This is a non-peer-reviewed preprint submitted to EarthArXiv.

This manuscript has been submitted for publication in *Geophysical Research Letters*. Please note the manuscript has yet to be formally accepted for publication. Subsequent versions of this manuscript may have slightly different content. If accepted, the final version of this manuscript will be available via the 'Peer-reviewed Publication DOI' link on the right-hand side of this webpage. Please feel free to contact any of the authors; we welcome feedback.

Integrating machine learning with a process-based model for estimating global wetland methane emissions

Chris C R Smith[1,2], Shuo Chen[3], Sparkle L. Malone[4], Gavin McNicol[5], Qing Zhu[6], Licheng Liu[7], and Youmi Oh[1,2]

[1]Cooperative Institute for Research in Environmental Sciences, University of Colorado Boulder

[2]NOAA Global Monitoring Laboratory

[3]Earth, Atmospheric, and Planetary Sciences, Purdue University

[4]Yale School of the Environment, Yale University

[5]University of Illinois Chicago

[6]Lawrence Berkeley National Laboratory

[7]University of Wisconsin Madison

Key Points

- Together, process-based model outputs and satellite imagery improved machine learning accuracy.
- Conditional model selection based on local environmental inputs circumvents model generalization limitations using machine learning.
- These advancements yield a global emissions estimate, in the upper range previous estimates.

Abstract

Estimates of methane emissions from natural wetlands are uncertain and depend on the modeling approach. Process-based models incorporate knowledge of the underlying biogeochemistry, but prediction accuracy is sensitive to parameterization and empirical equations. Machine learning models have potential to improve estimates, however they struggle to generalize to new prediction sites. We explore combined process-based machine learning strategies. Prediction accuracy improved when two new inputs were incorporated: the estimate from a process-based surrogate model and satellite imagery from each measurement site. We estimated global-scale emissions using conditional model selection: a machine learning model was applied to grid cells within the range of environmental conditions used for training, and a process-based model was used for more extreme regions. The hybrid approach resulted in a global emissions estimate of 180.9 teragrams of methane per year, which was unexpectedly in the upper range of estimates from either individual model.

Plain language summary

Wetlands produce methane, a critical greenhouse gas, but it is challenging to measure how much is produced globally. Various analytical methods are used to estimate such emissions, including process-based simulations of how wetlands naturally produce methane and newer machine learning methods that learn patterns from data. Each approach has strengths and weaknesses: process-based simulations are grounded in known processes but are not very accurate, while machine learning models can be more accurate but may fail when applied to new geographic regions. In this study, the two approaches are combined to improve emission estimates. In particular, we found that carefully choosing which model to use depending on local environmental conditions gave more reliable estimates across the globe. Overall, the combined approach gave a global estimate of 180.9 teragrams of methane per year from wetlands, which is higher than estimates from either method alone. By combining physical understanding with data-driven methods, our findings improve confidence in how much wetlands contribute to global methane emissions.

Introduction

Methane emissions from natural wetlands are a major contributor to the atmospheric methane budget and changing climate (Saunois et al., 2025). While atmospheric methane concentrations have increased since the onset of industrialization, concentrations were relatively stable during the brief period from 2000 to 2006. Since 2006, global emissions have increased sharply (Lan et al., 2022), with one of the main contributors believed to be increased emissions from wetlands driven by changes in meteorology and inundation (Qu et al.,

2024; Lin et al., 2024; Oh et al., 2026). These changes in wetland processes, together with climate–carbon feedbacks, are projected to play a key role in global methane dynamics (Kuhn et al., 2025; Treat et al., 2024). Therefore, to fully characterize the global methane cycle and ultimately guide mitigation strategies, a crucial goal is to accurately quantify emissions produced by wetlands. However, current estimates of global wetland emissions are highly uncertain—146 to 212 $\mathrm{TgCH_4\,yr^{-1}}$ (Liu et al., 2020; McNicol et al., 2023)—and depend strongly on which modeling approach is used (Lin et al., 2024).

One common approach for studying wetland methane dynamics is process-based modeling (Zhang et al., 2025). Process-based models typically consist of equations that incorporate scientific knowledge about the biogeochemical process being modeled and are parameterized using theory and empirical observations. For example, process-based ecosystem models integrate environmental variables such as temperature, soil moisture, and substrate availability to represent the mechanisms that control methane production, oxidation, and transport (Xu et al., 2016). Despite incorporating prior scientific knowledge, process-based models still have substantial knowledge gaps: for example, the relationships between soil properties (e.g., pH) and methane emissions are not fully understood (Knox et al., 2021). Some process-based models use relatively few model parameters, making them a natural choice for global-scale analyses because they provide a principled way to extrapolate beyond sparse observations. However, evaluation of process-based models on ground truth observations has shown systemic mismatches between model outputs and eddy covariance tower measurements, partly due to differences in spatial scale. For example, Liu et al. (2020) reported an R^2 value 0.41 on evaluation sites using the Terrestrial Ecosystem Model (TEM; Zhuang et al., 2004, 2013).

Machine learning may alleviate some limitations of process-based models. In particular, deep neural networks and random forests are well-suited for modeling complex interactions among environmental variables where the likelihood function is unavailable, and without prior scientific knowledge about the modeled process (Chen et al., 2024). However, machine learning models are often challenged to extrapolate beyond the training distribution, e.g, applied to new environments, and typically require large datasets for training. Despite the limited availability of measurement sites from natural wetlands, machine learning approaches perform surprisingly well at estimating methane emissions, achieving even better accuracy than process-based models. For example, McNicol et al. (2023) achieved an R^2 of 0.54 using random forests and estimated a global wetland emission of 146 Tg $\mathrm{CH_4}$ $\mathrm{yr^{-1}}$, although the evaluation sites used in this study only partly overlap with those of Liu et al. (2020). The McNicol et al. estimate is smaller than that of TEM—158$\pm$24 Tg $\mathrm{CH_4\,yr^{-1}}$ (Zhang et al., 2025)—which might be attributable to underestimated emissions in tropical regions where training data is most limited.

The most useful quantification of wetland emissions might be achieved by combining the two divergent modeling strategies (Beckh et al., 2021; Von Rueden et al., 2021; Zobeiry and Poursartip, 2021; Cui et al.,

2023; Karpatne et al., 2024). Ideally, knowledge from the process-based model would be retained, freeing the neural network to focus only on the unknown parts of the process. Such an approach would be more efficient than naive machine learning, and may improve spatial extrapolation and model interpretability. Similar approaches have been used in the environmental sciences. For example, Liu et al. (2024) used a "knowledge-guided" framework to model agricultural CO_2 emissions, which involved (i) pre-training on process-based model outputs, (ii) a customized neural network architecture, and (iii) incorporating prior knowledge into the loss function. In the context of natural methane, we explored several techniques integrating TEM outputs with machine learning—an example of transfer learning—and found marginal benefits of transfer learning for spatial upscaling under some parameter regimes (Sun et al., 2026). The current analysis was conducted in parallel with that of Sun et al., and involves different study sites, inputs, models, and other parameters.

Here, we develop combined process-based machine learning strategies for estimating wetland emissions globally. First, we design a deep neural network for use with limited training data and experiment with various input variables including (i) local tower measurements, (ii) reanalysis data products, and (iii) satellite images. Second, we evaluate four transfer learning approaches for combining a neural network with TEM: pre-training, domain adaptation, mixing simulations and real data in the training set, and stacking model outputs across domains. Last, we extrapolate to the global scale using a conditional model selection procedure and interpret our results in the context of previous global estimates.

Materials and Methods

Process-based model simulations

TEM is a commonly used biogeochemistry model that incorporates modules for methane dynamics, carbon and nitrogen cycling, soil thermal dynamics, and hydrology, and has been extensively evaluated in previous studies (Zhuang et al., 2004; Liu et al., 2020; Oh et al., 2022; Zhang et al., 2025). The CH_4 dynamics module simulates CH_4 production and oxidation, as well as three transport processes between soil and atmosphere: diffusion, ebullition, and plant-mediated transport.

We conducted three independent experiments using the TEM methane dynamics module to simulate: (1) gridded CH_4 emissions, (2) gridded CH_4 intensity, and (3) site-level CH_4 emissions. All simulations covered the study period from 2006 through 2018. Global, gridded emissions and intensity were generated with a spatial resolution of $0.5^\circ \times 0.5^\circ$, whereas the site-level emissions used some measurements from flux tower sites and gridded inputs otherwise. The first set of inputs were spatially-explicit drivers of ERA-Interim (Dee et al., 2011) daily climate data: vegetation type; topsoil bulk density; sand, silt, and clay fractions;

soil pH; wetland type; elevation; global annual CO_2 concentration; and global annual CH_4 concentration. Historical climate data from ERA-Interim were also used as input: air temperature, precipitation, vapor pressure, and cloudiness. All other parameters were the same as in Liu et al. (2020) except for inundation area fractionation; the gridded emissions simulation used the GLWD-SWAMPS inundation map (Jensen and Mcdonald, 2019; Lehner and Döll, 2004), whereas for methane intensity the grid cells were assumed to be 100% inundated. Both gridded TEM simulations used daily resolution. Site-level CH_4 emissions were simulated using climate variables available for measurement sites (described below) together with the other drivers listed above, and any missing climate data points filled using the gridded ERA-Interim daily climate data. Half-hourly measurements were converted to daily resolution for site-level simulations. All simulation outputs were averaged to monthly resolution for comparison with machine learning models.

In situ measurements

FLUXNET

Half-hourly methane fluxes and other tower measurements from January 2006 to December 2018 were downloaded from FLUXNET-CH4 Version 1.0 (`https://fluxnet.org/data/download-data/`) (Delwiche et al., 2021; Knox et al., 2021). Of the available measurement sites, we retained the 43 non-agricultural wetland sites for analysis, and manually assigned wetland type classifications to each tower location based on their site descriptions and metadata (Table S1).

For training machine learning models, we used methane fluxes that were gap-filled using the marginal distribution sampling method ("FCH4_F"); we purposefully avoided using data that were gap-filled using a neural network trained on the same data analyzed in the current study ("FCH4_F_ANNOPTLM"). For model evaluation, we only compare model outputs to real observations ("FCH4").

Spectral data

All seven spectral bands from the Moderate Resolution Imaging Spectroradiometer (MODIS) surface reflectance (8-day, 500 m) (Vermote, 2021) were downloaded from Google Earth Engine (Gorelick et al., 2017) in a 5 km $\times$ 5 km area centered on each tower site. The associated quality bitmask was used to retain time points with high quality scores for every band. Linear interpolation was used to fill missing time points, and missing pixels were filled using linear row-wise interpolation followed by column-wise interpolation.

Choice of machine learning inputs

We explored several potential inputs. FLUXNET provides several measurements in addition to methane flux that are shared across the evaluated sites: air temperature; precipitation; latent heat flux; presence of various moss, tree and shrub types; net ecosystem exchange; sensible heat flux; latent heat flux; shortwave radiation, incoming; gross primary productivity; ecosystem respiration; wind speed; longwave radiation, incoming; vapor pressure deficit; and atmospheric pressure. In addition, we downloaded several MODIS 500m-resolution, dynamic data products from Google Earth Engine: tree canopy height, leaf area index, land cover classification, vegetation cover, land cover dynamics, vegetation indices, evapotranspiration, and GPP.

Each variable was initially screened using a simple linear regression against methane flux. The strongest linear predictor was latent heat flux (Pearson correlation, $r = 0.63$), and temperature also showed a strong correlation ($r = 0.50$). This simple assessment was consistent with a previous FLUXNET analysis showing that temperature is a dominant predictor of methane flux at the seasonal scale, followed by latent heat flux and water table depth (Knox et al., 2021); however, water table depth was not available at all sites analyzed in the present study. Precipitation showed no correlation with methane flux ($r = -0.01$), because the analyzed sites are perpetually inundated.

Ultimately, the variables used for model comparison, transfer learning, and upscaling included: (1) latent heat flux, (2) temperature, (3) wetland type, and, for some models, (4) MODIS reflectance. It is possible that other, more complex interactions exist between the examined (or other) variables, however it is beyond the scope of the current study to explore all combinations of inputs.

Preprocessing inputs for machine learning

Monthly averages were calculated for all inputs, including TEM-generated synthetic data, FLUXNET-CH4 measurements, and MODIS images. Monthly data points were retained if each input and target variable had at least 14% non-missing data, similar to the threshold from McNicol et al. (2023) requiring one day of data per week. Otherwise, all inputs and the target for a time point were changed to a missing data value (-9999.0). TEM wetland types—(i) forested bog, (ii) unforested bog, (iii) forested swamp, (iv) unforested swamp, or (v) alluvial formation—were one-hot encoded, with each category represented as a separate binary input variable. Information on choosing inputs for machine learning is described in the Supporting Information.

Baseline machine learning model

Architecture

PyTorch (Ansel et al., 2024) was used to build several neural network models, each stemming from the following baseline architecture and using the same settings. Because emissions may respond to changes in predictor variables with a temporal lag, we implemented a recurrent neural network (RNN; Figure 1), which allows information to propagate towards later time points and can thereby remember past events in the time series. The input is a time series of predictor variables (described below) and the output is a methane emission for each time point (i.e., sequence-to-sequence prediction). Specifically, two gated recurrent unit (GRU; Cho et al., 2014) layers were used with eight neurons, followed by a dense layer also with eight neurons, a rectified linear unit (ReLU), and a final dense layer that outputs a single value for each time point. The small model size was designed to avoid overfitting, which was a challenge due to limited training examples.

Inputs were fed to the neural network as 24-month time series, however the first 12 monthly time points in a window served to initialize the recurrent hidden state ("spin-up") and we only evaluated model performance on the later 12 time points. To align seasons across hemispheres, each window from northern hemisphere sites was forced to start in January, whereas windows for southern hemisphere sites started in July. We retained windows with at least four months of non-missing input data in the first year, and at least four months of non-missing input and target data ("FCH4") in the second year; using four months for this setting allowed all sites to have at least one window represented. Sites typically had one or a few years of data, and no site had data from all years.

Model weights were initialized using defaults for the GRU layer, while the output linear layer was initialized uniformly in [-0.1, 0.1] with zero bias. Training used one hundred training iterations, batch size of 10, Adam optimizer, and an adaptive learning rate beginning at 0.001 that is halved every ten epochs without improvement in validation loss, and default parameter initializations. We used a mean squared error loss, modified to ignore missing data and normalized by the number of non-missing time points when calculating the squared error for each training example.

CNN branch

A convolutional neural network (CNN) branch was implemented to analyze the MODIS image input ("CNN branch" in Figure 1). CNNs are a commonly used architecture for analyzing image data (LeCun et al., 2002), because they use a localized receptive field that scans over an image to detect spatially confined features (as a hypothetical example: a lake, or a stand of trees) and can capture spatial relationships between such

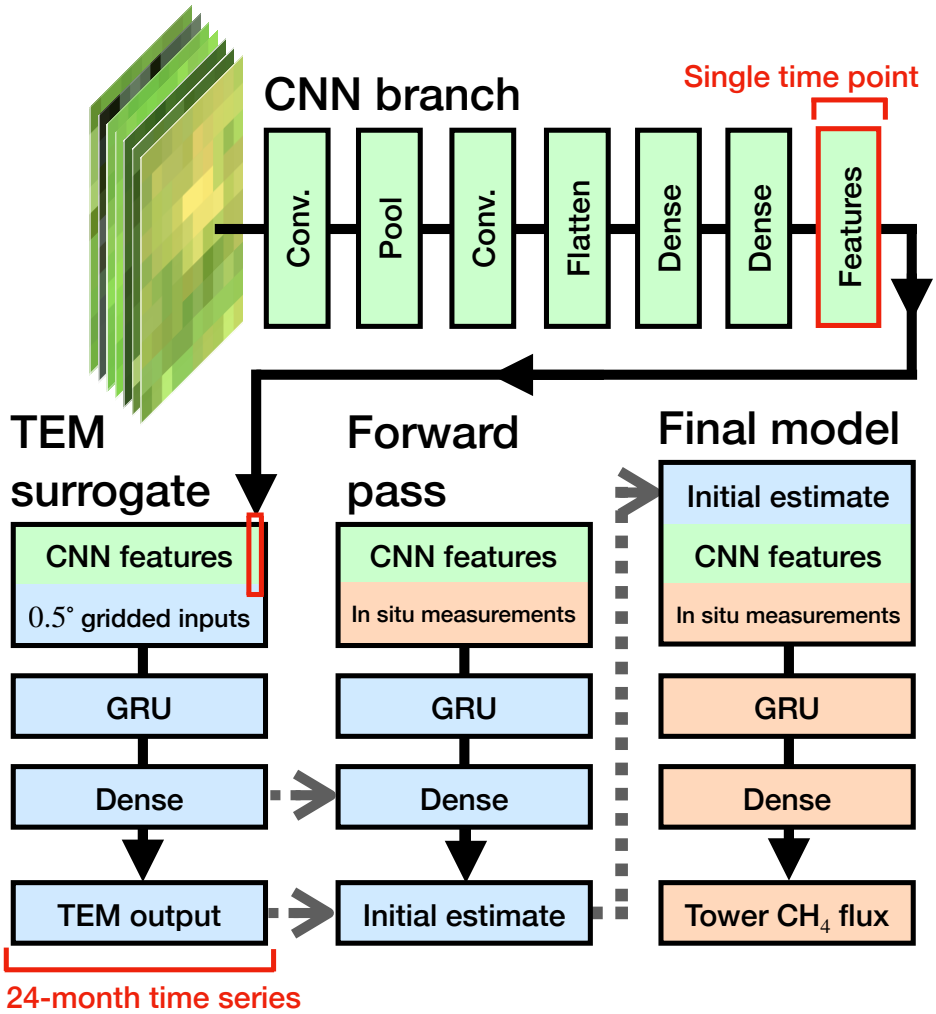

Figure 1: Neural network diagram for the best performing model: cross-domain model stacking with GRU + CNN branch. Boxes represent either model inputs or outputs from neural network layers. The top row conveys the CNN branch applied to a single monthly time point. Feature vectors output from the CNN are concatenated into a 24 month time series (angled black arrow) and combined with other gridded variables input to a TEM surrogate model. The TEM surrogate model is trained simultaneously with the CNN branch and uses a 24-month TEM-generated timeseries as the training target. Next, the pre-trained weights from the surrogate model used (small grey arrows) to obtain an initial estimate for each measurement site using in situ FLUXNET measurements as input. The initial estimates are used as input for training a second neural network (angled grey arrow), along with FLUXNET measurements and CNN features using local MODIS images, to obtain a final estimate for each site. Box colors track three different components: green layers or outputs are associated with the CNN branch; blue boxes are associated with 0.5° gridded data; and orange boxes are associated with local measurements specific to a flux tower.

features (e.g., a lake adjacent to the measurement site) (Zhu et al., 2017). In our case, a CNN was used to extract features from satellite images centered around the measurement site. Last, the information from MODIS is further condensed, as we tell the CNN branch to classify each image into one of three categories for the final output. The three categories represent abstract, machine-learned summary information, and the CNN branch is trained simultaneously with the GRU backbone of the model. Specifics of the CNN branch

are included in the Supplementary Material.

The following are specific parameters of the CNN architecture. Images from each time point were processed independently, applying a CNN consisting of a convolution (3×3 kernel; 64 neurons), ReLU, max pool (2×2), convolution (3×3 kernel; 128 neurons), ReLU, flatten channels, dense (256 neurons), ReLU, dense with three neurons, followed by a Gumbel-Softmax activation. The Gumbel-Softmax was used to produce a one-hot-like encoding in the forward pass while maintaining differentiability in backpropogation. This strategy was intended to alleviate overfitting.

Transfer learning

To combine information from the process-based model with our neural network we used a cross-domain model stacking approach (Figure 1). First, we trained a TEM surrogate model using the 0.5 degree TEM outputs as targets and 0.5 degree reanalysis data—temperature, latent heat flux, wetland type, MODIS reflectance—as input ("TEM surrogate" in Figure 1). To train the CNN branch we chose a random 5 km $\times$ 5 km tile within the grid cell, and generated new tiles each training iteration. Next, an initial methane estimate was obtained from the surrogate model for each FLUXNET site using local measurements as input ("Forward pass" in Figure 1). Last, we trained a second neural network using the initial estimate from the surrogate model, combined with the in situ measurements from the baseline model, as inputs ("Final model" in Figure 1). In other words, the final model uses the initial TEM estimate as an extra predictor variable. This approach conveys to the neural network what we expect methane flux should look like given what is known about the underlying biogeochemical processes. Other training parameters were the same as the basic GRU model, above. The surrogate model strategy takes advantage of fine-spatial resolution predictors, and avoids other inputs typically used by TEM that are only available with 0.5-degree resolution; the latter may blur the fine-resolution signal due to environmental heterogeneity within the grid cell. We also tried three additional transfer learning strategies that did not improve performance over the baseline (see Supplementary Material).

Model evaluation

Due to the small number of measurement sites, we used leave-one-out cross-validation: one site at a time was held out for testing, requiring independent training runs to evaluate performance on all sites. In addition, we repeated training 100 times per held-out site with random training and validation splits: 30% of sites were used for validation, i.e., for hyperparameter tuning. Repetitions were averaged to obtain an ensemble estimate for each site.

Models were evaluated and compared based on their mean annual total emissions estimate. For the 13 sites that each had at least one year of non-gap-filled (i.e., non-missing) data, we evaluated predictions from years without gap-filled data. The annual sum was calculated for each predicted 12-month time series, and the squared errors for the site were averaged across years. Finally, a model-wise RMSE in units of $gCH_4m^{-2}yr^{-1}$ was calculated that weights each site equally, irrespective of the number years of data.

A different procedure was used for sites with no years of complete data. We summed estimates from the m available months, and rescaled this value to obtain a yearly estimate, $S = \frac{12}{m}\sum_{i}^{m} x_i$. Squared errors were averaged across years. At this point, all sites were averaged together, including sites with complete years of data, to obtain a pseudo-RMSE, denoted RSME*. Therefore, the all-sites RMSE* may be used to glean at least a qualitative sense of performance on the full dataset, although the complete-sites RMSE is a more trustworthy metric for comparing models.

Bootstrapping was used to quantify model selection uncertainty. Specifically, predictions were sampled among measurement sites with replacement, using the same bootstrap resamples across models. For each iteration, RMSE was calculated using the precomputed site-level predictions and models were ranked by performance. The best-performing model from each iteration was recorded, yielding a frequency of winning for each model. Last, we calculated the difference in RMSE between pairs of models and obtained a confidence interval from the 2.5th and 97.5th percentiles of the paired bootstrap distribution. The latter serves as a hypothesis test: if the 95% confidence interval crosses zero, there is no significant difference between the models.

Global upscaling

To predict global emissions with machine learning, we trained an ensemble of 100 independent replicates of the best-performing model with random training-validation splits. After calculating a global sum for each replicate, the ensemble mean was used as the point estimate of global emission. Inputs included 0.5-degree gridded temperature and latent heat flux from ERA5 (Hersbach et al., 2023). To accommodate the CNN branch we predicted on every 5 km $\times$ 5 km tile in the grid cell and averaged the tiled predictions. Emissions estimates within each grid cell were multiplied by the estimated fraction of inundated area from the WAD2M $0.25^\circ \times 0.25^\circ$ product (Zhang et al., 2021). This inundation map was converted to $0.5^\circ \times 0.5^\circ$ by: multiplying the 0.25-resolution fractions by grid cell area to obtain the area inundated, averaging windows of 2×2 grid cells while ignoring missing data, multiplying by four, and dividing by the combined area. The first two years (2006 and 2007) of the study period were skipped during upscaling, because the initial year of each time series was used as spin-up, and we removed the first six months of input data in southern hemisphere

sites when shifting each time series to align seasonally with the northern hemisphere. In addition, we present results using alternative inundation fractions from SWAMPS-GLWD.

We quantified sampling uncertainty in the machine learning estimates using bootstrap resampling of the training sites. Sites were resampled with replacement for 100 replicates. Within each replicate, we trained an ensemble of three independent models using random training-validation splits and averaged the global emission estimate among the three runs. The 2.5th and 97.5th percentiles of the bootstrap distribution were used to define a 95% confidence interval.

A similar global-scale estimate was calculated from TEM output. Furthermore, we reused the aforementioned models to calculate a hybrid estimate: machine learning predictions were used for grid cells where each predictor variable fell within the training range, and TEM predictions were used for the remaining grid cells. To obtain a bootstrap distribution for the hybrid approach, the global estimate for each replicate used (i) predictions from a machine learning model trained on resampled sites and (ii) the point estimate from TEM.

Results

Baseline model performance

As a benchmark, we implemented a site-specific TEM run using local measurements as inputs when available and gridded reanalysis data for the remaining inputs. Evaluating TEM on 13 sites that had 12-month windows without missing data led to an RMSE of 84.9 $gCH_4m^{-2}yr^{-1}$ (Figure 2A; Table 1). TEM overestimated methane emissions at most of the evaluated sites, although the largest emitting site, Old Woman Creek (OWC), was predicted relatively well by TEM (Figure 2A). This overall poor performance suggests there are some parts of the biogeochemical process that are not modeled well by TEM. Alternatively, errors might accumulate due to mismatch in spatial scale between the gridded reanalysis data used for some inputs and the local tower measurements used for model evaluation.

Next, we trained a basic GRU on eddy covariance measurements, holding out one site at a time for testing (i.e., leave-one-out cross-validation; see Supplementary Material for model evaluation framework). The best performing GRU used three inputs: temperature, latent heat flux, and wetland type. We found it was not beneficial to include other variables measured at towers or reanalysis data products, even those expected to be informative for methane emissions; extra inputs hurt performance due to overfitting (however, see below discussion of reflectance data). Due to the small training set, training progress was heterogeneous between different training-validation splits, sometimes with training and validation losses both decreasing steadily

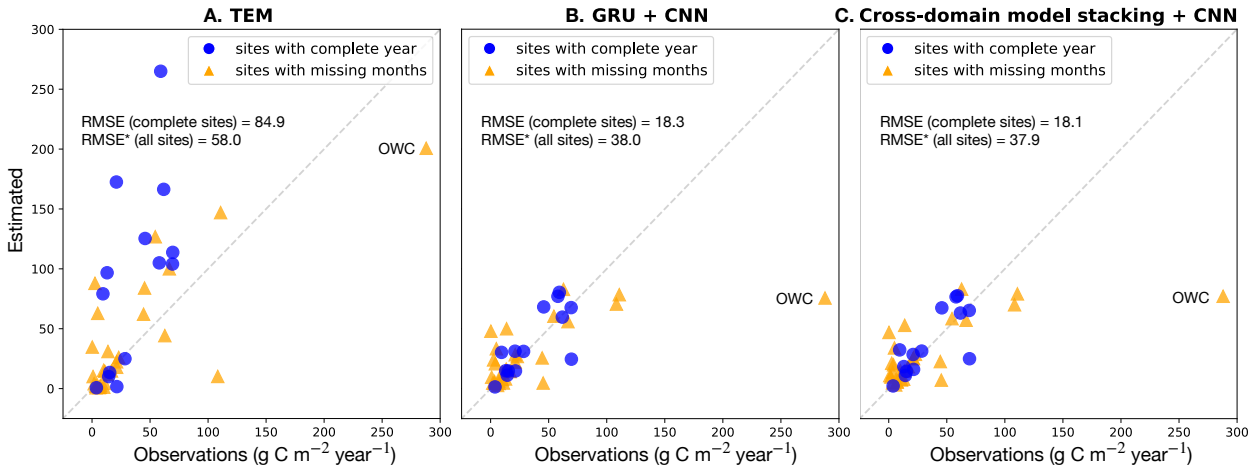

Figure 2: (A) Evaluation of TEM for predicting tower measurements, after training on local input data. (B) Machine learning model including CNN branch. (C) Machine learning model using initial TEM estimate as additional predictor. RMSE: root mean squared error ($gCH_4m^{-2}yr^{-1}$). RMSE*: psuedo-mean squared error, after extrapolating the site mean (including missing data for some sites) to 12-months.

Model	Model type	Spatial resolution	CNN	RMSE	RMSE*
TEM	process-based	0.5-degree		84.9	58.0
GRU	ML	tower observations		19.0	38.1
GRU, test on gridded inputs	ML	train with tower observations; test with 0.5-degree		24.0	47.8
pre-training	transfer learning	pre-train with 0.5-degree; fine-tune on tower observations		19.4	38.8
pre-train + domain adaptation	transfer learning	pre-train with 0.5-degree and unlabeled tower observations; fine-tune on tower observations		19.3	38.8
mixed training set	transfer learning	0.5-degree & tower observations		19.1	38.0
cross-domain model stacking	transfer learning	first model 0.5-degree; second model tower observations		18.6	38.1
GRU + CNN	ML	tower observations	Y	18.3	38.0
pre-training + CNN	transfer learning	pre-train with 0.5-degree; fine-tune on tower observations	Y	19.8	39.5
mixed training set + CNN	transfer learning	0.5-degree & tower observations	Y	18.4	38.3
model stacking + CNN	transfer learning	first model 0.5-degree; second model tower observations	Y	18.1	37.9

Table 1: Performance comparison between models for predicting fluxes measured at tower sites. The fourth column conveys whether the model includes a CNN branch. The root mean squared error (RMSE; in units of $gCH_4m^{-2}yr^{-1}$) was calculated using leave-one-out cross-validation on a subset of sites without missing data. RMSE* represents all sites, including those where missing months were interpolated.

for several training iterations, and other times with no apparent learning. The ensemble estimate from 100 training iterations of the GRU outperformed TEM with RMSE of 19.0 $gCH_4m^{-2}yr^{-1}$ (Table 1). The OWC

site was challenging to predict using machine learning, as the emissions from this site were larger than the other sites which were used to train the model. As a consequence, the estimates for OWC were visibly capped at the maximum emissions value from the other sites. Furthermore, OWC lacks measurements from winter months when emissions are expected to be smaller, potentially inflating the error for this site.

The initial comparison between TEM and the GRU was imperfect, because some of the inputs used for TEM are coarse spatial resolution data products, while the ground truth fluxes used for model evaluation were measured on-site. Consequently, we expect the prediction errors for TEM to be inflated. To contrive a more direct comparison, we trained the GRU as before, but tested the model using gridded reanalysis inputs. This experiment increased error for the GRU (RMSE=24.0 $gCH_4m^{-2}yr^{-1}$), but still outperformed TEM.

We found an additional input that improved accuracy: 5 km $\times$ 5 km MODIS reflectance images centered on each tower site. The images were processed using a CNN branch, and the extracted feature vector was merged with the above three inputs for the GRU. The GRU + CNN model reduced RMSE by 0.7 $gCH_4m^{-2}yr^{-1}$ relative to the baseline GRU (95% bootstrap CI: -0.361 to 1.962 $gCH_4m^{-2}yr^{-1}$) (Figure 2B). This improvement was not statistically significant, despite the GRU + CNN outperforming the GRU in a large proportion of bootstrap replicates (Table 2).

Model	Bootstrap support	Pairwise bootstrap
GRU	0.003	
GRU + CNN branch	0.196	$(-0.361, +1.962)$
cross-domain model stacking	0.204	$(-0.544, +1.159)$
model stacking + CNN branch	0.597	$(+0.033, +2.401)$

Table 2: Model-comparison uncertainty for four select models. The second column is the proportion of bootstrap replicates for which a model had the lowest error among the four models. The third column conveys a 95% bootstrap confidence interval for the magnitude of improvement in RMSE relative to the baseline GRU.

Transfer learning

We attempted to bridge the simulated and real domains using a model-stacking strategy: the site-specific output from the TEM surrogate model was included as an input feature alongside the other input variables. This cross-domain model stacking reduced RMSE relative to the baseline GRU by 0.4 $gCH_4m^{-2}yr^{-1}$ (95% bootstrap CI: -0.544 to 1.159 $gCH_4m^{-2}yr^{-1}$) (Table 1). Although the effect of model-stacking alone was not significant on its own, it outperformed the GRU in a large proportion of bootstrap replicates (Table 2). Combining model stacking with the CNN-branch led to a significantly better-performing model, reducing RMSE by 0.9 $gCH_4m^{-2}yr^{-1}$ relative to the baseline GRU (95% bootstrap CI: 0.033 to 2.401 $gCH_4m^{-2}yr^{-1}$) (Figure 2C), and outperforming other top models in the majority of bootstrap replicates (Table 2).

In addition, we explored three other transfer learning strategies—pre-training, domain-adaptation, and

training with a mixture of observed and simulated data. Unexpectedly, these techniques led the final model to under-perform relative to the baseline GRU, with or without the CNN branch (Table 1), presumably due to poor information from the process based model. Information on the alternative transfer learning strategies is described in the Supplementary Material.

In a followup experiment, we subsetted the 43 tower sites for seven sites with the best TEM estimates—these sites represent a best case scenario for transfer learning, as TEM is relatively good at predicting methane emissions at these sites. The subsetted sites shared the same wetland type, unforested bog. In this experiment, model stacking (without CNN) showed a more substantial effect, improving the RMSE by 1.84 $gCH_4m^{-2}yr^{-1}$ relative to the baseline GRU. This indicates that the overall weak transfer learning outcome using the full dataset is at least partly due to inaccurate outputs from the process-based model.

Global upscaling

We estimated global wetland emissions using three different approaches (Figure 3). First, we calculated the global sum directly from the TEM output, accounting for the inundated area within each grid cell. Using the process-based model and the WAD2M inundation map the mean annual emission was 171.2 $TgCH_4\,yr^{-1}$ globally (Figure 3). This value is larger than, but comparable to, a recent estimate using an ensemble of process-based models 158 ± 24 Tg (Zhang et al., 2025).

The global estimate from the GRU + CNN model was not far from that of TEM: 167.4 $TgCH_4\,yr^{-1}$ (95% bootstrap CI: 128.9 to 213.2 $TgCH_4\,yr^{-1}$) (Figure 3B). However, the relative emission intensity across geographic regions varied between the two methods. For example, at the continental scale TEM estimated larger emissions in North America, Europe, and Asia, while the machine learning model estimated higher emissions in South America and Africa (Figure 3C). Overall, the machine learning model predicted higher emissions than TEM for 60% of grid cells, however the difference in magnitude between the two methods was 32% greater for grid cells where TEM predicted larger emissions.

One limitation of machine learning is that models often struggle to generalize beyond the training distribution (Meyer and Pebesma, 2021). In particular, the temperature range that our GRU was exposed to during training was relatively moderate—e.g., maximum 29.6°C—and does not represent the full range of temperatures among global wetlands—e.g., up to 42.9°C (Figure 3D). Additional data points fell outside the training distribution for latent heat flux. The proportion of grid cells with at least one month of temperature or latent heat flux measurements outside the training range make up 23.1% of wetlands and are concentrated in the world's hottest regions. Machine learning models trained on currently available FLUXNET data will likely underestimate in such regions, as they are not used to seeing such high temperatures. For example,

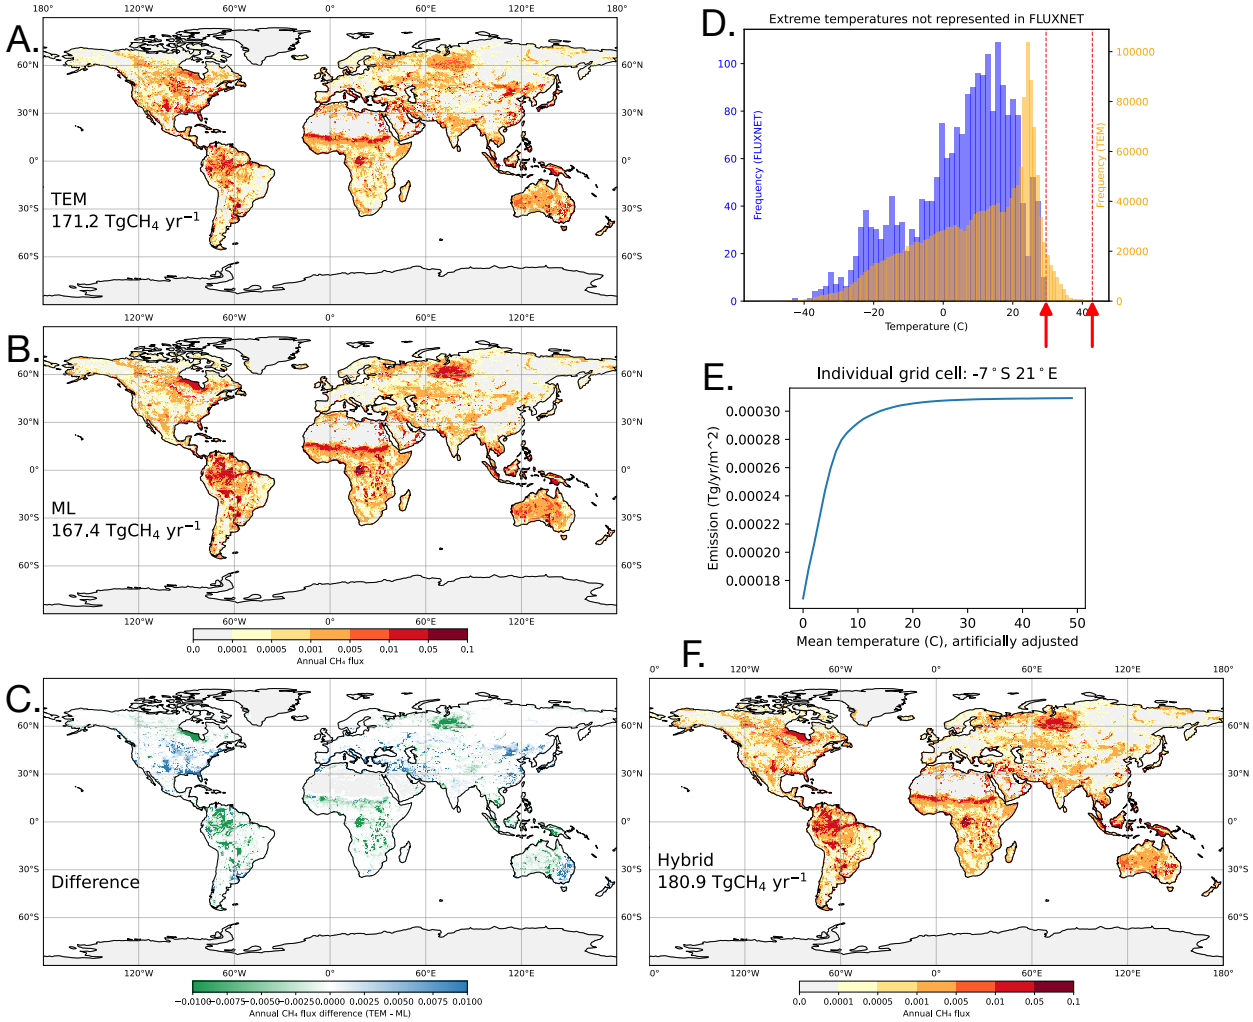

Figure 3: Panels A, B, C, and F show global emission predictions using WAD2M inundation fractions. (A) TEM emission intensity. (B) GRU with CNN branch and cross-domain model stacking. (C) Subtraction of the machine learning estimate from the TEM estimate. (D) Temperature distribution from FLUXNET combining all sites (blue), and global reanalysis inputs used by TEM (orange). Red arrows and dashed lines delineate the range of values in the global dataset that are larger than the maximum FLUXNET measurement. (E) Simulated analysis focusing on a randomly chosen site. A range of predictions were recorded from forward passes of the GRU, artificially re-centering the mean of the temperature inputs to values between 0°C and 50°C. (F) Hybrid approach implementing the machine learning model where inputs fall within the marginal ranges of predictors from training (76.9% of grid cells) and TEM otherwise (23.1%). Heatmap bins used arbitrary cutoffs to improve visualization.

in a simulation experiment we recorded outputs from a pre-trained GRU model across an array of input temperatures and observed a maximum-emission plateau (Figure 3E), assumed to be a machine learning artifact and not a biogeochemical effect.

To address this concern, we present a hybrid map: the machine learning model was implemented only on grid cells with inputs falling within the range of environmental conditions represented in the training data, and TEM was conservatively applied to the remaining grid cells (Figure 3F). Construction of the

hybrid map did not require additional training. Unexpectedly, the global emissions estimate using the hybrid approach was greater than either individual method: 180.9 $TgCH_4\,yr^{-1}$ (95% bootstrap CI: 141.6 to 204.8 $TgCH_4\,yr^{-1}$). This is due to the hybrid map retaining some of the largest emitting regions predicted from both the machine learning model—e.g., in the Hudson Bay Lowlands, Amazon, northern Siberia, and Congo Basin—and TEM—e.g., in the Southwestern United States and Western Asia. However, the WAD2M inundation map includes significant wetlands in the Southwestern United States and Australia where wetlands are unlikely.

Last, we implemented an alternative inundation map—GLWD-SWAMPS—that included 21% more wetland grid cells (Figure S1). However, GLWD-SWAMPS resulted in significant emissions over the Sahara desert, which is particularly suspect. Corresponding to the new inundation map, TEM estimated higher total emissions (203.8 $TgCH_4\,yr^{-1}$), while the new ML estimate was considerably lower (175.7 $TgCH_4\,yr^{-1}$; 95% bootstrap CI: 134.3 to 225.4). Under the new inundation map, the hybrid approach led to an intermediate global estimate (187.2 $TgCH_4\,yr^{-1}$; 95% bootstrap CI: 154.5 to 227.4). Together, these results highlight the need for reliable inundation maps for studying natural methane emissions.

Comparison with previous studies

Next, we compared our findings with previous machine learning-based global wetland models from McNicol et al. (2023) and Chen et al. (2024) that also used the WAD2M inundation map. McNicol et al. used random forests to model the same FLUXNET sites analyzed in the current study. Likewise, we examined the output from Chen et al. who analyzed 47 FLUXNET sites and 35 chamber sites using random forests.

Compared to McNicol et al. (151.1 TgCH4 yr-1), both Chen et al. (178.2 TgCH4 yr-1) and our hybrid map (180.9 TgCH4 yr-1) produced higher global mean wetland emissions (Figure 4, Table 3). One difference is that McNicol et al. used WorldClim meteorology (Fick and Hijmans, 2017), whereas Chen et al. and our study used ERA meteorology (Hersbach et al., 2023). Regionally, our machine learning and hybrid maps showed larger emissions in tropical forests and temperate coniferous forests but smaller emissions in woodlands and shrublands relative to McNicol et al. and Chen et al. (Table 3). These differences likely arise from the incorporation of process information from TEM within our analysis.

While the process-based TEM model showed an expected positive trend in global wetland emissions over 2008-2018, the machine learning-based outputs were relatively flat (Figure 4). This leveling-out may have occurred because the machine learning model could not extrapolate beyond the training distribution (Figure 3). Furthermore, our machine learning analysis showed a slight decreasing trend in mean annual emissions. This is partly due to the model's dependence on the latent heat flux predictor variable, which

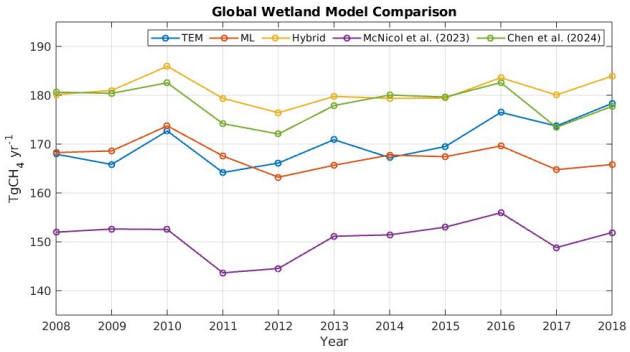

Figure 4: Comparing mean annual emission estimates across studies.

Vegetation Type	TEM	ML	Hybrid	McNicol et al. (2023)	Chen et al. (2024)
Alpine Tundra	0.9	1.5	1.5	0.9	1.3
Wet Tundra	2.0	7.0	6.9	3.9	6.2
Boreal Forest	11.4	19.7	19.8	14.0	22.2
Temp. Coniferous Forest	21.6	5.0	8.2	4.3	4.2
Temp. Deciduous Forest	6.4	2.8	3.5	2.5	2.7
Grasslands	54.5	46.4	50.8	43.5	54.8
Shrublands	16.3	22.9	18.9	34.7	31.7
Tropical Forest	40.4	55.3	59.3	33.4	40.3
Woodlands	1.7	4.3	3.4	11.4	13.1
Temp. Evergreen Forest	16.0	2.4	8.6	2.3	2.0
Sum	171.2	167.4	180.9	150.9	178.4

Table 3: Mean annual emissions (Tg CH_4 yr^{-1}) by vegetation type.

has significant interannual fluctuations and a slight downward trend during the study period (Figure S2). However, our hybrid map showed a more even trend, reflecting the use of TEM's information. The hybrid approach also exhibited stronger seasonality relative to both McNicol et al. and Chen et al. (Figure 5). In conclusion, integrating process-based information with machine learning-based upscaling improved consistency with established trend and seasonal behaviors while maintaining or enhancing agreement with independent machine learning-based global estimates.

Discussion

We set out to advance the estimation of methane emissions from natural wetlands by combining machine learning with a process-based model. This was motivated by the expectation of more accurate flux estimates, in particular by supplementing the machine learning training set which currently includes a limited number of geographic locations with *in situ* flux measurements. One technique in particular, cross-domain model stacking, incrementally improved performance and is the most computationally lightweight transfer learning method we examined (Table 1). However, the 2.1% improvement in accuracy with this technique does

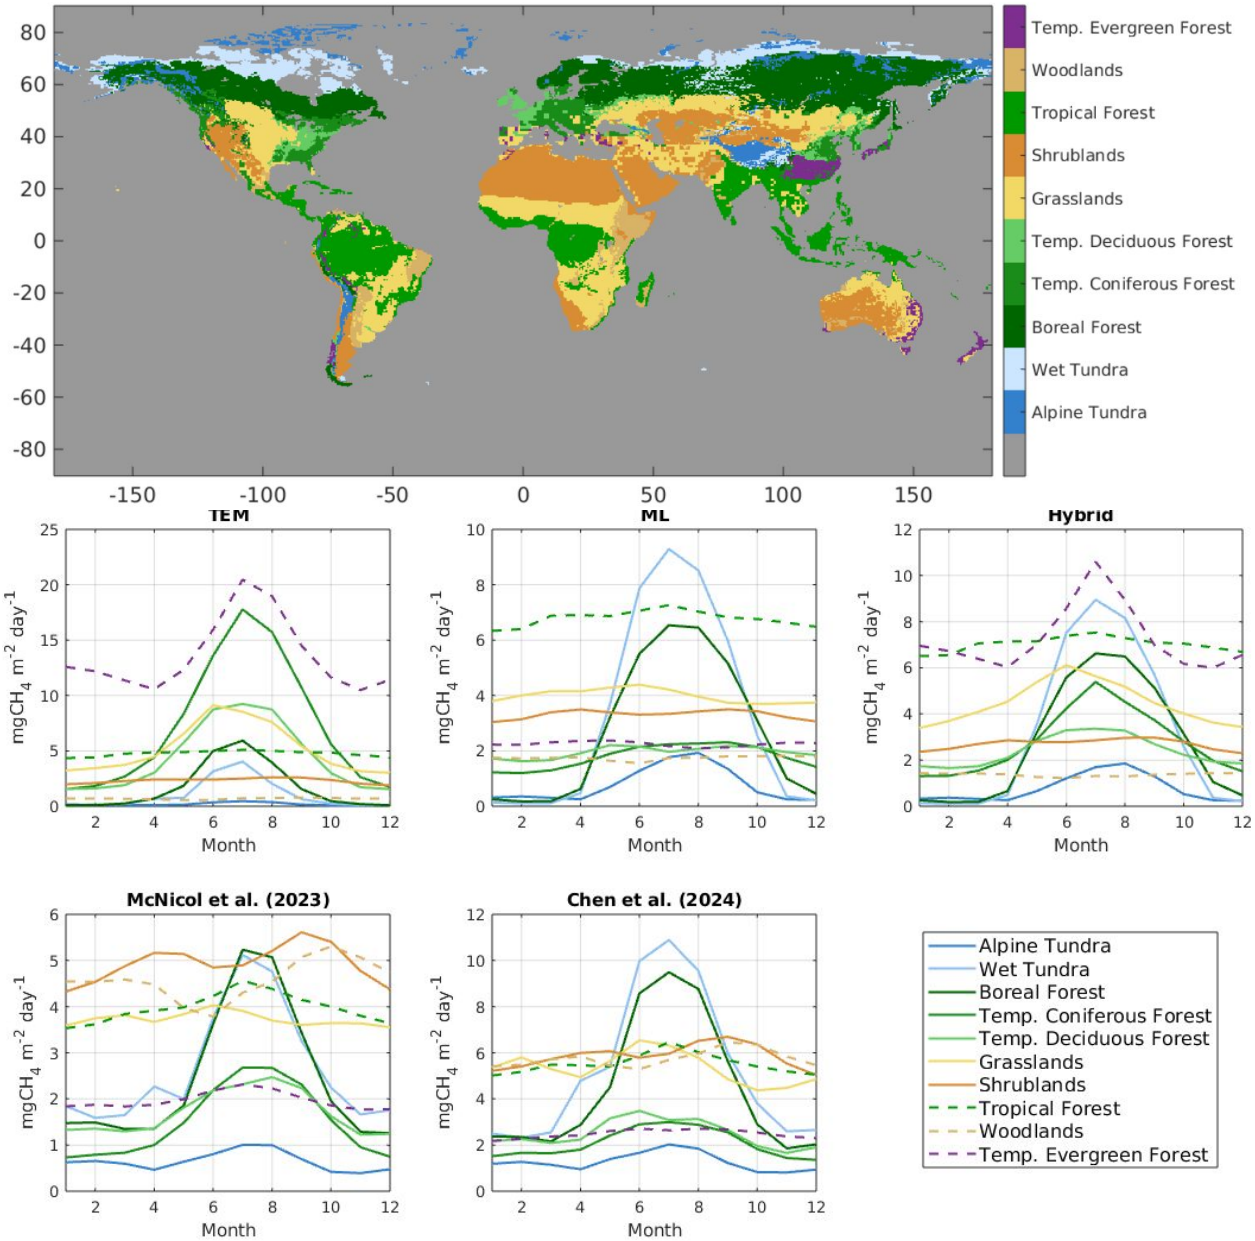

Figure 5: (Top) Global vegetation type map. (Bottom) Visualizing relative emissions across months (mean daily units), with curves colored by vegetation type, and separate panels for each model.

not represent a full solution to the challenging problem of estimating wetland emissions. Other examined transfer learning experiments did not improve performance (Table 1); however, we observed a larger effect using transfer learning for sites with more accurate TEM estimates, which provides hope for using transfer learning with improved process-based models in the future. We note that it is challenging to evaluate TEM on eddy covariance measurements with relatively small flux footprints when TEM uses coarse, gridded data as input, especially when only a small fraction of the grid cell is wetland. The error reported for TEM is

overestimated due to this mismatch. An additional limitation of our cross-validation framework is spatial autocorrelation between sites. Although the sites are distributed globally, some are co-located; for example, CA-SCB and CA-SCC are separated by less than 1 km. Consequently, nearby sites appear in different cross-validation folds, likely causing RMSE to be underestimated.

Our machine learning framework differed qualitatively from those used in recent studies, but performed similarly. The recurrent neural network architecture serves to *automatically* remember past events, for example, temperatures and latent heat fluxes from the previous month that may lag behind a corresponding change in methane flux. Similar architectures have been used for carbon dioxide flux estimation (Liu et al., 2024). In contrast, McNicol et al. (2023) address temporal lag by manually shifting the timing of some predictors; however Chen et al. (2024) included recurrent neural networks in their ensemble analysis. Excluding the OWC site to facilitate comparison with McNicol et al. (2023), whose model achieved an R^2 of 0.54, our best model achieved an R^2 of 0.580; however, unlike McNicol et al. (2023), we did not take steps to account for spatial autocorrelation. Chen et al. (2024) analyzed additional eddy covariance sites (47 in total) and 36 chamber sites, and reported R^2 values ranging from 0.54 to 0.7 depending on the choice of model.

We also used different predictors than the aforementioned studies. While most predictors we examined led to overfitting and impaired performance, the addition of satellite imagery from MODIS marginally improved prediction accuracy. Previous methane studies have used MODIS-derived data products, e.g., land cover type and vegetation indices (Malone et al., 2021; McNicol et al., 2023), however we use a computer vision approach to extract useful information from raw, two-dimensional satellite images in an automated fashion. A key step in implementing image data was constraining the output from the CNN branch to a small number of features to alleviate overfitting.

To estimate global-scale emissions, we used conditional model selection, applying the best performing neural network only to grid cells where the air temperature and other model inputs were moderate. As neural networks use large numbers of parameters, they often struggle to generalize beyond the training distribution (Tobin et al., 2017; Gulrajani and Lopez-Paz, 2020). By avoiding pushing the neural network to generalize beyond the range of inputs it is used to seeing, we can be more confident about its output. Meanwhile, TEM, representing a conservative option in this context, was applied to sites with extreme temperatures. Surprisingly, the resulting global emissions estimate was larger than either individual method, and larger than previous studies (McNicol et al., 2023). Such hybrid modeling may be required for accurate, global upscaling, especially considering the limited number of flux tower sites. Future studies may obtain even more trustworthy machine learning predictions by analyzing the similarity between grid cells and the training data in multivariate space (Meyer and Pebesma, 2021). However, this approach may further limit the number of regions where machine learning is applied, due to the small number of measurement sites representative

of natural wetlands. Our hybrid-modeling approach fits naturally for Earth-systems research that involves process-based models, because process-based models can often extrapolate reliably, allowing prediction at the global scale.

We intend for our findings—including useful strategies, and ones that didn't work—to guide future modeling efforts in wetland methane modeling. Our study also highlights the need for refined process-based models that more accurately reflect heterogeneous global environments. As outlined in Malone et al. (2021), additional measurement sites will also improve our modeling capabilities using data-driven approaches like neural networks.

Conflict of Interest Disclosure

The authors declare there are no conflicts of interest for this manuscript.

Acknowledgments

This work was supported by DOE Environmental System Science (DE-SC0024360).

Data availability

Code for reproducing the full analysis is available from Zenodo (Smith, 2026).

References

Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, et al. Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 929–947, 2024.

Katharina Beckh, Sebastian Müller, Matthias Jakobs, Vanessa Toborek, Hanxiao Tan, Raphael Fischer, Pascal Welke, Sebastian Houben, and Laura von Rueden. Explainable machine learning with prior knowledge: an overview. *arXiv preprint arXiv:2105.10172*, 2021.

Shuo Chen, Licheng Liu, Yuchi Ma, Qianlai Zhuang, and Narasinha J Shurpali. Quantifying global wetland

methane emissions with in situ methane flux data and machine learning approaches. *Earth's Future*, 12 (11):e2023EF004330, 2024.

Zijun Cui, Tian Gao, Kartik Talamadupula, and Qiang Ji. Knowledge-augmented deep learning and its applications: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.

Dick P Dee, SM Uppala, Adrian J Simmons, Paul Berrisford, Paul Poli, Shinya Kobayashi, U Andrae, MA Balmaseda, G Balsamo, d P Bauer, et al. The era-interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the royal meteorological society*, 137(656):553–597, 2011.

Kyle B Delwiche, Sara Helen Knox, Avni Malhotra, Etienne Fluet-Chouinard, Gavin McNicol, Sarah Feron, Zutao Ouyang, Dario Papale, Carlo Trotta, Eleonora Canfora, et al. Fluxnet-ch4: A global, multi-ecosystem dataset and analysis of methane seasonality from freshwater wetlands. *Earth System Science Data Discussions*, 2021:1–111, 2021.

Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 2017. doi: 10.1016/j.rse.2017.06.031. URL `https://doi.org/10.1016/j.rse.2017.06.031`.

Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.

H Hersbach, B Bell, P Berrisford, G Biavati, A Horányi, J Muñoz Sabater, J Nicolas, C Peubey, R Radu, I Rozum, et al. Era5 monthly averaged data on pressure levels from 1940 to present, copernicus climate change service (c3s) climate data store (cds)[data set], 2023.

Katherine Jensen and Kyle Mcdonald. Surface water microwave product series version 3: A near-real time and 25-year historical global inundated area fraction time series from active and passive microwave remote sensing. *IEEE Geoscience and Remote Sensing Letters*, 16(9):1402–1406, 2019.

Anuj Karpatne, Xiaowei Jia, and Vipin Kumar. Knowledge-guided machine learning: Current trends and future prospects. *arXiv preprint arXiv:2403.15989*, 2024.

Sara H Knox, Sheel Bansal, Gavin McNicol, Karina Schafer, Cove Sturtevant, Masahito Ueyama, Alex C Valach, Dennis Baldocchi, Kyle Delwiche, Ankur R Desai, et al. Identifying dominant environmental predictors of freshwater wetland methane fluxes across diurnal to seasonal time scales. *Global Change Biology*, 27(15):3582–3604, 2021.

McKenzie Kuhn, David Olefeldt, Kyle A Arndt, David Bastviken, Lori Bruhwiler, Patrick Crill, Tonya DelSontro, Etienne Fluet-Chouinard, Guido Grosse, Mikael Hovemyr, et al. Current and future methane emissions from boreal-arctic wetlands and lakes. *Nature Climate Change*, pages 1–6, 2025.

X. Lan, K.W. Thoning, and E.J. Dlugokencky. Trends in globally-averaged ch4, n2o, and sf6 determined from noaa global monitoring laboratory measurements. Version 2025-11, `https://doi.org/10.15138/P8XG-AA10`, 2022. Accessed 28 November 2025.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 2002.

Bernhard Lehner and Petra Döll. Development and validation of a global database of lakes, reservoirs and wetlands. *Journal of hydrology*, 296(1-4):1–22, 2004.

Xin Lin, Shushi Peng, Philippe Ciais, Didier Hauglustaine, Xin Lan, Gang Liu, Michel Ramonet, Yi Xi, Yi Yin, Zhen Zhang, et al. Recent methane surges reveal heightened emissions from tropical inundated areas. *Nature Communications*, 15(1):10894, 2024.

Licheng Liu, Qianlai Zhuang, Youmi Oh, Narasinha J Shurpali, Seungbum Kim, and Ben Poulter. Uncertainty quantification of global net methane emissions from terrestrial ecosystems using a mechanistically based biogeochemistry model. *Journal of Geophysical Research: Biogeosciences*, 125(6):e2019JG005428, 2020.

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.

Sparkle Malone, Youmi Oh, Kyle Arndt, George Burba, Roisin Commane, Alexandra Contosta, Jordan Goodrich, Henry Loescher, Gregory Starr, and Ruth Varner. Gaps in network infrastructure limit our understanding of biogenic methane emissions in the united states. *Biogeosciences Discussions*, 2021:1–24, 2021.

Gavin McNicol, Etienne Fluet-Chouinard, Zutao Ouyang, Sara Knox, Zhen Zhang, Tuula Aalto, Sheel Bansal, Kuang-Yu Chang, Min Chen, Kyle Delwiche, et al. Upscaling wetland methane emissions from the fluxnet-ch4 eddy covariance network (upch4 v1. 0): Model development, network assessment, and budget comparison. *AGU Advances*, 4(5):e2023AV000956, 2023.

Hanna Meyer and Edzer Pebesma. Predicting into unknown space? estimating the area of applicability of spatial prediction models. *Methods in Ecology and Evolution*, 12(9):1620–1633, 2021.

Youmi Oh, Qianlai Zhuang, Lisa R Welp, Licheng Liu, Xin Lan, Sourish Basu, Edward J Dlugokencky, Lori Bruhwiler, John B Miller, Sylvia E Michel, et al. Improved global wetland carbon isotopic signatures support post-2006 microbial methane emission increase. *Communications Earth & Environment*, 3(1):159, 2022.

Youmi Oh, Lori Bruhwiler, Xin Lan, Santanu Halder, Ben Riddell-Young, Sourish Basu, John B Miller, Sylvia Michel, Ken Schuldt, Arlyn Andrews, et al. Isotopic constraints in methane inversions reveal larger trends in wetland emissions with improved linkage to terrestrial water storage. *Nature Communications*, 2026.

Zhen Qu, Daniel J Jacob, A Anthony Bloom, John R Worden, Robert J Parker, and Hartmut Boesch. Inverse modeling of 2010–2022 satellite observations shows that inundation of the wet tropics drove the 2020–2022 methane surge. *Proceedings of the National Academy of Sciences*, 121(40):e2402730121, 2024.

Marielle Saunois, Adrien Martinez, Benjamin Poulter, Zhen Zhang, Peter A Raymond, Pierre Regnier, Josep G Canadell, Robert B Jackson, Prabir K Patra, Philippe Bousquet, et al. Global methane budget 2000–2020. *Earth System Science Data*, 17(5):1873–1958, 2025.

Chris Smith. Kgml-ch4/wetland_methane_ml: preprint, 2026. URL `https://doi.org/10.5281/zenodo.20403861`. [Software].

Yiming Sun, Shuo Chen, Shengyu Chen, Chonghao Qiu, Licheng Liu, Youmi Oh, Sparkle L Malone, Gavin McNicol, Qianlai Zhuang, Chris Smith, et al. X-methanewet: A cross-scale global wetland methane emission benchmark dataset for advancing science discovery with ai. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pages 2782–2793, 2026.

Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017.

Claire C Treat, Anna-Maria Virkkala, Eleanor Burke, Lori Bruhwiler, Abhishek Chatterjee, Joshua B Fisher, Josh Hashemi, Frans-Jan W Parmentier, Brendan M Rogers, Sebastian Westermann, et al. Permafrost carbon: Progress on understanding stocks and fluxes across northern terrestrial ecosystems. *Journal of Geophysical Research: Biogeosciences*, 129(3):e2023JG007638, 2024.

E. Vermote. MODIS/Terra Surface Reflectance 8-Day L3 Global 500m SIN Grid V061 [Data set]. `https://doi.org/10.5067/MODIS/MOD09A1.061`, 2021. Accessed: 2025-08-06.

Laura Von Rueden, Sebastian Mayer, Katharina Beckh, Bogdan Georgiev, Sven Giesselbach, Raoul Heese, Birgit Kirsch, Julius Pfrommer, Annika Pick, Rajkumar Ramamurthy, et al. Informed machine learning–a taxonomy and survey of integrating prior knowledge into learning systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):614–633, 2021.

Xiaofeng Xu, Fengming Yuan, Paul J Hanson, Stan D Wullschleger, Peter E Thornton, William J Riley, Xia Song, David E Graham, Changchun Song, and Hanqin Tian. Reviews and syntheses: Four decades of modeling methane cycling in terrestrial ecosystems. *Biogeosciences*, 13(12):3735–3755, 2016.

Zhen Zhang, Etienne Fluet-Chouinard, Katherine Jensen, Kyle McDonald, Gustaf Hugelius, Thomas Gumbricht, Mark Carroll, Catherine Prigent, Annett Bartsch, and Benjamin Poulter. Development of the global dataset of wetland area and dynamics for methane modeling (wad2m). *Earth System Science Data*, 13(5):2001–2023, 2021.

Zhen Zhang, Benjamin Poulter, Joe R Melton, William J Riley, George H Allen, David J Beerling, Philippe Bousquet, Josep G Canadell, Etienne Fluet-Chouinard, Philippe Ciais, et al. Ensemble estimates of global wetland methane emissions over 2000–2020. *Biogeosciences*, 22(1):305–321, 2025.

Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE geoscience and remote sensing magazine*, 5(4):8–36, 2017.

Qianlai Zhuang, Jerry M Melillo, David W Kicklighter, Ronald G Prinn, A David McGuire, Paul A Steudler, Benjamin S Felzer, and Shaomin Hu. Methane fluxes between terrestrial ecosystems and the atmosphere at northern high latitudes during the past century: A retrospective analysis with a process-based biogeochemistry model. *Global Biogeochemical Cycles*, 18(3), 2004.

Navid Zobeiry and Anoush Poursartip. Theory-guided machine learning for process simulation of advanced composites. *arXiv preprint arXiv:2103.16010*, 2021.

Supplementary Material

Additional transfer learning implementations

Pre-training. As an initial transfer learning strategy, we pre-trained the GRU on TEM outputs and reanalysis data as inputs. Here, we pre-trained a GRU on abundant TEM outputs, plus corresponding gridded reanalysis data for temperature, latent heat flux, and wetland type as inputs. Observed data were normalized independently from the TEM data. Pre-training used the same parameters as before, except with a batch size of 1,000. After pre-training on TEM data, we fine-tuned the network—that is, without re-initializing its weights but only resetting the learning rate to 0.001—by performing additional training on observed tower measurements, using the same settings as above. Unexpectedly, this procedure led the final model to under-perform relative to the baseline GRU without transfer learning (Table 1). Similarly, pre-training hurt performance when using the CNN branch.

Domain adaptation. Next, we tried a more sophisticated pre-training procedure involving domain adaptation. Specifically, we showed the neural network examples of tower measurements ("target domain") during pre-training, and attached a domain classifier branch to the network that implements a gradient reversal layer. The basic GRU model was separated into a feature extractor and output branch, and the domain classifier included a gradient reversal layer. This strategy is designed to extract features from inputs that are shared between the two domains, i.e., between TEM and flux tower measurements. The observed data were normalized using the TEM mean and standard deviation for pre-training. We pre-trained the domain adaptation model on a mixture of observed and TEM data, feeding three training examples of each to the domain classifier branch each training iteration. Next, the in situ measurements were re-normalized independently from TEM, and the model was fine-tuned on the observed data as before. However, this strategy hurt performance relative to the baseline, as with the simpler pre-training run. We skipped trying domain adaptation with an added CNN branch because domain adaptation did not help the basic GRU.

Mixed training set. Next, we tried a separate transfer learning approach: supplementing the observed training set with TEM outputs and the corresponding gridded reanalysis inputs. Due to the low proportion of real observations we normalized the combined dataset using a weighted Z-normalization with one part observed and one part TEM using: $\mu = \frac{\sum w_i x_i}{\sum w_i}$ and $\sigma = \sqrt{\frac{\sum w_i (x_i - \mu)^2}{\sum w_i}}$. Thus, the new training set is a mixture of: (i) limited flux tower observations, and (ii) abundant "simulated" data from TEM. Likewise, we used weighted mean squared error loss with one part observed and one part TEM; thus, the observed and simulated examples are weighted equally in the loss function despite the training set being dominated by the simulated domain. Each training batch consisted of five observed examples and five randomly chosen TEM examples, until all observed training data were covered; as a result, each training iteration sees a different

subset of TEM data. This approach also did not improve performance, with or without the CNN branch.

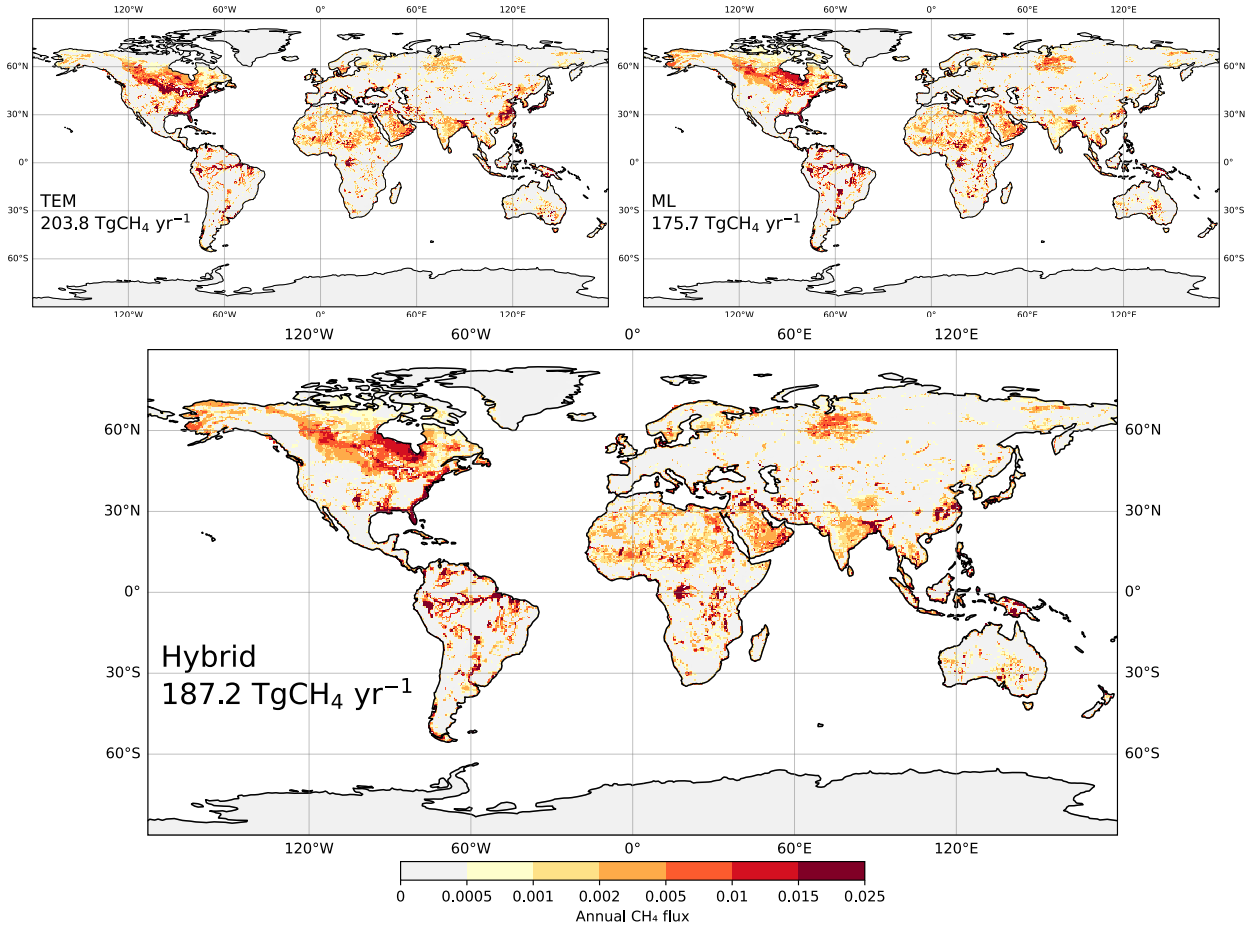

Figure S1: Upscaling with SWAMPS-GLWD inundation map.

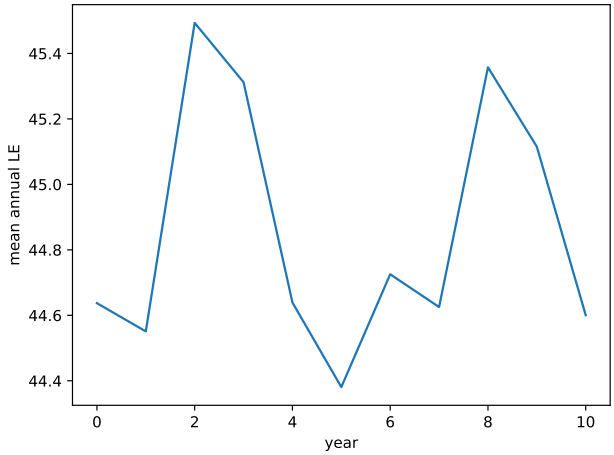

Figure S2: Trend in latent heat flux (LE; Wm^{-2}) over the full study period.

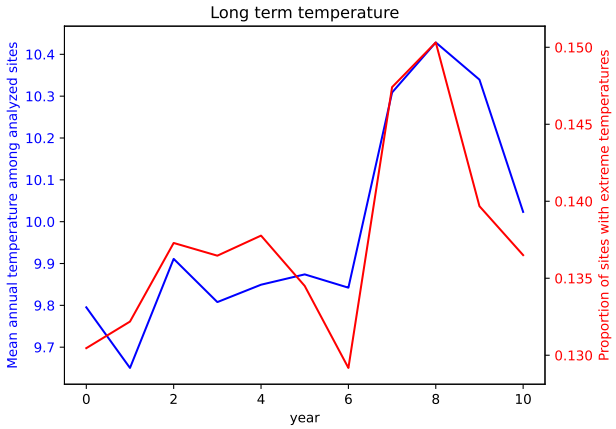

Figure S3: Trend in mean annual temperature (°C) over the full study period.

Site Identifier	Wetland type
CA-SCB	Unforested bog
CA-SCC	Forested bog
DE-Hte	Unforested bog
DE-Zrk	Unforested bog
DE-SfN	Unforested bog
FI-Lom	Unforested bog
FI-Si2	Unforested bog
FI-Sii	Unforested bog
FR-LGt	Unforested bog
JP-BBY	Unforested bog
NZ-Kop	Unforested bog
RU-Ch2	Unforested bog
RU-Cok	Unforested bog
SE-Deg	Unforested bog
US-A03	Unforested bog
US-ICs	Unforested bog
US-A10	Unforested bog
US-Atq	Unforested bog
US-Beo	Unforested bog
US-Bes	Unforested bog
US-NGB	Unforested bog
US-BZB	Forested bog
US-BZF	Unforested bog
US-Uaf	Forested bog
US-Ivo	Unforested bog
US-Los	Unforested bog
US-NGC	Unforested bog
BR-Npw	Forested swamp
BW-Gum	Unforested swamp
BW-Nxr	Unforested swamp
ID-Pag	Forested swamp
MY-MLM	Forested swamp
US-NC4	Forested swamp
US-DPW	Unforested swamp
US-LA2	Unforested swamp
US-Myb	Unforested swamp
US-ORv	Forested swamp
US-OWC	Unforested swamp
US-Sne	Unforested swamp
US-Tw1	Unforested swamp
US-Tw4	Unforested swamp
US-Tw5	Unforested swamp
US-WPT	Unforested swamp

Table S1: Wetland classifications for the analyzed sites.