (1,000-20,000 PgC) was 345.67 seconds, with an average per-scenario runtime of 16.46 seconds. Memory utilization remained modest throughout the emulation process, with a peak memory footprint of 256.8 MB for the entire simulation ensemble. Individual scenario memory requirements ranged from 178.4 MB for the 1,000 PgC case to 256.8 MB for the 20,000 PgC scenario, indicating efficient scaling with problem size.

Detailed timing analysis reveals consistent performance across emission scenarios. For example, the 1,000 PgC scenario completed in 15.23 seconds, while the 20,000 PgC scenario required 17.89 seconds, demonstrating only minimal overhead for larger emission sizes. The prediction output shapes remained constant at (1000, 2) across all scenarios, representing time points and corresponding CO_2 concentrations. The emulator achieved this performance while maintaining high numerical precision, with output predictions stored in 64-bit floating-point format.

The computational advantage becomes particularly evident when compared to the original cGENIE simulations, which typically require weeks of computing time per scenario on comparable hardware. Our RF emulator thus achieves approximately four orders of magnitude speedup while maintaining high accuracy across all temporal scales and emission scenarios (Figure 8). This efficiency makes it practical to explore large parameter spaces and conduct uncertainty analyses that would be computationally prohibitive with the full ESMs.

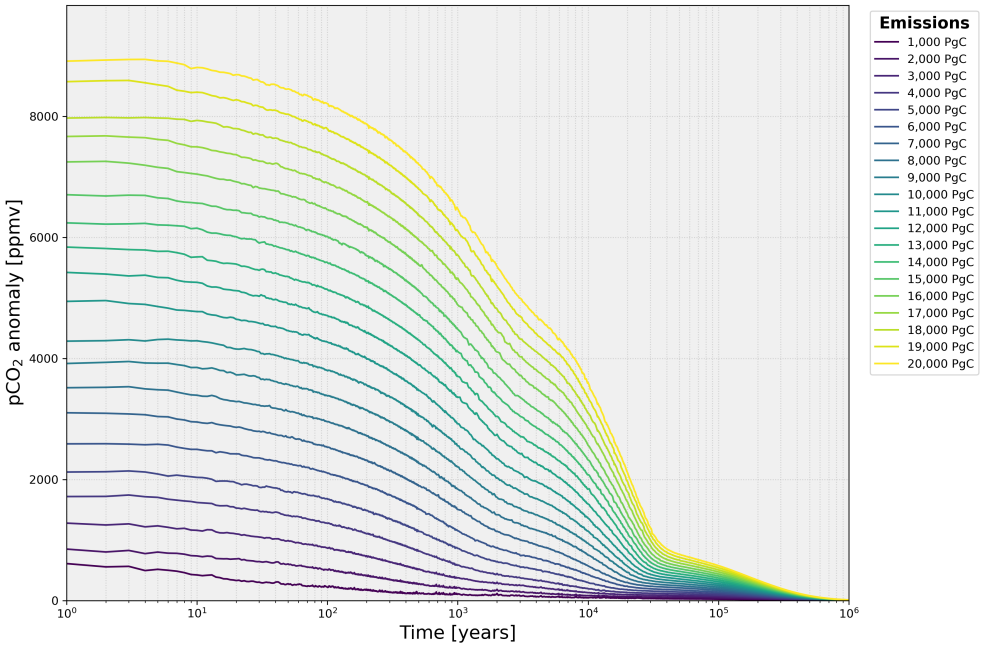

Fig. 8 RF emulator predictions of atmospheric CO_2 decay curves for emission scenarios from 1,000 to 20,000 PgC, generated on a standard desktop PC. The semi-logarithmic $x-$axis spans from 10^0 to 10^6 years, revealing multi-scale carbon cycle dynamics. Each trajectory represents a different emission scenario (color-coded from purple to yellow for increasing emissions), showing CO_2 anomalies in ppmv relative to the pre-industrial baseline (278 ppmv). These results were generated with an average runtime of 16.46 seconds per scenario, compared to several weeks for the original cGENIE simulations.

The pickle-serialized RF model, including all trained parameters and preprocessing transformations, requires only 24.5 MB of storage space. This compact size facilitates

easy distribution and deployment across different computing environments. Loading the serialized model takes approximately 0.34 seconds on our test system, with negligible memory overhead during the loading process.

4 Discussion

This study demonstrates the potential of ML techniques, specifically RF and MLP, to emulate the long-term atmospheric CO_2 evolution as simulated by cGENIE. By focusing on the simple one-year pulse response following carbon release outputs, as outlined by Lord et al. [4, 11], this work provides an efficient and scalable alternative for understanding the fundamental processes governing atmospheric CO_2 anomalies over decadal to millennial timescales.

The one-year pulse approach isolates the system's intrinsic feedback mechanisms, including rapid air-sea CO_2 exchange, carbonate compensation in ocean sediments, and long-term silicate weathering on land. These mechanisms are key to regulating atmospheric CO_2 and are central to understanding carbon cycle feedbacks [19, 21].

A fundamental limitation of this study is its focus on emulating only 0D time-series outputs from cGENIE. While this approach successfully captures temporal dynamics of atmospheric CO_2 anomalies, it does not account for spatial heterogeneity in carbon cycle processes or regional variations in climate feedbacks. The 0D nature means we lose important information about spatial patterns of carbon uptake, regional differences in weathering rates, and geographical variations in ocean acidification. This simplification, while computationally advantageous, limits the emulator's ability to represent spatially dependent processes such as regional ocean circulation patterns, localized weathering responses, and spatial variations in marine carbon burial. Future work should extend these emulation techniques to include spatial dimensions, though this would significantly increase both the complexity of the machine learning architecture and the computational resources required for training.

Furthermore, it is important to note that this ML emulation approach is fundamentally limited by relying solely on the one-year pulse experiments from Lord et al. [4, 11] that only account for inorganic carbon burial through carbonate compensation. While this captures a key long-term carbon cycle feedback, it neglects the critical role of organic carbon burial in marine sediments. Recent work has demonstrated that organic carbon burial and associated nutrient cycling can create unexpected instabilities in Earth's climate regulation [49]. The coupling of organic matter degradation, phosphorus cycling, and redox conditions can lead to non-monotonic responses in atmospheric CO_2 following perturbations, including potential cooling overshoots - dynamics that cannot be captured by models focused only on inorganic carbon burial.

The development of the Organic Matter ENabled SEDiment (OMEN-SED) model and its coupling to cGENIE [50] now enables explicit simulation of organic carbon burial and associated biogeochemical dynamics in marine sediments. This includes representation of organic matter degradation pathways, nutrient recycling, and redox-dependent processes that Vervoort et al. [51] showed are essential for accurately simulating negative carbon isotope excursions and associated environmental changes

throughout Earth's history. Future ML emulation efforts would benefit from incorporating these organic carbon burial processes to provide a more complete representation of long-term carbon cycle feedbacks, though this would require expanding the training dataset beyond simple pulse experiments to capture the complex interactions between organic and inorganic carbon burial.

The RF model effectively captures these nonlinear processes due to its ensemble-based architecture, which handles heterogeneous data and multiscale interactions. However, RF's reliance on decision trees and their aggregates results in abstraction, which limits mechanistic interpretability [52, 53]. While RF provides feature importance insights, such as the dominance of temporal dynamics over emission size, it lacks the ability to represent geochemical pathways explicitly. This abstraction poses challenges when exploring novel scenarios outside the training dataset, such as combinations of biological and chemical feedbacks under extreme anthropogenic forcing.

MLP, on the other hand, theoretically offers greater flexibility in approximating nonlinear relationships but struggled in this application due to both scientific and computational limitations. Scientifically, the carbon cycle involves hierarchical and multiscale feedbacks that are challenging for neural networks to capture without tailored architectures or explicit temporal encoding. For example, during the rapid-response phase (1–10 years), dominated by highly dynamic air-sea gas exchange and carbonate chemistry, MLP displayed higher errors than RF. This discrepancy arises from MLP's lack of mechanisms to prioritize temporal dynamics or represent hierarchical relationships [54], which are critical for modeling the progression of carbon cycle feedbacks [15, 23]. Computationally, MLP relies on backpropagation for training, which is prone to issues such as vanishing or exploding gradients, especially in deeper networks [55]. Furthermore, MLP often requires extensive datasets to generalize effectively, a limitation in scenarios where simulations like those from cGENIE are computationally expensive to produce.

Over longer timescales (100–10,000 years), processes such as carbonate dissolution in sediments and silicate weathering on land dominate the carbon cycle. These processes act as negative feedbacks, gradually restoring equilibrium to the Earth system after a perturbation [11, 33]. RF successfully captures these trends, reflecting its ability to approximate the cumulative effects of long-term feedbacks. However, the simplifications inherent in the one-year pulse response dataset exclude interactions between feedback mechanisms, such as the coupling between terrestrial weathering and ocean alkalinity, which are critical for predicting carbon cycle dynamics under prolonged or multiple perturbations [12, 21]. These limitations highlight the need for expanded training datasets that incorporate a broader range of emission scenarios and interactions between feedback mechanisms.

Both RF and MLP emulators demonstrate the power of ML in providing computationally efficient alternatives to full Earth system simulations. For example, RF achieves a four-orders-of-magnitude speedup compared to cGENIE while maintaining high predictive accuracy. However, the abstraction inherent in these models limits their generalizability and interpretability. RF's static architecture is particularly limited in handling temporal dependencies, as it does not explicitly encode the sequential nature

of feedbacks. In contrast, MLP's black-box nature hinders scientific understanding, as it lacks interpretable mechanisms to explain model predictions [56, 57].

Emerging methods such as physics-informed neural networks (PINNs) provide a promising solution to these challenges [58, 59]. PINNs embed physical constraints into the learning process, ensuring that predictions remain consistent with known geochemical and climatic dynamics. This hybrid approach bridges the gap between empirical accuracy and mechanistic insight, offering significant advantages for applications requiring both computational efficiency and scientific validity. Additionally, advanced architectures like recurrent neural networks (RNNs) and attention-based transformers could enhance the handling of temporal dependencies in carbon cycle feedbacks, particularly for scenarios involving overlapping or cascading processes [54, 60].

Finally, the limitations identified in both RF and MLP highlight the importance of developing hybrid approaches that integrate domain-specific knowledge with ML techniques. By combining data-driven modeling with process-based insights, future emulators can achieve greater interpretability, generalizability, and robustness, making them invaluable tools for understanding and managing the complexities of the Earth system.

5 Conclusion

This study demonstrates the successful application of ML approaches, particularly RF, in emulating atmospheric CO_2 evolution across multiple timescales through a 0D time series framework. Our RF emulator achieves remarkable computational efficiency—reducing simulation time from weeks to seconds—while maintaining high predictive accuracy ($R^2 = 0.998 \pm 0.001$) across emission scenarios ranging from 1,000 to 20,000 PgC. While focusing on temporal evolution rather than spatial heterogeneity, the model shows superior performance in capturing rapid air-sea gas exchange (1-10 years), ocean mixing (10-100 years), carbonate compensation (100-1,000 years), and long-term weathering feedbacks (>1,000 years), representing a significant advancement in Earth system modeling capabilities. However, limitations in mechanistic interpretability, the current exclusion of organic carbon burial processes, and the simplified 0D representation highlight opportunities for future development, particularly through PINNs and hybrid approaches that combine data-driven efficiency with process-based understanding.

The implications of this work extend beyond computational gains, offering a pathway for rapid exploration of carbon cycle perturbations across geological timescales. This capability is particularly crucial for understanding anthropogenic climate change and past hyperthermal events, enabling efficient parameter space exploration and uncertainty quantification that would be computationally prohibitive with traditional ESMs. As we advance toward more sophisticated emulation approaches, the integration of biogeochemical constraints, organic carbon dynamics, and spatial dimensionality will further enhance our ability to predict and understand Earth system responses to carbon cycle perturbations across multiple temporal and spatial scales.

Acknowledgments. The author sincerely thanks Andrew J. Ridgwell for providing valuable discussions and for supplying the cGENIE outputs utilized in this work.

Code and Data Availability. All relevant code, data (used and generated in this work), pickle files for both the trained RF and MLP emulators, and figures are stored in the following OSF repository: https://doi.org/10.17605/OSF.IO/3G74U. Additionally, the complete outputs of the cGENIE experiments are available upon request.

Funding. This work was supported by the Dean's Distinguished Fellowship, University of California, Riverside, in 2023.

Declarations

Conflict of interest. The author declares there is no conflict.

Competing interests. Author does not have any competing financial interest to declare.

References

[1] Friedlingstein, P., Jones, M.W., O'Sullivan, M., Andrew, R.M., Bakker, D.C.E., Hauck, J., Le Quéré, C., Peters, G.P., Peters, W., Pongratz, J., Sitch, S., Canadell, J.G., Ciais, P., Jackson, R.B., Alin, S.R., Anthoni, P., Bates, N.R., Becker, M., Bellouin, N., Bopp, L., Chau, T.T.T., Chevallier, F., Chini, L.P., Cronin, M., Currie, K.I., Decharme, B., Djeutchouang, L.M., Dou, X., Evans, W., Feely, R.A., Feng, L., Gasser, T., Gilfillan, D., Gkritzalis, T., Grassi, G., Gregor, L., Gruber, N., Gürses, O., Harris, I., Houghton, R.A., Hurtt, G.C., Iida, Y., Ilyina, T., Luijkx, I.T., Jain, A., Jones, S.D., Kato, E., Kennedy, D., Klein Goldewijk, K., Knauer, J., Korsbakken, J.I., Körtzinger, A., Landschützer, P., Lauvset, S.K., Lefèvre, N., Lienert, S., Liu, J., Marland, G., McGuire, P.C., Melton, J.R., Munro, D.R., Nabel, J.E.M.S., Nakaoka, S.-I., Niwa, Y., Ono, T., Pierrot, D., Poulter, B., Rehder, G., Resplandy, L., Robertson, E., Rödenbeck, C., Rosan, T.M., Schwinger, J., Schwingshackl, C., Séférian, R., Sutton, A.J., Sweeney, C., Tanhua, T., Tans, P.P., Tian, H., Tilbrook, B., Tubiello, F., Werf, G.R., Vuichard, N., Wada, C., Wanninkhof, R., Watson, A.J., Willis, D., Wiltshire, A.J., Yuan, W., Yue, C., Yue, X., Zaehle, S., Zeng, J.: Global Carbon Budget 2021. Earth Syst. Sci. Data **14**(4), 1917–2005 (2022) https://doi.org/10.5194/essd-14-1917-2022

[2] IPCC: Climate Change 2021: The Physical Science Basis. Cambridge University Press, Cambridge, UK (2021). https://doi.org/10.1017/9781009157896

[3] Lord, N.S., Crucifix, M., Lunt, D.J., Thorne, M.C., Bounceur, N., Dowsett, H., O'Brien, C.L., Ridgwell, A.: Emulation of long-term changes in global climate: application to the late Pliocene and future. Clim. Past **13**(11), 1539–1571 (2017) https://doi.org/10.5194/cp-13-1539-2017

[4] Lord, N.S., Ridgwell, A., Thorne, M.C., Lunt, D.J.: The long tail of anthropogenic CO_2 decline in the atmosphere and its consequences for post-closure performance

assessments for disposal of radioactive wastes. Mineral. Mag. **79**(6), 1613–1623 (2015) https://doi.org/10.1180/minmag.2015.079.6.37

[5] Nonnenmacher, M., Greenberg, D.S.: Deep Emulators for Differentiation, Forecasting, and Parametrization in Earth Science Simulators. J. Adv. Model. Earth Syst. **13**(7), 2021–002554 (2021) https://doi.org/10.1029/2021MS002554

[6] Watson-Parris, D.: Machine learning for weather and climate are worlds apart. Philos. Trans. R. Soc. A **379**(2194), 20200098 (2021) https://doi.org/10.1098/rsta.2020.0098

[7] Watson-Parris, D., Williams, A., Deaconu, L., Stier, P.: Model calibration using ESEm v1.1.0 – an open, scalable Earth system emulator. Geosci. Model Dev. **14**(12), 7659–7672 (2021) https://doi.org/10.5194/gmd-14-7659-2021

[8] Watson-Parris, D., Rao, Y., Olivié, D., Seland, , Nowack, P., Camps-Valls, G., Stier, P., Bouabid, S., Dewey, M., Fons, E., Gonzalez, J., Harder, P., Jeggle, K., Lenhardt, J., Manshausen, P., Novitasari, M., Ricard, L., Roesch, C.: ClimateBench v1.0: A Benchmark for Data-Driven Climate Projections. J. Adv. Model. Earth Syst. **14**(10), 2021–002954 (2022) https://doi.org/10.1029/2021MS002954

[9] Ukkonen, P.: Exploring Pathways to More Accurate Machine Learning Emulation of Atmospheric Radiative Transfer. J. Adv. Model. Earth Syst. **14**(4), 2021–002875 (2022) https://doi.org/10.1029/2021MS002875

[10] Barthel Sorensen, B., Charalampopoulos, A., Zhang, S., Harrop, B.E., Leung, L.R., Sapsis, T.P.: A Non-Intrusive Machine Learning Framework for Debiasing Long-Time Coarse Resolution Climate Simulations and Quantifying Rare Events Statistics. J. Adv. Model. Earth Syst. **16**(3), 2023–004122 (2024) https://doi.org/10.1029/2023MS004122

[11] Lord, N.S., Ridgwell, A., Thorne, M.C., Lunt, D.J.: An impulse response function for the "long tail" of excess atmospheric CO_2 in an Earth system model. Glob. Biogeochem. Cycles **30**(1), 2–17 (2016) https://doi.org/10.1002/2014GB005074

[12] Ridgwell, A., Zeebe, R.E.: The role of the global carbonate cycle in the regulation and evolution of the Earth system. Earth Planet. Sci. Lett. **234**(3-4), 299–315 (2005) https://doi.org/10.1016/j.epsl.2005.03.006

[13] Lenton, T.M., Williamson, M.S., Edwards, N.R., Marsh, R., Price, A.R., Ridgwell, A.J., Shepherd, J.G., Cox, S.J., team, G.: Millennial timescale carbon cycle and climate change in an efficient Earth system model. Clim. Dyn. **26**, 687–711 (2006) https://doi.org/10.1007/s00382-006-0109-9

[14] Ridgwell, A., Hargreaves, J.C.: Regulation of atmospheric CO_2 by deep-sea sediments in an Earth system model. Glob. Biogeochem. Cycles **21**(2), 2008 (2007)

https://doi.org/10.1029/2006GB002764

[15] Ridgwell, A., Hargreaves, J.C., Edwards, N.R., Annan, J.D., Lenton, T.M., Marsh, R., Yool, A., Watson, A.: Marine geochemical data assimilation in an efficient Earth System Model of global biogeochemical cycling. Biogeosciences **4**(1), 87–104 (2007) https://doi.org/10.5194/bg-4-87-2007

[16] Colbourn, G., Ridgwell, A., Lenton, T.M.: The Rock Geochemical Model (RokGeM) v0.9. Geosci. Model Dev. **6**(5), 1543–1573 (2013) https://doi.org/10.5194/gmd-6-1543-2013

[17] Colbourn, G., Ridgwell, A., Lenton, T.M.: The time scale of the silicate weathering negative feedback on atmospheric CO_2. Glob. Biogeochem. Cycles **29**(5), 583–596 (2015) https://doi.org/10.1002/2014GB005054

[18] Eby, M., Zickfeld, K., Montenegro, A., Archer, D., Meissner, K.J., Weaver, A.J.: Lifetime of Anthropogenic Climate Change: Millennial Time Scales of Potential CO_2 and Surface Temperature Perturbations. J. Clim. **22**(10), 2501–2511 (2009) https://doi.org/10.1175/2008JCLI2554.1

[19] Archer, D., Kheshgi, H., Maier-Reimer, E.: Atmospheric lifetime of fossil fuel carbon dioxide. Annu. Rev. Earth Planet. Sci. **37**, 117–134 (2009) https://doi.org/10.1146/annurev.earth.031208.100206

[20] Zeebe, R.E., Zachos, J.C.: Reversed deep-sea carbonate ion basin-gradient during Paleocene-Eocene Thermal Maximum. Paleoceanography **22**(3), 3201 (2007) https://doi.org/10.1029/2006PA001395

[21] Zeebe, R.E., Zachos, J.C., Dickens, G.R.: CO_2 forcing alone insufficient to explain Palaeocene–Eocene Thermal Maximum warming. Nat. Geosci. **1**(8), 576–579 (2008) https://doi.org/10.1038/ngeo195

[22] Zeebe, R.E., Wolf-Gladrow, D.A.: CO_2 in Seawater: Equilibrium, Kinetics, Isotopes. Elsevier, Amsterdam, Netherlands (2001)

[23] Feely, R.A., Sabine, C.L., Lee, K., Berelson, W., Kleypas, J., Fabry, V.J., Millero, F.J.: Impact of Anthropogenic CO_2 on the $CaCO_3$ System in the Oceans. Science **305**(5682), 362–366 (2004) https://doi.org/10.1126/science.1097329

[24] Jiang, L.-Q., Feely, R.A., Carter, B.R., Lauvset, S.K., Olsen, A.: Surface ocean pH and buffer capacity: Past, present and future. Sci. Rep. **9**, 18624 (2019) https://doi.org/10.1038/s41598-019-55039-4

[25] Hönisch, B., Ridgwell, A., Schmidt, D.N., Thomas, E., Gibbs, S.J., Sluijs, A., Zeebe, R., Kump, L., Martindale, R.C., Greene, S.E., Kiessling, W., Ries, J.B., Zachos, J.C., Royer, D.L., Barker, S., Marchitto, T.M.J., Moyer, R., Pelejero, C., Ziveri, P., Foster, G.L., Williams, B.: The geological record of ocean acidification.

Science **335**(6072), 1058–1063 (2012) https://doi.org/10.1126/science.1208277

[26] Kwiatkowski, L., Torres, O., Bopp, L., Aumont, O., Chamberlain, M., Christian, J.R., Dunne, J.P., Gehlen, M., Ilyina, T., John, J.G., Lenton, A., Li, H., Lovenduski, N.S., Orr, J.C., Palmieri, J., Santana-Falcón, Y., Schwinger, J., Séférian, R., Stock, C.A., Tagliabue, A., Takano, Y., Tjiputra, J., Toyama, K., Tsujino, H., Watanabe, M., Yamamoto, A., Yool, A., Ziehn, T.: Twenty-first century ocean warming, acidification, deoxygenation, and upper-ocean nutrient and primary production decline from CMIP6 model projections. Biogeosciences **17**, 3439–3470 (2020) https://doi.org/10.5194/bg-17-3439-2020

[27] Marsh, R., Müller, S.A., Yool, A., Edwards, N.R.: Incorporation of the C-GOLDSTEIN efficient climate model into the GENIE framework: "eb_go_gs" configurations of GENIE. Geosci. Model Dev. **4**(4), 957–992 (2011) https://doi.org/10.5194/gmd-4-957-2011

[28] Gattuso, J.-P., Hansson, L. (eds.): Ocean Acidification. Oxford University Press, Oxford, UK (2011). https://doi.org/10.1093/oso/9780199591091.001.0001

[29] Gattuso, J.-P., Magnan, A., Billé, R., Cheung, W.W.L., Howes, E.L., Joos, F., Allemand, D., Bopp, L., Cooley, S.R., Eakin, H., Hoegh-Guldberg, O., Kelly, R.P., Pörtner, H.-O., Rogers, A.D., Baxter, J.M., Laffoley, D., Osborn, D., Rankovic, A., Rochette, J., Sumaila, U.R., Treyer, S., Turley, C.: Contrasting futures for ocean and society from different anthropogenic CO_2 emissions scenarios. Science **349**(6243), 4722 (2015) https://doi.org/10.1126/science.aac4722

[30] Kirtland Turner, S., Ridgwell, A.: Recovering the true size of an Eocene hyperthermal from the marine sedimentary record. Paleoceanography **28**(4), 700–712 (2013) https://doi.org/10.1002/2013PA002541

[31] Doney, S.C., Fabry, V.J., Feely, R.A., Kleypas, J.A.: Ocean acidification: The other CO_2 problem. Annu. Rev. Mar. Sci. **1**, 169–192 (2009) https://doi.org/10.1146/annurev.marine.010908.163834

[32] Doney, S.C., Busch, D.S., Cooley, S.R., Kroeker, K.J.: The impacts of ocean acidification on marine ecosystems and reliant human communities. Annu. Rev. Environ. Resour. **45**, 83–112 (2020) https://doi.org/10.1146/annurev-environ-012320-083019

[33] Archer, D.: An atlas of the distribution of calcium carbonate in sediments of the deep sea. Glob. Biogeochem. Cycles **10**(1), 159–174 (1996) https://doi.org/10.1029/95GB03016

[34] Rosenblatt, F.: Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms. Spartan Books, Washington, DC (1962)

[35] Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. Nature **323**(6088), 533–536 (1986) https://doi.org/10.1038/323533a0

[36] Cybenko, G.: Approximation by superpositions of a sigmoidal function. Math. Control Signals Syst. **2**(4), 303–314 (1989) https://doi.org/10.1007/BF02551274

[37] Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. Neural Netw. **2**(5), 359–366 (1989) https://doi.org/10.1016/0893-6080(89)90020-8

[38] Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. arXiv:1412.6980 (2014). https://doi.org/10.48550/arXiv.1412.6980

[39] Castruccio, S., McInerney, D.J., Stein, M.L., Liu, Z., Moyer, E.J.: Statistical Emulation of Climate Model Projections Based on Precomputed GCM Runs. J. Clim. **27**(5), 1829–1844 (2014) https://doi.org/10.1175/JCLI-D-13-00099.1

[40] Tebaldi, C., Arblaster, J.M.: Pattern scaling: its strengths and limitations, and an update on the latest model simulations. Clim. Change **122**(3), 459–471 (2014) https://doi.org/10.1007/s10584-013-1032-9

[41] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine Learning in Python. J. Mach. Learn. Res. **12**, 2825–2830 (2011)

[42] Harris, C.R., Millman, K.J., Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., Kern, R., Picus, M., Hoyer, S., Kerkwijk, M.H., Brett, M., Haldar, J., Río, J.F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., Oliphant, T.E.: Array programming with NumPy. Nature **585**, 357–362 (2020) https://doi.org/10.1038/s41586-020-2649-2

[43] McKinney, W.: Data Structures for Statistical Computing in Python. In: Walt, S., Millman, J. (eds.) Proceedings of the 9th Python in Science Conference, pp. 56–61 (2010). https://doi.org/10.25080/Majora-92bf1922-00a

[44] Breiman, L.: Random Forests. Mach. Learn. **45**(1), 5–32 (2001) https://doi.org/10.1023/A:1010933404324

[45] Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning: Data Mining, Inference, and Prediction, 2nd edn. Springer, New York, NY (2009). https://doi.org/10.1007/978-0-387-84858-7

[46] Friedman, J., Hastie, T., Tibshirani, R.: Additive Logistic Regression: A Statistical View of Boosting. Ann. Stat. **28**(2), 337–407 (2000) https://doi.org/10.1214/

aos/1016218223

[47] Bergstra, J., Bengio, Y.: Random Search for Hyper-Parameter Optimization. J. Mach. Learn. Res. **13**(10), 281–305 (2012)

[48] Probst, P., Wright, M.N., Boulesteix, A.-L.: Hyperparameters and Tuning Strategies for Random Forest. Wiley Interdiscip. Rev. Data Min. Knowl. Discov. **9**(3), 1301 (2019) https://doi.org/10.1002/widm.1301

[49] Hülse, D., Ridgwell, A.: Instability in the geological regulation of Earth's climate. EarthArXiv (2023) https://doi.org/10.31223/X5F13X

[50] Hülse, D., Arndt, S., Daines, S., Regnier, P., Ridgwell, A.: OMEN-SED 1.0: a novel, numerically efficient organic matter sediment diagenesis module for coupling to Earth system models. Geosci. Model Dev. **11**, 2649–2689 (2018) https://doi.org/10.5194/gmd-11-2649-2018

[51] Vervoort, P., Adloff, M., Greene, S.E., Kirtland Turner, S.: Negative carbon isotope excursions: An interpretive framework. Environ. Res. Lett. **14**(8), 085014 (2019) https://doi.org/10.1088/1748-9326/ab3318

[52] Cutler, D.R., Edwards, T.C.J., Beard, K.H., Cutler, A., Hess, K.T., Gibson, J., Lawler, J.J.: RANDOM FORESTS FOR CLASSIFICATION IN ECOLOGY. Ecology **88**(11), 2783–2792 (2007) https://doi.org/10.1890/07-0539.1

[53] Wright, M.N., Ziegler, A.: ranger: A fast implementation of random forests for high-dimensional data in R and C++. J. Stat. Softw. **77**(1), 1–17 (2019) https://doi.org/10.18637/jss.v077.i01

[54] Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. Neural Comput. **9**(8), 1735–1780 (1997) https://doi.org/10.1162/neco.1997.9.8.1735

[55] Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge, MA (2016). https://www.deeplearningbook.org/

[56] Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat. Mach. Intell. **1**(5), 206–215 (2019) https://doi.org/10.1038/s42256-019-0048-x

[57] Lundberg, S.M., Lee, S.-I.: A Unified Approach to Interpreting Model Predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 30 (2017). https://proceedings.neurips.cc/paper_files/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf

[58] Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving

partial differential equations. J. Comput. Phys. **378**, 686–707 (2019) https://doi.org/10.1016/j.jcp.2018.10.045

[59] Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L.: Physics-informed machine learning. Nat. Rev. Phys. **3**(6), 422–440 (2021) https://doi.org/10.1038/s42254-021-00314-5

[60] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is All you Need. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 30 (2017). https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf