



**Titre:** Multimodal Deep Learning for Carbon Flux Modelling  
Title:

**Auteur:** Matthew Fortier  
Author:

**Date:** 2024

**Type:** Mémoire ou thèse / Dissertation or Thesis

**Référence:** Fortier, M. (2024). Multimodal Deep Learning for Carbon Flux Modelling [Mémoire de maîtrise, Polytechnique Montréal]. PolyPublie.  
Citation: <https://publications.polymtl.ca/59450/>

 **Document en libre accès dans PolyPublie**  
Open Access document in PolyPublie

**URL de PolyPublie:** <https://publications.polymtl.ca/59450/>  
PolyPublie URL:

**Directeurs de recherche:** Christopher J. Pal, & Olivier Sonnetag  
Advisors:

**Programme:** Génie informatique  
Program:

**POLYTECHNIQUE MONTRÉAL**

affiliée à l'Université de Montréal

**Multimodal Deep Learning for Carbon Flux Modelling**

**MATTHEW FORTIER**

Département de génie informatique et génie logiciel

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*

Génie informatique

Septembre 2024

**POLYTECHNIQUE MONTRÉAL**

affiliée à l'Université de Montréal

Ce mémoire intitulé :

**Multimodal Deep Learning for Carbon Flux Modelling**

présenté par **Matthew FORTIER**

en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*

a été dûment accepté par le jury d'examen constitué de :

**Sarath CHANDAR ANBIL PARTHIPAN**, président

**Christopher J. PAL**, membre et directeur de recherche

**Oliver SONNENTAG**, membre et codirecteur de recherche

**Markus REICHSTEIN**, membre externe

**DEDICATION**

*To my loving wife Katherine,  
whose boundless support and encouragement enables me to be the best version of myself,  
even when it means uprooting our lives to follow a dream.*

*This world may be cold and indifferent,  
but there is still time to leave it better than we found it.*

## ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my research supervisors, Chris Pal and Oliver Sonnentag. Their unwavering support and guidance through innumerable projects allowed me to find and pursue a research path I am truly passionate about.

I am also immensely grateful to Joe Melton, a long-time mentor whose insight and encouragement were invaluable for not just this project, but for finding my way through academia and my early career.

The staff and students at Mila Quebec have provided an incredible community that made my research journey enriching and collaborative. In particular I would like to thank Mats L. Richter - his technical expertise helped me get my research off the ground on more than one occasion.

I extend my sincere thanks to the IVADO AI, Biodiversity, and Climate Change initiative, whose funding made this research possible.

Thank you also to Compute Canada for providing the computational resources necessary for this research.

Lastly, I am grateful to Polytechnique Montréal for providing the academic scaffolding upon which I could build my work. Votre personnel, votre infrastructure et vos ressources ont joué un rôle fondamental dans mon développement académique et dans la réalisation de cette thèse.

To all those mentioned, and to everyone who has supported me along the way, I offer my heartfelt thanks.

## RÉSUMÉ

Les flux de carbone terrestre fournissent des informations vitales sur la santé de notre biosphère et sur sa capacité à absorber les émissions anthropiques de  $\text{CO}_2$ . L'importance de la prévision des flux de carbone a conduit à l'émergence du domaine de la modélisation des flux de carbone guidée par les données (DDCFM), qui utilise des techniques statistiques pour prévoir les flux de carbone à partir de données biophysiques. La DDCFM permet aux chercheurs de produire des cartes globales de l'activité photosynthétique et respiratoire, mais elle est freinée par des modèles d'apprentissage automatique dépassés qui ne peuvent pas tirer parti de la multimodalité des données biophysiques.

Dans cette thèse, nous présentons deux avancées significatives pour la DDCFM : CarbonSense, un ensemble complet de données multimodales, et EcoPerceiver, une architecture neuronale multimodale profonde. Nous menons des expériences approfondies comparant EcoPerceiver à un modèle de référence de pointe utilisant CarbonSense comme ensemble d'entraînement. Nos résultats démontrent qu'EcoPerceiver est toujours plus performant que le modèle de référence dans une large gamme de types d'écosystèmes, soulignant le potentiel de l'apprentissage multimodal profond pour faire progresser la DDCFM. En publiant à la fois l'ensemble de données et le modèle, nous visons à fournir une base solide pour la recherche future et à encourager l'innovation dans ce domaine.

## ABSTRACT

Terrestrial carbon fluxes provide vital information about our biosphere’s health and its capacity to absorb anthropogenic CO<sub>2</sub> emissions. The importance of predicting carbon fluxes has led to the emerging field of data-driven carbon flux modelling (DDCFM), which uses statistical techniques to predict carbon fluxes from biophysical data such as meteorological and soil variables. DDCFM allows researchers to produce global maps of photosynthetic and respiratory activity, but it is held back by the use of dated machine learning techniques such as decision tree models which cannot capitalize on the multimodality of the biophysical data.

In this thesis, we present two significant advancements for DDCFM: CarbonSense, a comprehensive multimodal dataset, and EcoPerceiver, a deep multimodal neural architecture. We conduct extensive experiments comparing EcoPerceiver to a baseline state-of-the-art model using CarbonSense as a training set. Our results demonstrate that EcoPerceiver consistently outperforms the baseline across a wide range of ecosystem types, highlighting the potential of multimodal deep learning to drive advances in DDCFM. By releasing both the dataset and the model, we aim to provide a robust foundation for future research and encourage further innovation in this field.

## TABLE OF CONTENTS

DEDICATION . . . . .	iii
ACKNOWLEDGEMENTS . . . . .	iv
RÉSUMÉ . . . . .	v
ABSTRACT . . . . .	vi
TABLE OF CONTENTS . . . . .	vii
LIST OF TABLES . . . . .	ix
LIST OF FIGURES . . . . .	x
LIST OF SYMBOLS AND ACRONYMS . . . . .	xii
LIST OF APPENDICES . . . . .	xiii
CHAPTER 1 INTRODUCTION . . . . .	1
1.1 Motivation . . . . .	1
1.2 Research Objectives . . . . .	2
1.3 Thesis Outline . . . . .	2
CHAPTER 2 LITERATURE REVIEW . . . . .	4
2.1 Data-Driven Carbon Flux Modelling . . . . .	4
2.1.1 Measuring Carbon Fluxes: The Eddy Covariance Method . . . . .	4
2.1.2 Common Flux Predictors . . . . .	5
2.1.3 Previous Works . . . . .	6
2.2 Deep Multimodal Learning . . . . .	7
2.2.1 Transformers . . . . .	7
2.2.2 Multimodal Learning With Transformers . . . . .	10
2.2.3 The Perceiver Architecture . . . . .	12
CHAPTER 3 DEEP MULTIMODAL CARBON FLUX MODELLING . . . . .	14
3.1 The CarbonSense Dataset . . . . .	14
3.1.1 Data Collection . . . . .	15



3.1.2	Data Pipeline . . . . .	15
3.1.3	Using the Dataset . . . . .	17
3.2	The EcoPerceiver Architecture . . . . .	18
3.2.1	Data Ingestion . . . . .	19
3.2.2	Windowed Cross-Attention . . . . .	22
3.3	XGBoost Baseline . . . . .	24
3.3.1	Image Processing . . . . .	25
CHAPTER 4 EXPERIMENTAL RESULTS . . . . .		26
4.1	Experiment Setup . . . . .	26
4.1.1	Data Splitting . . . . .	26
4.1.2	Compute Infrastructure . . . . .	27
4.1.3	Model Configurations and Hyperparameters . . . . .	28
4.1.4	Reproducibility and Reliability . . . . .	29
4.1.5	Metrics and Reporting . . . . .	29
4.2	Core Experiments . . . . .	30
4.2.1	Quantitative Results . . . . .	30
4.2.2	Qualitative Analysis . . . . .	30
4.2.3	Discussion . . . . .	34
4.3	Ablation Studies . . . . .	35
4.3.1	Context Window Length . . . . .	35
4.3.2	Core Architectural Decisions . . . . .	35
4.3.3	Ecosystem-Specific DDCFM . . . . .	38
CHAPTER 5 CONCLUSION . . . . .		39
5.1	Summary of Works . . . . .	39
5.2	Limitations . . . . .	39
5.3	Future Research . . . . .	40
REFERENCES . . . . .		41
APPENDICES . . . . .		60

## LIST OF TABLES

Table 3.1	Meteorological Variables in CarbonSense . . . . .	17
Table 4.1	Train / test split distribution by IGBP type . . . . .	26
Table 4.2	XGBoost Hyperparameters . . . . .	29
Table 4.3	NSE and RMSE by model and IGBP type, aggregate mean across 10 seeds. Bold numbers indicate better performance. Pairwise t-test results are given for NSE values with 9 degrees of freedom. . . . .	31
Table 4.4	NSE and RMSE by model ablation. $t\_mean$ and $t\_std$ represent trun- cated means and standard deviations to account for the outlier biome ("WAT"). . . . .	38
Table 4.5	EcoPerceiver DBF performance when trained on DBF data vs trained on all sites . . . . .	38
Table A.1	EC Sites . . . . .	60
Table A.2	EC Sites (cont'd) . . . . .	61
Table A.3	EC Sites (cont'd) . . . . .	62

## LIST OF FIGURES

Figure 2.1	A simplified eddy covariance (EC) station. An EC system comprises a gas analyzer and a three-dimensional sonic anemometer to measure net turbulent CO <sub>2</sub> flux across eddies. . . . .	5
Figure 2.2	MODIS (left) and Sentinel-2 (right) images of the EC tower site at Scotty Creek, Northwest Territories. RGB bands here are approximated, but the spectral signatures differ between these constellations. Image covers approximately 25km <sup>2</sup> . . . . .	6
Figure 2.3	Scaled dot-product attention (left) and multi-head attention (right) as presented in [1] . . . . .	9
Figure 2.4	Original Transformer architecture as presented in [1] . . . . .	11
Figure 2.5	General architecture of Perceiver [2]. . . . .	12
Figure 3.1	Global map of EC sites used in CarbonSense, with corresponding source networks. Some sites were present in multiple networks. . . . .	16
Figure 3.2	Simplified depiction of the CarbonSense processing pipeline. . . . .	18
Figure 3.3	Overview of EcoPerceiver architecture. . . . .	19
Figure 3.4	Fourier feature encoding (left) and spectral band encoding (right) for EcoPerceiver. . . . .	21
Figure 3.5	Comparison of standard (a) and windowed (b) cross-attention, where each latent token matches an independent subset of the input. Windowed cross-attention reduces computation by a factor of $T$ . . . . .	23
Figure 3.6	Windowed Cross-Attention (WCA) block. Encoded inputs are cross-attended onto the latent space with a modality mask to indicate missing values. The time dimension is pushed into the batch dimension, so this operation is performed $B \cdot T$ times per batch. Causal self-attention proceeds as normal. . . . .	24
Figure 3.7	Simplified depiction of band compression for XGBoost. . . . .	25
Figure 4.1	Comparison of common data types in deep learning. bfloat16 allows for the same exponential range as a single-precision float at the size of a half-precision float. . . . .	27
Figure 4.2	NSE (a) and RMSE (b) scores of EcoPerceiver and XGBoost. Each chart represents 10 experiments with different seeds. Whiskers indicate approximately 1.5 x interquartile range. . . . .	32

Figure 4.3	Hourly data and model results for GF-Guy (a), an evergreen broadleaf forest station in French Guiana, and CA-LP1 (b), a pine beetle-attacked evergreen needleleaf forest in northern British Columbia. . . . .	33
Figure 4.4	Effects of context window length on NSE (a) and RMSE (b) of test site inference for EcoPerceiver. . . . .	36
Figure 4.5	Effects of various architectural choices on NSE (a) and RMSE (b) of test site inference for EcoPerceiver. . . . .	37

**LIST OF SYMBOLS AND ACRONYMS**

CC	Compute Canada
CV	Computer vision
CPU	Central processing unit
CNN	Convolutional neural network
DDCFM	Data-driven carbon flux modelling
GPP	Gross primary production
GPU	Graphics processing unit
ML	Machine learning
MML	Multimodal learning
MSE	Mean squared error
NEE	Net ecosystem exchange
NLP	Natural language processing
NSE	Nash-Sutcliffe modelling efficiency
PFT	Plant functional type
RAM	Random access memory
RECO	Ecosystem respiration
RL	Reinforcement learning
RNN	Recurrent neural network
SOTA	State-of-the-art

## LIST OF APPENDICES

Appendix A	Eddy Covariance Site Details . . . . .	60
------------	--	----

## CHAPTER 1 INTRODUCTION

### 1.1 Motivation

Our biosphere plays a critical role in regulating Earth’s climate. Since the mid-20th century, terrestrial ecosystems have absorbed up to a third of anthropogenic carbon emissions [3], with the other two thirds being distributed between the ocean and atmosphere. Terrestrial absorption has therefore acted as a buffer against more severe climate change impacts by sequestering carbon that would otherwise contribute to atmospheric greenhouse gas concentrations. However, climate change introduces uncertainty about the future resilience and capacity of these ecosystems. Understanding how the carbon dynamics of our biosphere are changing in response to both climate change and increasing anthropogenic pressures will give crucial insight into the health of our ecosystems and their continued ability to sequester carbon.

One of the best methods for quantifying ecosystem health and carbon dynamics is measuring its carbon fluxes [4]. Carbon fluxes describe the movement of carbon into and out of an ecosystem resulting from processes like photosynthesis and cellular respiration. This gives a good indication of when an ecosystem is a carbon sink and when it is a source. However, measuring carbon fluxes at ecosystem scale requires field researchers to deploy and maintain expensive sensory equipment. The most common method is eddy covariance where the sensors are mounted on towers above an ecosystem, which covers an area of 1000m<sup>2</sup> at most [4]. This presents a data bottleneck for our understanding of carbon dynamics, since it is infeasible to blanket the earth in these towers.

This bottleneck has given rise to the field of data-driven carbon flux modelling (DDCFM). DDCFM loosely refers to the inference of carbon flux data based on biophysical predictors such as meteorological data, topography, or surface reflectance. Researchers use these predictors to train statistical models on measured carbon fluxes and then run inference on other areas to produce synthetic carbon flux data, sometimes across the entirety of Earth’s surface [5, 6]. This yields global carbon flux maps that can provide crucial information to ecologists and decision-makers alike.

But the quality of this information is dependent on the ability of the model to effectively predict carbon fluxes. DDCFM is still in its infancy and has mostly been attempted with dated general-purpose machine learning (ML) techniques like random forests [7–9], gradient boosting [6], or ensembles of similar methods [10, 11]. These methods fail to capitalize on the

multimodal nature of the underlying biophysical data. Recently, multimodal deep learning has exploded in popularity [12, 13] and may offer a more appropriate framework for DDCFM through effective data integration and advanced neural architectures. If we can leverage these advances to more accurately model carbon fluxes, we will significantly enhance the quality of information available to decision-makers, thereby improving our ability to address climate change.

## 1.2 Research Objectives

The objective of this thesis is to build a deep multimodal architecture for DDCFM. Our approach will include the following steps:

1. Construct a comprehensive multimodal dataset for DDCFM.
  - (a) This dataset must contain data which is publicly available with permissive licensing in order to encourage standardization.
  - (b) It must structurally resemble the previous ad-hoc datasets in DDCFM to encourage fair comparison.
2. Construct a deep multimodal model for DDCFM.
  - (a) This model must be flexible in its ability to ingest data so as not to rely on gap-filled data<sup>1</sup>.
  - (b) It must perform well across a wide variety of ecosystems.
3. Construct a baseline DDCFM model for comparison.
  - (a) This model must use current state-of-the-art (SOTA) techniques from recent publications in DDCFM to present a fair comparison
4. Evaluate the performance of both models on the new dataset.

## 1.3 Thesis Outline

This thesis is structured to give an overview of carbon flux modelling for deep learning researchers, and to explore the development and evaluation of our model as a case study for deep multimodal learning for DDCFM. A summary of subsequent chapters is presented here:

---

<sup>1</sup>Gap-filled data is created with a model, but we want to stay as close to the raw measurements as possible. This reduces the risk of compounding model errors.



**Literature Review** We provide an overview and foundational knowledge for several topics including eddy covariance, DDCFM, transformers, and multimodal learning.

**Methods** This section details the construction of our dataset (“CarbonSense”) and our novel model architecture (“EcoPerceiver”), as well as the SOTA baseline we compare it to.

**Experiments** Here we provide details on our experiments. This includes computational infrastructure, hyperparameter tuning, experimental results, and ablation studies.

**Conclusion** Finally, we discuss the implications of our research, its limitations, and future work to build on our successes.

## CHAPTER 2 LITERATURE REVIEW

### 2.1 Data-Driven Carbon Flux Modelling

At its core, DDCFM is a regression problem. The target (carbon flux) depends on many factors including ecosystem composition, meteorological conditions, local topography and geology, and disturbances (fires, animal activity, etc). Meteorological data is relatively easy to obtain, but the other predictors are challenging to measure and represent, especially at a global scale. Remote sensing and semantic data are commonly employed as a proxy for these other predictors.

DDCFM stands in contrast to process-based modelling, where models use known physical, chemical, and biological processes to simulate carbon fluxes (and many other variables) from initial conditions and forcing data. Process-based carbon flux modelling is a parallel field beyond the scope of this work, but as of this publication DDCFM tends to produce more accurate flux predictions in most ecosystem types [14]. Here we give an overview of the data used in DDCFM, and previous works in the field.

#### 2.1.1 Measuring Carbon Fluxes: The Eddy Covariance Method

The most common technique for measuring fluxes at ecosystem scale is eddy covariance (EC) [15]. This is a micrometeorological technique where researchers erect a tower (typically above canopy height) and mount sensors that measure high-frequency fluctuations in atmospheric gas concentrations across small turbulent vortices (eddies). The idea is that the resulting fluxes will - in aggregate - represent the movement of gases into and out of the surrounding ecosystem. EC towers (interchangeably EC stations) most commonly measure  $\text{CO}_2$  and water vapour, but it is becoming increasingly common to measure methane ( $\text{CH}_4$ ) [7, 8] or nitrous oxide ( $\text{N}_2\text{O}$ ) [16] depending on the researcher's objectives. Our work focuses on  $\text{CO}_2$  due to its importance in understanding our climate and the prevalence of standardized data collections. We give a simplified depiction of an EC tower in Figure 2.1.

These towers can only measure the net turbulent  $\text{CO}_2$  flux, known as net ecosystem exchange (NEE). NEE is expressed in  $\mu\text{mol} \cdot \text{m}^{-2} \cdot \text{s}^{-1}$ . Researchers will often be interested in the component fluxes of NEE such as gross primary production (GPP, the total carbon uptake from photosynthesis) or ecosystem respiration (RECO, the total release from cellular respiration). Partitioning NEE into its component fluxes can be done using a variety of methods [17], but are beyond the scope of this thesis.




Figure 2.1 A simplified eddy covariance (EC) station. An EC system comprises a gas analyzer and a three-dimensional sonic anemometer to measure net turbulent  $\text{CO}_2$  flux across eddies.

### 2.1.2 Common Flux Predictors

**Auxilliary EC Data** In addition to carbon fluxes, EC systems measure heat fluxes and the station will typically measure local meteorological and soil conditions such as four-component net radiation, air temperature and relative humidity, precipitation, soil moisture and temperature, etc. The exact number and type of variables depends on the site and the researcher’s goals. Regional flux networks<sup>1</sup> specify a minimum mandatory set for researchers wishing to submit their data [18]. For the purpose of this work, we will use the term EC data to refer to the wealth of data variables collected by EC stations.

**Global Reanalysis Data** Biophysical variables like those collected by EC stations can be obtained on a global grid from reanalysis products such as ERA5 [19]. ERA5 provides hourly historic data by integrating observations with modern atmospheric models. This means reanalysis data is technically a model output and not a direct measurement, but it provides consistent and comprehensive coverage from January 1940 to present. Reanalysis data can be used for DDCFM in two ways: as gap-filling values for missing EC data [18], or as predictors for areas where there is no EC tower at all. The latter is useful for taking a trained DDCFM model and running regional or global inference, a process known as upscaling [6, 7, 11, 20, 21].

**Remote Sensing Data** Satellite imagery of the area surrounding an EC tower can give useful information about the land cover and ecosystem composition. The most widely used

---

<sup>1</sup>Researchers typically submit their data to a regional network for aggregation. Examples include Ameriflux (North and South America), Integrated Carbon Observation System (Europe), or CarboAfrica.

remote sensing products for DDCFM are based on Moderate Resolution Imaging Spectroradiometer (MODIS) data [22]. This satellite pair (“Aqua” and “Terra”) produce new imagery for Earth’s surface every 1-2 days and have 36 spectral bands with resolutions varying between 250m and 1km. Landsat and Sentinel-2 are also common choices for satellite data; they have a higher spatial resolution (between 15 and 30 meters depending on spectral band), but have a lower temporal frequency. This makes it more appropriate for modelling fluxes at the monthly or annual level, rather than daily or hourly. We contrast MODIS and Sentinel-2 data in Figure 2.2.




Figure 2.2 MODIS (left) and Sentinel-2 (right) images of the EC tower site at Scotty Creek, Northwest Territories. RGB bands here are approximated, but the spectral signatures differ between these constellations. Image covers approximately 25km<sup>2</sup>.

**Semantic Data** Some models ingest semantic data such as land cover (“Croplands”, “Evergreen needleleaf forest”, “Snow and ice”, etc). Land cover classifications follow standardized schemes such as the International Geosphere-Biosphere Programme (IGBP). Some MODIS products coarsely approximate this information on a global grid [23], allowing this data to also be used for global inference.

### 2.1.3 Previous Works

The application of machine learning techniques to model carbon fluxes was present as early as 2003 [20] where researchers used rudimentary neural networks for regional flux upscaling. Global upscaling models were first seen in 2011 [21], facilitated by the global research network FLUXNET which pooled EC observations from sub-networks around the world. Subsequent

releases of FLUXNET increased the data quality and number of sites [18].

Recent models such as those developed by FLUXCOM [5,6] use a combination of EC data and satellite data to improve model conditioning. This is when DDCFM became a multimodal regression task, but we note that the models being used by FLUXCOM and similar efforts were intended only for tabular data, not multimodal data. The current SOTA model for global flux modelling is arguably<sup>2</sup> XBASE [6], which uses the XGBoost [24] architecture. Satellite data is compressed down to a single pixel per spectral band, allowing it to be fed into the model alongside tabular meteorological data from EC sites. We further interrogate this process in Section 3.3.1.

Targeted modelling of specific regions such as subtropical wetlands [11] or arctic and boreal areas [7,9,10] is still common today. Focusing on a narrow ecological distribution simplifies model fitting and reduces the amount of compute resources needed for training. It also allows for the study of region-specific phenomena; for example, methane is commonly recorded in the boreal wetland regions where it is of particular scientific interest owing to its potency as a greenhouse gas [7].

## 2.2 Deep Multimodal Learning

### 2.2.1 Transformers

The Transformer architecture, proposed in 2017 by Vaswani et al. [1], is considered a foundational advancement in modern deep learning. Originally designed for machine translation tasks, Transformers have since been widely adopted across numerous deep learning disciplines, including natural language processing (NLP), computer vision (CV), and reinforcement learning (RL) [25]. The key innovation of the Transformer lies in its self-attention mechanism, which allows for the parallel processing of sequential data, overcoming the limitations of recurrent neural networks (RNNs) and convolutional neural networks (CNNs) in capturing long-range dependencies. This architecture has significantly improved both the efficiency and performance of models across various tasks, leading to breakthroughs in applications ranging from text generation to image analysis [1,25–27].

---

<sup>2</sup>We say “arguably” because researchers are not currently using a standardized benchmark to measure performance. While quantitative metrics are reported, they are not directly comparable as the train and test data can differ significantly.

## The Self-Attention Mechanism

In sequence modelling, “attention” is the ability of a model to weigh the importance of different tokens in a set of data with respect to other tokens. In NLP, these tokens may represent words or word pieces in a sentence; in CV they may represent spatially contiguous patches of pixels. Tokens simply refer to discrete pieces of a larger data structure being analyzed. Transformers use attention as their primary processing operation [1], and it takes two primary forms. In self-attention, each token in the sequence attends to every other token. In cross-attention, each token in a sequence attends to every token in a different sequence. Both self- and cross-attention are used extensively in our work.

To use self-attention as an example, we take a sequence of tokens  $X \in \mathbb{R}^{(l,h)}$  with sequence length  $l$  and token dimension  $h$  (also called the “hidden size”). Each token is linearly mapped to a query, key, and value space with

$$Q = f_Q(X) \quad K = f_K(X) \quad V = f_V(X). \quad (2.1)$$

Each of  $Q$ ,  $K$ , and  $V$  are now matrices where each row is an encoded token. By multiplying the query and key matrices we get attention logits, which are then scaled and processed with a softmax operation to produce attention weights. Intuitively, this allows each query token to determine which key tokens have the information it needs. The result is a matrix in  $\mathbb{R}^{(l,l)}$  (the attention weights) indicating how much attention each token will pay to every other token.

Once the attention weights are calculated, they are multiplied with the value matrix to determine the final attention values. In matrix notation, this is expressed with:

$$\text{Attention}(Q, K, V) = \text{softmax} \left( \frac{QK^T}{\sqrt{d_k}} \right) V \quad (2.2)$$

where  $\sqrt{d_k}$  is a scaling factor dependent on the length of the key tokens. A computation graph of this operation is shown in 2.3. While this operation is powerful on its own, attention is typically used in multi-head form where many attention operations are run in parallel with different weights for  $f_{Q/K/V}$  on each head. This allows each token to attend to every other token in multiple contexts.

$$\begin{aligned} \text{MultiHead}(Q, K, V) &= \text{Concat}(\text{head}_1, \dots, \text{head}_2) W^O \\ \text{where } \text{head}_i &= \text{Attention}(QW_i^Q, KW_i^K, VW_i^V) \end{aligned} \quad (2.3)$$

Here  $W_i^Q$ ,  $W_i^K$ , and  $W_i^V$  are learned projection matrices for  $Q$ ,  $K$ , and  $V$  respectively, and  $W^O$  is a final projection matrix to return the output to the original dimension of  $X$ .




Figure 2.3 Scaled dot-product attention (left) and multi-head attention (right) as presented in [1]

### The Cross-Attention Mechanism

Cross-attention is another operation introduced in [1]. During cross-attention, each token in a target sequence attends to all tokens in a separate sequence. The classic example of this is translation, where each word in the target sentence will be influenced by its attention to all words in the source sequence. This helps to overcome the semantic ordering of ideas in language, such as noun-adjective ordering.

Mathematically, cross-attention is identical to self-attention apart from the initial encoding in 2.1 which is revised as

$$Q = f_Q(X_t) \quad K = f_K(X_s) \quad V = f_V(X_s). \quad (2.4)$$

where  $X_t$  is the target sequence and  $X_s$  is the source sequence. Again, we intuitively see this as each token in the target sequence looking for information, with the key and value matrices providing said information from the source sequence. Cross-attention is less ubiquitous than self-attention generally, but it is extremely relevant to the use of transformers in multimodal learning [12].

## Transformer Architectures

The original Transformer architecture was used for machine translation. To achieve this, the Vaswani et al. created an encoder and a decoder in a similar manner to previous RNNs [1]. The encoder used self-attention on the input (source) sequence, while the decoder alternated between masked self-attention on the output (target) sequence and cross-attention with the input sequence. This is shown in Figure 2.4.

Because attention is positionally agnostic (the order of tokens does not matter), the original Transformer paper added sinusoidal positional embeddings to the input and output sequences prior to processing, allowing the algorithm to learn where each token belongs in the sequence. This is still common in most Transformer architectures, though many researchers opt for simple learned embeddings rather than sinusoidal ones.

Transformers do not always require an encoder and a decoder; they often need only one or the other depending on the task. Encoder-only Transformers are strictly self-attentive and do not have positional restrictions on attention. The Vision Transformer [27] is a good example of this, where each token is a patch of pixels and can attend to any other patch. Decoder-only Transformers are also self-attentive, but each token may only attend to tokens at previous positions. This is achieved by masking the attention logits. Decoder-only Transformers are common in language modelling such as GPT [26].

### 2.2.2 Multimodal Learning With Transformers

Multimodal learning (MML) is a subfield of ML that integrates and processes information from multiple modalities (such as vision and audition) and is thought to imitate the way humans learn and understand the world [12, 28, 29]. Neural networks have been used to study MML as early as 1989 [30], but the advances in deep learning over the last decade have vastly improved our multimodal models.

Transformers have been particularly effective for MML tasks. They do not have the same degree of inductive bias we see in CNNs and RNNs and can operate in a modality-agnostic manner. This flexibility allows Transformers to process different types of data uniformly, facilitating the integration of heterogeneous information [12]. There are two categories of Transformer models for MML: pretrained models which seek strong generalization to a diverse set of downstream tasks, and application-specific models where data fusion is applied to a targeted problem [12, 29].



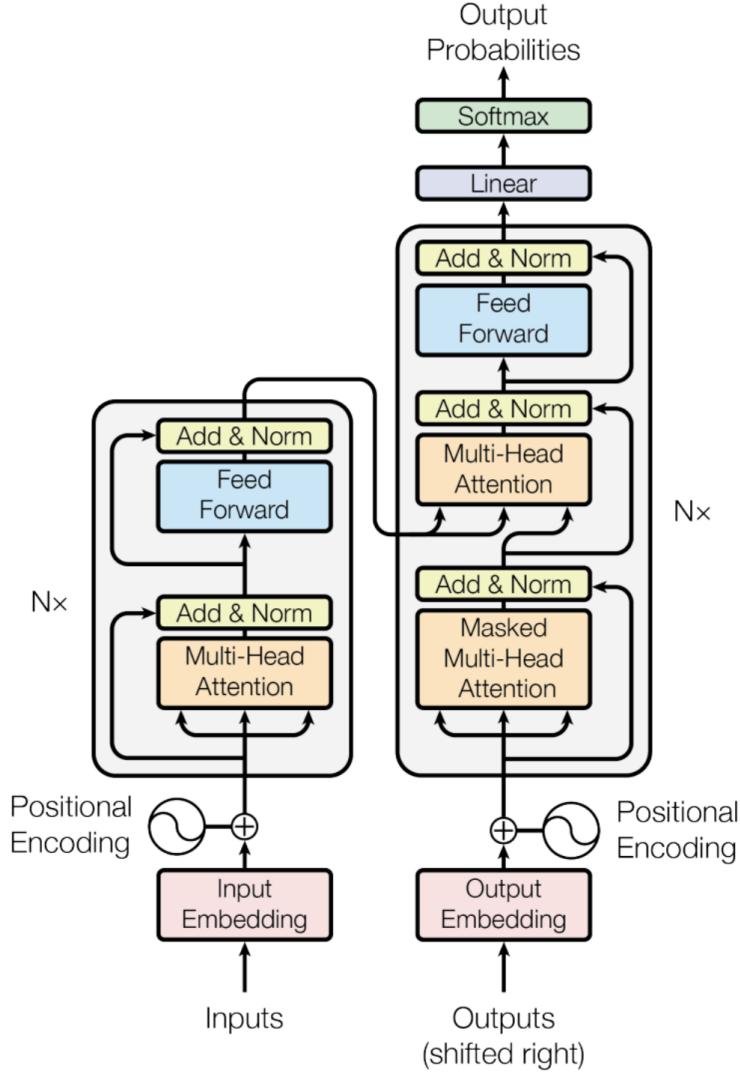


Figure 2.4 Original Transformer architecture as presented in [1]

## Pretrained Multimodal Transformers

Pretrained multimodal Transformers are pretrained on large corpora of data - typically images and text - to achieve strong few-shot performance on downstream tasks [12,29]. This is similar to how large language models are used; pretraining on virtually the entire internet’s worth of text to achieve strong zero-shot performance on language tasks [26]. Models trained in this way are often referred to as “foundation models”.

As an example, the multimodal vision-language learning paradigm with context-based compact Transformer (MPCCT) is a general framework for training Transformers for visual-language tasks [31]. The authors train their compact Transformer model on common image-

text datasets such as COCO and CLEVR to obtain SOTA performance on visual grounding and visual question answering. This type of text and image fusion appears to be the most common type of pretrained multimodal Transformer, likely owing to the availability of appropriate training data.

## Application-Specific Multimodal Transformers

Application-specific MML applies the same principles to more focused tasks [12]. Generally, models are built to ingest data modalities specific to a task such as video and text for sentiment analysis [32] or radiographs and lab results for clinical diagnostics [33]. The idea is to achieve higher performance by leveraging the relationship between data modalities, but without the expectation of generalizing to other downstream tasks.

Geospatial applications in particular have benefited from advances in multimodal Transformers [13]. The authors of ExVIT [34] fuse hyperspectral imagery and LiDAR to achieve SOTA land use cover classification. More recently, Benson et al. developed Contextformer [35], a multimodal temporal Transformer which fuses satellite imagery, meteorological data, and elevation maps for vegetation forecasting. Importantly, Contextformer was released alongside its bespoke training dataset (GreenEarthNet) to encourage further research.

### 2.2.3 The Perceiver Architecture

Perceiver is a modern Transformer architecture for modality-agnostic processing [2]. It was built to handle arbitrary input modalities and to avoid strong architectural priors. The authors describe this as “encouraging the data to speak for itself”. The general idea is to do most of the processing in a compact latent space, with inputs being mapped onto this space using cross attention. Figure 2.5 depicts this architecture at-a-glance.

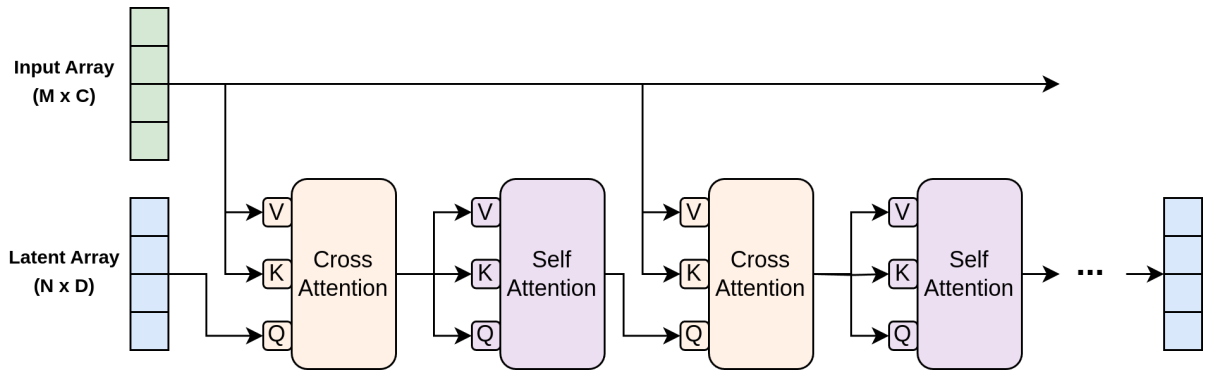


Figure 2.5 General architecture of Perceiver [2].

The latent array is an entirely learned parameter and can have different dimensions from the inputs (this being one of the advantages of cross-attention). The first part of the model alternates cross-attending inputs onto the latent array and self-attention on the latent array. This is repeated many times, refining the latent array and then re-querying the input again. Once complete, a linear layer can be used to obtain logits from the latent array (without loss of generality - this can change as needed by the task).

While the architecture itself contains minimal inductive bias, the input array must be constructed in a way that is sensible to the task. All inputs must be tokenized, and different modalities can be given embeddings to concatenate onto the tokens. For images this will include a positional embedding, and such learned parameters help the model identify what type of input it is handling during the forward pass.

Perceiver presents the multimodal equivalent of a “blank slate”; the authors avoided inductive bias in its architecture, but such bias can be acceptable and even beneficial for application-specific modelling. While the original intent may have been to frame Perceiver as a pretrained model for downstream tasks (hence the tagline “General Perception”), the architecture’s simplicity allows it to be adapted for application-specific MML, as we will explore in the following chapters.

## CHAPTER 3 DEEP MULTIMODAL CARBON FLUX MODELLING

In this section we present our primary contributions. First we introduce CarbonSense, an ML-ready dataset designed specifically for DDCFM. This dataset incorporates EC data and satellite imagery from 385 sites worldwide in a format suitable for training and benchmarking ML models. We will discuss the data sources, the processing pipeline, and the final dataset structure.

Following this, we present EcoPerceiver, a novel multimodal Transformer model for DDCFM, inspired by the Perceiver architecture. This section will cover the motivations behind developing EcoPerceiver, its architectural design, and the implementation details that make it adept at handling the multimodal nature of the CarbonSense dataset. We also discuss the dataloader used as the interface between our dataset and model.

Finally, we discuss the recreation of a baseline model based on the current SOTA. This will be an important comparative tool for gauging the effectiveness of deep learning techniques for DDCFM.

### 3.1 The CarbonSense Dataset

The first step in training an ML model is obtaining an appropriate dataset. When operating in a unimodal domain such as image classification, large pretraining datasets like ImageNet [36] and its successors are commonly used. Multimodal pretraining datasets are becoming more accessible for common modality pairs, such as LAION-5b [37] for image and text. However, niche applications such as DDCFM do not have off-the-shelf datasets due to lower historical research interest. Previous works have therefore created datasets ad-hoc according to their research needs [5–7, 11, 20, 21].

Our solution is to create an ML-ready dataset using similar sources and methods as described in previous works [5, 6, 22]. CarbonSense uses a combination of hourly EC data<sup>1</sup> and daily satellite data as predictors, with hourly carbon fluxes (NEE, GPP, and RECO) as regression targets. CarbonSense is unique in that it is openly available as an end-product for other researchers. By structuring the dataset similar to previous works and releasing it freely, we aim to reduce the barrier to entry for other researchers in DDCFM.

---

<sup>1</sup>In this work, we will use "EC data" as shorthand to refer to CO<sub>2</sub> and heat fluxes, along with all ancillary meteorological data collected by the EC stations.

### 3.1.1 Data Collection

All EC data are aggregated from major networks. Gathering data from individual stations would be intractable and would require extensive processing of each site to ensure common predictor names, units, and gapfilling. For this reason we draw from networks which use the ONEFlux processing pipeline: FLUXNET 2015 [18], the Integrated Carbon Observation System (ICOS) 2023 release [38], ICOS Warm Winter release [39], and Ameriflux 2023 release [40]. A map of EC sites from these source networks is shown in Figure 3.1. North America and Europe are over-represented in this site list due greater data accessibility, and we discuss the implications of this in Section 3.1.3.

Geospatial data in CarbonSense are sourced from MODIS products. Specifically, we use the MCD43A4 and MCD43A2 products [41, 42]. These products fuse MODIS data in a 16-day sliding window to produce a single image each day. This helps to address cloud coverage and produces images which remove angle effects from directional reflectance. Each image therefore appears as it would from directly overhead at solar noon. MCD43A4 contains 7 spectral bands between 620nm and 2155nm, at 500 meters per pixel. MCD43A2 contains categorical values for each pixel indicating snow and water cover. The decision to use MCD43A4 and MCD43A2 was based on the success of these products in XBASE [6].

Following the guidelines from [22], we extract images in a 4km by 4km square centered on each EC station. Given a spatial resolution of 500m per pixel, this yields an 8x8 pixel image with 9 channels for every site-day. Refer back to Figure 2.2 for an example image (though that one had not yet been cropped to 8x8 pixels).

### 3.1.2 Data Pipeline

We begin by fusing site data from the source networks. Many sites had overlapping data from different sources. For example, the site Degero in Sweden (SE-Deg) had data from 2001-2020 in the ICOS Warm Winter release, and data from 2019-2022 in the ICOS 2023 release. Data are fused with overlapping values taken from the more recent release as in previous DDCFM work [6]. Any sites which report half-hourly data are downsampled to hourly at this stage, and daily and monthly recordings are discarded.

Once fused, we extract the relevant time blocks for each EC station along with its geographic location. We then calculate the latitude and longitude coordinates for a 4km square centered on each station. This metadata is used to obtain the appropriate MODIS data for each site. We pull this data procedurally from Google Earth Engine [43]. Each MODIS image is reprojected to the Universal Transverse Mercator (UTM) projection zone appropriate for its

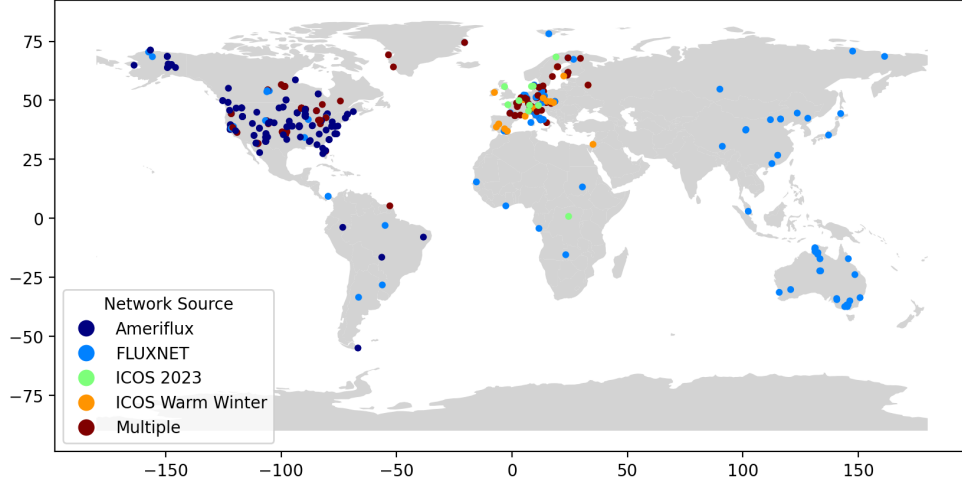


Figure 3.1 Global map of EC sites used in CarbonSense, with corresponding source networks. Some sites were present in multiple networks.

site and cropped to 8x8 pixels centered on the EC tower. This reprojection is necessary to correct for the distortion in MODIS’s native sinusoidal projection at extreme latitudes.

Meteorological data is pruned to remove unwanted variables. Some, like soil moisture and temperature, are either unavailable for most sites or are heavily gap-filled. We remove these variables to reduce the risk of compounding errors on the underlying pipeline gapfilling techniques. A full list of variables at this stage is given in Table 3.1.

As a final stage in the pipeline, we apply a min-max normalization on predictor variables. Data normalization is ubiquitous in ML, but we choose to use min-max normalization because some of the predictor variables are cyclic, such as wind direction or time of day. We map cyclic variables to the range  $[-1, 1)$  and acyclic variables to the range  $[-0.5, 0.5)$ . This normalization scheme is tailored toward Fourier feature encoding, which we discuss in detail in section 3.2.1.

We offer CarbonSense as a finished dataset but also provide the raw data<sup>2</sup>. The full pipeline code (written in Python) is available so that researchers can run and modify it freely. Our pipeline can be configured to include other variables or to have different “leniency” for gap-filled values. For example, those who wish to use CarbonSense with strictly observed values may do so at the cost of a smaller number of samples. A diagram of the entire pipeline is shown in Figure 3.2.

<sup>2</sup>CarbonSense can be downloaded at <https://zenodo.org/records/11403428>

Table 3.1 Meteorological Variables in CarbonSense

Code	Description	Units
<b>Predictors</b>		
TA_F	Air temperature	deg C
PA_F	Atmospheric pressure	kPa
P_F	Precipitation	mm
RH	Relative humidity	%
VPD_F	Vapor pressure deficit	hPa
WS_F	Wind speed	m s <sup>-1</sup>
USTAR	Wind shear	m s <sup>-1</sup>
WD	Wind direction	decimal degrees
NETRAD	Net radiation	W m <sup>-2</sup>
SW_IN_F	Incoming shortwave radiation	W m <sup>-2</sup>
SW_OUT	Outgoing shortwave radiation	W m <sup>-2</sup>
SW_DIF	Incoming diffuse shortwave radiation	W m <sup>-2</sup>
LW_IN_F	Incoming longwave radiation	W m <sup>-2</sup>
LW_OUT	Outgoing longwave radiation	W m <sup>-2</sup>
PPFD_IN	Incoming photosynthetic photon flux density	μmol Photon m <sup>-2</sup> s <sup>-1</sup>
PPFD_OUT	Outgoing photosynthetic photon flux density	μmol Photon m <sup>-2</sup> s <sup>-1</sup>
PPFD_DIF	Incoming diffuse photosynthetic photon flux density	μmol Photon m <sup>-2</sup> s <sup>-1</sup>
CO <sub>2</sub> _F_MDS	CO <sub>2</sub> atmospheric concentration	μmol CO <sub>2</sub> mol <sup>-1</sup>
G_F_MDS	Soil heat flux	W m <sup>-2</sup>
LE_F_MDS	Latent heat flux	W m <sup>-2</sup>
H_F_MDS	Sensible heat flux	W m <sup>-2</sup>
<b>Targets</b>		
NEE_VUT_REF	Net Ecosystem Exchange (variable USTAR)	μmol CO <sub>2</sub> m <sup>-2</sup> s <sup>-1</sup>
GPP_DT_VUT_REF	Gross Primary Production (daytime partitioning)	μmol CO <sub>2</sub> m <sup>-2</sup> s <sup>-1</sup>
GPP_NT_VUT_REF	Gross Primary Production (nighttime partitioning)	μmol CO <sub>2</sub> m <sup>-2</sup> s <sup>-1</sup>
RECO_DT_VUT_REF	Ecosystem Respiration (daytime partitioning)	μmol CO <sub>2</sub> m <sup>-2</sup> s <sup>-1</sup>
RECO_NT_VUT_REF	Ecosystem Respiration (nighttime partitioning)	μmol CO <sub>2</sub> m <sup>-2</sup> s <sup>-1</sup>

### 3.1.3 Using the Dataset

**Site Sampling** The biased geographic and ecological distribution of sites remains a challenge in DDCFM, and CarbonSense is no different. Given the significant overrepresentation of certain regions (North America, Europe) and ecosystems (evergreen needleleaf forests, grasslands), we maintain a partitioned structure where each site has its own directory containing EC data, satellite data, and metadata. Researchers are encouraged to select sites for training and testing based on their experiment objectives such as high performance on particular ecosystems, or out-of-distribution generalization. Our experiments in section 4.2.1 give an example of the latter.

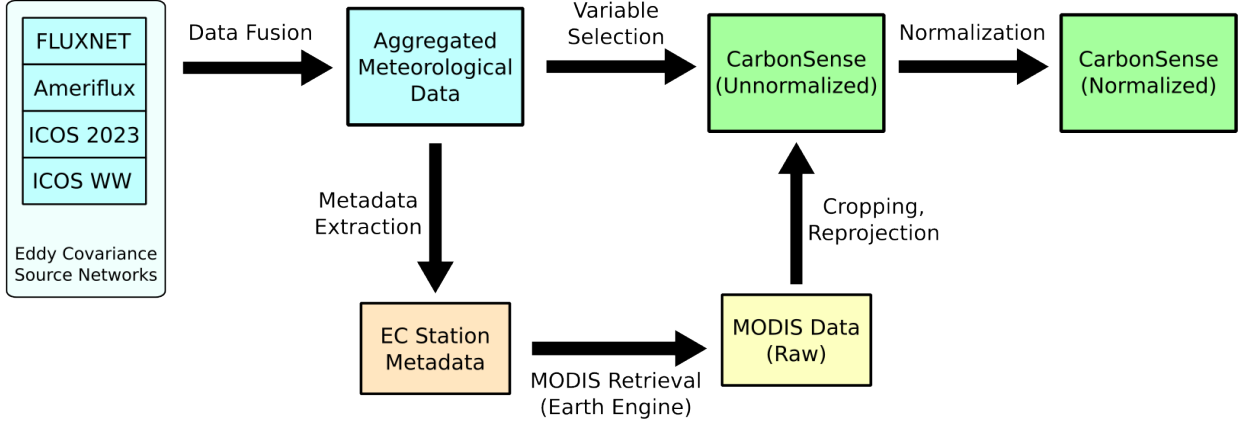


Figure 3.2 Simplified depiction of the CarbonSense processing pipeline.

**Dataloader** We supply an example PyTorch dataloader for CarbonSense specifically tailored to our model. Using the dataloader requires specifying which carbon flux to use as the target (i.e. NEE, GPP, RECO), which sites to include in each dataloader instance, and the context window length for multi-timestep training.

**Licensing** CarbonSense is available under the CC-BY-4.0 license, meaning it can be shared, transformed, and used for any purpose given proper attribution. This is an extension of the same license for all three source networks, and MODIS data is provided under public domain. We feel that permissive licensing is essential in order to foster greater scientific interest in DDCFM.

### 3.2 The EcoPerceiver Architecture

Here we present an overview of EcoPerceiver, our deep MML architecture for DDCFM. It ingests EC and satellite data in a moving temporal window to predict hourly carbon fluxes. To our knowledge it is the first DDCFM model using any kind of Transformer-based architecture, and we designed it with the following objectives:

**Input Flexibility** Different EC stations will measure different meteorological variables. Additionally, sensors will often fail and leave coverage gaps, or outlier values will be removed during post-processing. Rather than rely on gapfilling techniques, we design EcoPerceiver to be robust to missing inputs.



**Non-Markovian Processing** All available DDCFM models treat carbon dynamics as a Markovian process; they assume an ecosystem’s  $\text{CO}_2$  uptake and respiration are determinable using only immediate information. However, biological processes do not follow this assumption. A plant’s photosynthesis may depend on temperature, solar radiation, and precipitation levels spanning hours or days into the past (or further!). We design EcoPerceiver to condition its output on all data in a fixed context window, and explore the effect of varying window lengths in our ablation experiments.

**Efficiency** The downstream use of DDCFM is to produce local, regional, or global predictions of carbon fluxes. These are compute-heavy tasks often run by researchers who do not have access to large GPU clusters. While our analysis is limited to site-level experiments, we design EcoPerceiver to have a modest parameter count and runtime to accommodate researchers who wish to use it for upscaling.

Figure 3.3 gives an overview of EcoPerceiver’s end-to-end data flow. In this section we will explore how it ingests and processes data in an effective and ecologically valid way.

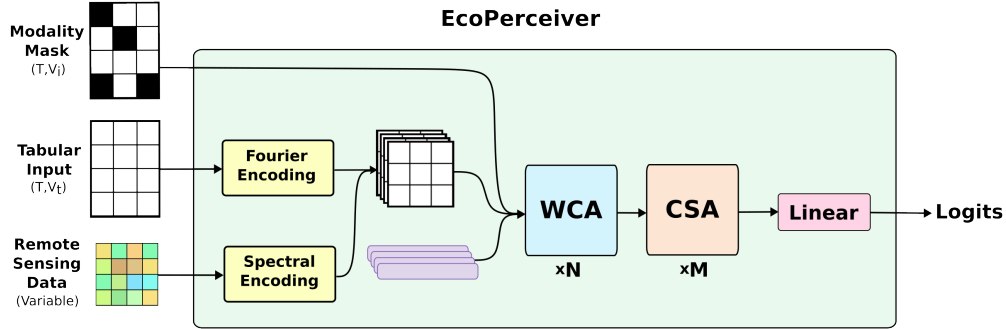


Figure 3.3 Overview of EcoPerceiver architecture.

### 3.2.1 Data Ingestion

As discussed in Section 2.2.1, Transformers take tokenized data as input. This presents an immediate problem for EC data, which is tabular by default. This means each predictor is a single value, and must be projected into a vector for tokenization. Additionally, neural networks are biased toward learning low frequency functions [44]. This is problematic because small changes in meteorological variables can have a large impact on an ecosystem’s productivity. We need a model which is sensitive to small changes in input, i.e. it learns high frequency functions. We solve both of these problems with our encoding scheme.

## Fourier Feature Encoding

EcoPerceiver uses Fourier feature encoding (FFE) [45] to convert tabular inputs into vectors, in a similar way to Neural Radiance Fields [46]. FFE maps continuous values to higher dimensional space by passing them through high frequency sinusoids. Suppose we wish to encode a variable  $x$ . FFE produces a vector with:

$$f(x; K) = \left[ \dots, \sin(2^k \pi x), \cos(2^k \pi x), \dots \mid k \in [0, K) \right], \quad (3.1)$$

where  $K$  is a hyperparameter indicating the maximum sampling frequency. Higher values of  $K$  allow the model to better discern between small differences in input. This has the added benefit of converting single values into vectors, which can be used as input tokens.

As an example, suppose we have two samples with air temperature values of 0.16 and 0.20 after normalization (this would equate to 12.8 and 16 degrees celsius). The difference between 0.2 and 0.3 is very small to a neural network, but a difference of 3.2 degrees celsius may affect ecosystem productivity. Under our encoding scheme with  $K = 3$ , we would achieve the following vectors:

$$\begin{aligned} f(0.16; 3) &= [0.482, 0.876, 0.844, 0.536, 0.905, -0.426] \\ f(0.20; 3) &= [0.588, 0.809, 0.951, 0.309, 0.588, -0.809] \end{aligned}$$

As the value of  $K$  increases, similar inputs will start to have diverging vector values, allowing the network to easily pick up on the difference. With this encoding scheme, input values of  $-1$  and  $1$  will produce identical vectors. We can show this algebraically. Given:

$$\begin{aligned} \sin(n\pi) &= 0 & \forall n \in \mathbb{Z} \\ \therefore \sin(2^k \pi) &= \sin(-2^k \pi) & \forall k \in \mathbb{N} \end{aligned} \quad (3.2)$$

$$\begin{aligned} \text{and} \quad \cos(n\pi) &= \cos(-n\pi) & \forall n \in \mathbb{Z} \\ \therefore \cos(2^k \pi) &= \cos(-2^k \pi) & \forall k \in \mathbb{N} \end{aligned} \quad (3.3)$$

we can show that:

$$\begin{aligned} f(-1; K) &= \left[ \dots, \sin(-1 * 2^k \pi), \cos(-1 * 2^k \pi), \dots \mid k \in [0, K) \right] \\ &= \left[ \dots, \sin(2^k \pi), \cos(2^k \pi), \dots \mid k \in [0, K) \right] & (\text{By 3.2, 3.3}) \\ &= f(1; K) \blacksquare \end{aligned}$$

This is why CarbonSense uses min-max normalization with cyclic values mapped to  $[-1, 1)$ ; a wind direction of 0 degrees and 359 degrees should be treated almost identically. Conversely, acyclic values are mapped to  $[-0.5, 0.5)$  so that the lower and upper extremes will have different encoded values that the model can latch onto.

After this transformation, each input is given a learned embedding specific to the underlying variable. This is then concatenated with the FFE vector to produce a final input vector of length  $H_i = 2K + l_{emb}$  for each input. Figure 3.4 depicts this encoding procedure.

Satellite data is similarly processed, except that each spectral band is flattened into a one-dimensionally array which is then linearly projected to length  $2K$ . Each band is then given an embedding to produce a vector of length  $H_i$  similar to the tabular inputs. The tabular inputs are then stacked with the encoded spectral inputs to create a matrix of shape  $(T, V_t, H_i)$ , where  $V_t$  is the total number of inputs, and  $T$  is the context window length.

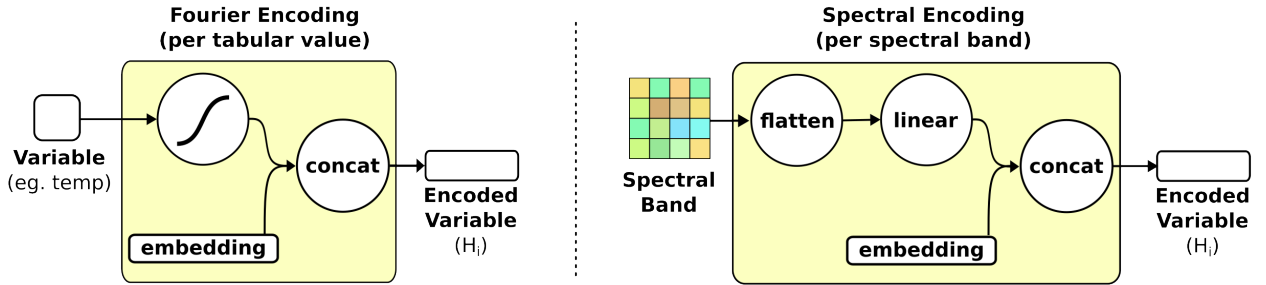


Figure 3.4 Fourier feature encoding (left) and spectral band encoding (right) for EcoPerceiver.

EcoPerceiver takes a modality mask as input in addition to the tabular and spectral data. EC data tends to have coverage gaps from equipment failures or the removal of outlier data during preprocessing. Rather than rely on gapfilling, we use the modality mask to indicate missing variables in the input. This is necessary when the model is batch processing since the input is a matrix with rows representing specific variables at every timestep. During the cross-attention operation, the modality mask is used to zero out attention logits before the softmax operation, so the model will simply ignore masked tokens.

An added bonus of the modality mask is that it can be used to randomly mask out variables as a form of dropout during training. We refer to this practice as “observational dropout” and explore its impact in our ablation studies.

### 3.2.2 Windowed Cross-Attention

Once the input is properly encoded, EcoPerceiver begins cross-attending it onto the latent array as seen in Perceiver. Suppose we have a total of  $V_i$  inputs (tabular variables + spectral bands) and a context window of  $T$  timesteps. Assuming our query, key, and value hidden size is  $H_L$ , then our key and value arrays would have size  $(TV_i, H_L)$  and our query array would have size  $(L, H_L)$  where  $L$  is the number of tokens in our latent space. Our expected runtime for cross attention (Formula 2.2) would then be:

$$O(LH_LTV_i)$$

This makes sense, as the runtime of a matrix multiplication  $A^{(l,m)} \times B^{(m,n)}$  scales with  $O(lmn)$ , and the attention operation is bounded by two matrix multiplications (first of  $QK^T$ , then  $(QK^T)V$ ) [1]. We choose a latent array with  $T$  tokens for EcoPerceiver; this is sensible because larger context windows will scale the latent space capacity, and it gives a convenient and ecologically valid inductive bias to our model. Each token in the latent space can be thought of as representing the ecosystem’s “state” at a given timestep.

This means  $L = T$  and our runtime is now bounded by  $O(T^2H_LV_i)$ , scaling quadratically with the context window length. This is in line with self-attention which scales quadratically with the number of tokens [1], but it presents us with a computational bottleneck. We want to use as large a context window as possible to maximize the information given to the model, but quadratic scaling makes that incompatible with our efficient design objective.

To solve this dilemma, we take inspiration from SWin Transformer’s windowed processing [47]. SWin uses self-attention only within isolated subsets (“windows”) of its latent space. To improve efficiency, it pushes the window index into the batch dimension. We apply the same principle here. Each timestep can be thought of as a “window” in the cross attention operation, so the latent token at time  $\tau$  only needs to attend to inputs at time  $\tau$ . We show this distinction in Figure 3.5. Using vanilla cross attention, each token is only concerned with at most  $\frac{1}{T}$  observations with an attention mask removing the rest.

By pushing these windows into the batch dimension, we are left with a set of  $T$  cross-attention operations with an key array of size  $(H_L, V_i)$  and a single latent token of size  $(1, H_L)$ . The total runtime is therefore reduced to

$$O(TH_LV_i). \tag{3.4}$$

We refer to this operation as windowed cross-attention (WCA), and it allows EcoPerceiver’s


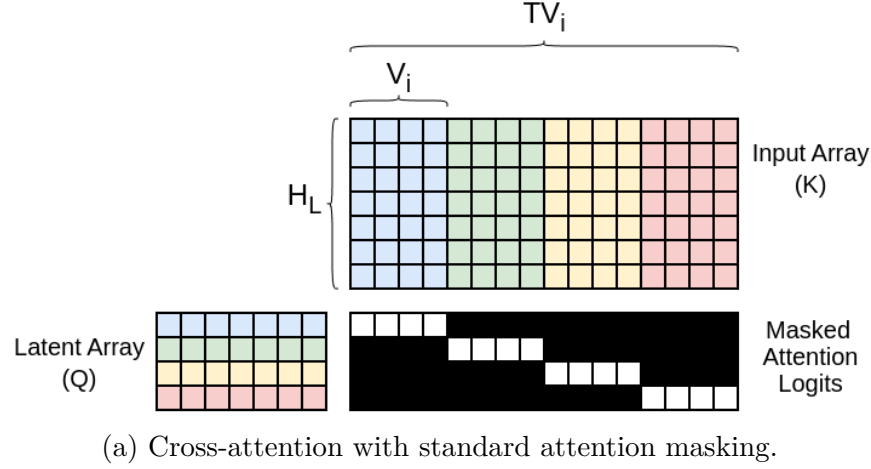


Figure 3.5 Comparison of standard (a) and windowed (b) cross-attention, where each latent token matches an independent subset of the input. Windowed cross-attention reduces computation by a factor of  $T$ .

runtime to scale linearly with context window length rather than quadratically, and lends an intuitive quality to the way the model processes data.

### Windowed Cross-Attention Blocks

In keeping with Perceiver, each WCA operation is followed by a self-attention operation in the latent space. This self-attention is given a causal attention mask, meaning each latent token can only attend to latent tokens from previous timesteps. We refer to this as causal self-attention (CSA) and it is the other primary operation in EcoPerceiver. The sequence of WCA from the input and CSA on the latent space comprises a WCA block. A high level view of the WCA block is shown in Figure 3.6.

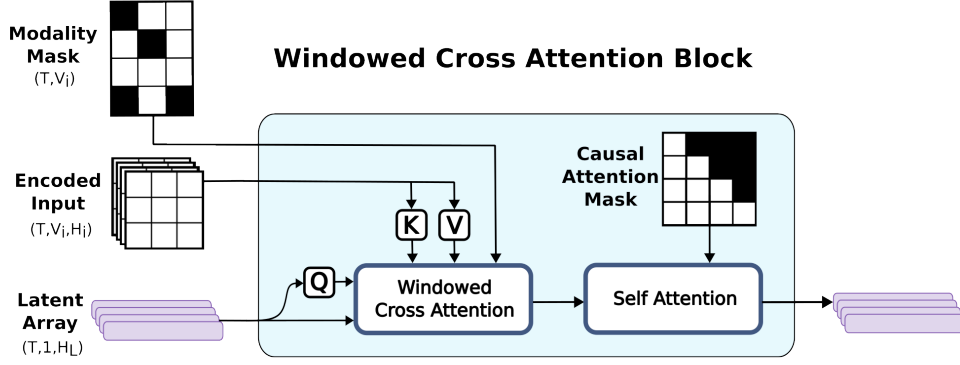


Figure 3.6 Windowed Cross-Attention (WCA) block. Encoded inputs are cross-attended onto the latent space with a modality mask to indicate missing values. The time dimension is pushed into the batch dimension, so this operation is performed  $B \cdot T$  times per batch. Causal self-attention proceeds as normal.

CSA serves two purposes in the WCA block. First, it refines the latent tokens before allowing them to re-query the input as previously discussed in Section 2.2.3. Second, it allows for information passing across timesteps. As mentioned before, the carbon flux of an ecosystem can depend on prior conditions, but WCA prevents any latent token from directly observing conditions from other timesteps. Alternating between WCA and CSA operations allows a latent token to attend to previous ecosystem states to inform its own state and how it might further query its input.

WCA blocks are repeated  $N$  times as shown in Figure 3.3, followed by  $M$  blocks of just CSA operations. We find that having a period of self attention toward the end of the model improves performance. Carbon fluxes are then obtained by applying a linear transform to the final latent state.

### 3.3 XGBoost Baseline

In order to properly test the performance of EcoPerceiver, we need a baseline model as well. Many tabular ML methods are used in the recent DDCFM literature, but the most comprehensive study comes from the FLUXCOM group [6] where they use XGBoost [24]. This study provided the inspiration for CarbonSense’s construction; the authors also used EC data from Ameriflux, ICOS, and FLUXNET, and also used MCD43A4 and MCD43A2 as their satellite data. For this reason, it makes sense to use the same ML algorithm as it will provide the closest possible comparison for our work.

XGBoost - a contraction of eXtreme Gradient Boosting - is a gradient boosting decision tree algorithm built to be highly flexible and efficient. Like most decision tree methods, XGBoost operates best with tabular data, and it can handle both continuous values and categorical values. It is highly regarded in the data science community as a go-to choice for processing this kind of data, in both classification and regression domains.

### 3.3.1 Image Processing

We encounter a challenge when trying to feed image data into XGBoost. The naive solution is to flatten the image and pass each pixel in as a predictor. However, even for our small 8x8 pixel images, this results in 64 additional features per spectral band. While XGBoost is generally resilient to sparse feature spaces, this proportion of pixel predictors would result in many of the decision trees entirely consisting of loosely correlated pixels. This can quickly drown out the predictive power of the tabular data.

The authors of [6] address this by compressing the images into a single value per spectral band. They employ an algorithm for averaging pixel values within a band depending on the number of viable pixels and the quality check values for each pixel. We take a similar but more simplified approach. For each spectral band, we take a weighted average of pixels based on their Euclidean distance from the center of the image. Empty pixels are removed before this processing, so an image with several empty pixels will result in higher weights for the remaining pixels. Any image with greater than 50% empty pixels is discarded outright. Figure 3.7 shows this process at a high level.




Figure 3.7 Simplified depiction of band compression for XGBoost.

For the purposes of our experiments, all images are processed ahead of time and the averaged pixels are added to the other predictors in comma separated value (“csv”) files. This effectively creates a parallel dataset that is used for XGBoost experiments. Processing this ahead of time vastly reduces the time spent during hyperparameter tuning and training.

## CHAPTER 4 EXPERIMENTAL RESULTS

In this chapter we present a series of experiments which evaluate the effectiveness of deep MML for DDCFM. We train our EcoPerceiver architecture on the CarbonSense dataset and compare the results against our XGBoost model. As mentioned in Section 3.3.1, the images in CarbonSense are reduced to single-pixel bands when training on XGBoost in a similar manner to [6], but the data is otherwise identical between the experiments.

This chapter begins with an overview of the data splitting based on plant functional types (PFTs; a way of categorizing ecosystems). We then detail the compute infrastructure and model configurations for the core experiments, including hyperparameter tuning. Next we present and discuss the quantitative and qualitative results of the main experiments. Finally, we show a collection of ablation experiments which provide context to some of EcoPerceiver’s architectural choices.

### 4.1 Experiment Setup

#### 4.1.1 Data Splitting

CarbonSense contains 385 EC sites from around the world, but it would be naive to use a random subsample as a holdout set due to the imbalanced distribution of PFTs. Instead we leverage the partitioned structure of the dataset to split the data based on PFT. In this case we use the IGBP classification system for PFTs, which we will refer to as the “IGBP type” in this work. We mostly refer to IGBP types by their acronyms for brevity, but a list of IGBP types with expanded names is found in Appendix A.

Despite the imbalance of IGBP types in CarbonSense, we wanted the test set to be as balanced as possible. The number of sites in the test set were determined with  $\min(5, \lceil 0.2 * \text{num\_sites} \rceil)$ . This provided between 1 and 5 sites per IGBP type as shown in Table 4.1. As a consequence, SNO and DNF provide information about zero-shot generalization (since they have zero sites in the train set), and CVM and WAT provide information about one-shot generalization.

Table 4.1 Train / test split distribution by IGBP type

IGBP	CRO	CSH	CVM	DBF	DNF	EBF	ENF	GRA	MF	OSH	SAV	SNO	WAT	WET	WSA
<b>Train</b>	44	5	1	42	0	10	80	59	10	25	11	0	1	42	8
<b>Test</b>	5	2	1	5	1	3	5	5	3	5	3	1	1	5	2



The main focus of this research is on models trained *across* different IGBP types, as opposed to other research studying DDCFM *within a single* type (ex [7, 10, 11]). However, the partitions in CarbonSense make it flexible for different modelling objectives, such as individual ecosystems. We give an example of this type of experiment in our ablation studies.

#### 4.1.2 Compute Infrastructure

All experiments were run on the Compute Canada (CC) Narval cluster. XGBoost experiments were run on compute nodes using 16 CPU cores and 64GB of RAM. EcoPerceiver experiments were each run on 4 A100 GPUs using dataset parallelization. This allowed our model to train to convergence within 12 hours. Each GPU was controlled by an independent process and had access to 8 CPU cores and 64GB of RAM. A large amount of RAM is necessary because our dataloader loads the entire dataset into memory to cut down on processing time.

The use of A100 GPUs is significant. NVIDIA GPUs from the Ampere generation (2020) or later support the use of the bfloat16 datatype. This is a variation of the half-precision (16-bit) floating point standard which allocates more bits to the exponent at the cost of fraction bits. We show this in Figure 4.1. Half-precision floats cause our model to suffer overflow problems, while single-precision floats take up to four times longer to train. The benefits of bfloat16 for our model were instrumental in choosing our compute infrastructure.




Figure 4.1 Comparison of common data types in deep learning. bfloat16 allows for the same exponential range as a single-precision float at the size of a half-precision float.

### 4.1.3 Model Configurations and Hyperparameters

#### EcoPerceiver

Hyperparameter tuning for EcoPerceiver comprised the bulk of the experiment efforts. Where possible, we started with our best guesses and then adjusted hyperparameters via random perturbation. A true random search of the parameter space would have been extremely sparse given the available compute resources.

We set our latent hidden size to 128, our input embedding size to 16, and the number of FFE frequencies to 12. This gave a total input hidden size of 40. Our context window is 32, meaning our model sees the previous 32 hours of observations. We use 8 WCA blocks followed by 4 CSA blocks. We set our observational dropout ratio to 0.3 (so 30% of all inputs are dropped out before WCA) and use causal masking in all self-attention blocks. In keeping with [2], we employ weight sharing between all WCA blocks. With these hyperparameters our model weighs in at a very reasonable 988,633 parameters.

Heavier configurations were considered, but performance gains were minimal (see ablation studies) and the compute tradeoff made it impractical for anyone without multi-GPU cluster access to use the model. This is especially true for increasing the context window or latent hidden dimension.

In order to prevent data leakage, we used only the sites from the train split during hyperparameter tuning. The train sites were further divided into train and validation splits at a 0.8 / 0.2 ratio respectively. Our loss function was mean squared error (MSE) between the model output and target NEE. We used the AdamW optimizer [48] with a learning rate of  $8e-5$  and a batch size of 4096. A single warm-up epoch was performed followed by a cosine annealing learning rate schedule over 20 epochs, but all experiments converged between 6 and 13 epochs.

#### XGBoost

Hyperparameters were found by random search. We used the same train/test split as the EcoPerceiver experiment; the train set was used in a 5-fold cross validation framework with 50 iterations. Once hyperparameters were found, we retrained XGBoost on all training data before running inference on the test set. Table 4.2 details the parameterization of our final model.

Table 4.2 XGBoost Hyperparameters

Parameter	Value
learning_rate	0.1
alpha	0.1
gamma	0.4
lambda	0.0
max_depth	9
min_child_weight	9
n_estimators	150
subsample	0.7
scale_pos_weight	0.5
colsample_bytree	0.7
colsample_bylevel	0.8

#### 4.1.4 Reproducibility and Reliability

All experiments were designed with reproducibility in mind. Once optimal hyperparameters were found, we performed 10 experiments with each model in order to obtain a reliable measure of performance (inspired by [49]). Set seeds were provided to all frameworks utilizing RNG, and distributed dataloader workers were also seed-controlled to ensure full reproducibility of our results.

The seeds for our experiments were simple integer values (0, 10, 20, ... , 90) and were provided for the final training runs after hyperparameters had already been chosen.

#### 4.1.5 Metrics and Reporting

The most commonly used performance metric in DDCFM (and any form of hydrologic modelling) is the Nash-Sutcliffe Modelling Efficiency (NSE) [50], described with the following equation:

$$\text{NSE}(x, y) = 1 - \frac{\sum_i (y_i - x_i)^2}{\sum_i (y_i - \bar{y})^2} \quad (4.1)$$

where a value of 1 represents perfect correlation between  $x$  and  $y$ . A value of 0 represents the same performance as guessing the mean of  $y$ , and negative values indicate that the mean of  $y$  is a better predictor than  $x$ . NSE is more challenging to use directly as a loss function since it would require the dataloader to also provide the mean of the data for a given site or IGBP type. This is why MSE is used in our loss function. However, NSE is heavily correlated with MSE and is a better metric for judging model performance. We report both in our experiments.

When reporting results, data balance is also a concern. At first glance, the data appears very imbalanced with respect to ecosystem prevalence. CarbonSense contains 64 grasslands (GRA) sites, but only 1 deciduous needleleaf forest (DNF). While this is an extreme gap, ecosystems are more diverse than IGBP types can capture; grasslands in central North America will differ significantly from those in Europe or Asia. Still, it is prudent to separate results by IGBP type to give a better picture of model performance.

## 4.2 Core Experiments

The core experiments were designed to evaluate EcoPerceiver against the XGBoost baseline across a wide range of IGBP types. We ran 10 experiments with each model and examine the NSE and MSE on the test sets.

### 4.2.1 Quantitative Results

EcoPerceiver outperformed the XGBoost baseline across most IGBP types. Table 4.4 gives a summary of performance at-a-glance, and more detailed performance breakdowns are shown in Figures 4.2a (NSE) and 4.2b (RMSE).

**Partition Analysis** EcoPerceiver performed better than the XGBoost model in 12 out of 15 IGBP types. XGBoost performed better in permanent wetlands (WET), water bodies (WAT), and evergreen needleleaf forests (ENF) by a slim margin. In WAT, both models do substantially worse than simply guessing the mean, which makes those results challenging to interpret which model would perform better on other WAT sites. The other 2, ENF and WET, are not significantly better than EcoPerceiver’s performance and have mean NSE advantages of +0.0071 and +0.0393 respectively in favour of XGBoost.

**Low-Shot Performance** Besides WAT, EcoPerceiver did substantially better on zero- and one-shot tests. The NSE differential between the two models was +0.0486 on CVM, +0.1533 on DNF, and +0.3268 on SNO in favour of EcoPerceiver. High predictive power on out-of-distribution sites like this is especially important for researchers wishing to run inference on global data, where each grid cell is likely to be quite different from any of the training sites.

### 4.2.2 Qualitative Analysis

While error metrics are useful for assessing the aggregate performance of the models, it is always important to inspect the model outputs in comparison with the observed data. As an

Table 4.3 NSE and RMSE by model and IGBP type, aggregate mean across 10 seeds. Bold numbers indicate better performance. Pairwise t-test results are given for NSE values with 9 degrees of freedom.

IGBP	XGBoost		EcoPerceiver		t-test (NSE)	
	NSE	RMSE	NSE	RMSE	t-statistic	p-value
CRO	0.8066	3.2381	<b>0.8482</b>	<b>2.8677</b>	13.4689	0.0000
CSH	0.7510	1.5224	<b>0.7670</b>	<b>1.4709</b>	1.9947	0.0772
CVM	0.5277	5.5157	<b>0.5763</b>	<b>5.2236</b>	9.6586	0.0000
DBF	0.7250	4.0959	<b>0.7547</b>	<b>3.8678</b>	10.5993	0.0000
DNF	0.2803	4.0974	<b>0.4336</b>	<b>3.6322</b>	8.6338	0.0000
EBF	0.7966	4.6050	<b>0.8220</b>	<b>4.3070</b>	8.1990	0.0000
ENF	<b>0.7765</b>	<b>2.8141</b>	0.7694	2.8579	-2.3853	0.0409
GRA	0.7461	3.2487	<b>0.7967</b>	<b>2.9059</b>	13.7609	0.0000
MF	0.7559	3.8633	<b>0.7717</b>	<b>3.7361</b>	8.3540	0.0000
OSH	0.5451	1.8796	<b>0.6060</b>	<b>1.7475</b>	3.9356	0.0034
SAV	0.5802	1.6514	<b>0.7368</b>	<b>1.3070</b>	28.0814	0.0000
SNO	-0.0370	1.4291	<b>0.2898</b>	<b>1.1816</b>	16.3974	0.0000
WAT	<b>-11.0524</b>	<b>3.1838</b>	-14.4010	3.5802	-2.4809	0.0349
WET	<b>0.4530</b>	<b>2.2073</b>	0.4137	2.2830	-2.1005	0.0651
WSA	0.6132	2.5153	<b>0.6267</b>	<b>2.4706</b>	2.6798	0.0252

example, consider Figures 4.3a and 4.3b below. Both of these were randomly selected 4-day stretches of data from their respective sites. Both models appear able to model GF-Guy very well, but not CA-LP1, and this may be counterintuitive at first glance.

GF-Guy is an evergreen broadleaf forest in the tropics. This is not highly prevalent IGBP type in CarbonSense, but its carbon fluxes appear quite stable from day to day. We found that ecosystems with this interseasonal stability tend to be more easily modelled in our experiments, though this is not an easy metric to quantify. It should be noted that the y-axis has a much larger scale, so while the models appear close to the ground truth, they often have an error in excess of  $5 \mu\text{mol CO}_2 \text{ m}^{-2}$ . This highlights the importance of using NSE as an error metric - RMSE will unfairly punish highly active ecosystems like this due to higher natural variance in carbon fluxes.

CA-LP1 is an evergreen needleleaf forest in the temperate region, which is one of the best represented ecosystems in CarbonSense, yet both models struggle with it (despite low *absolute* error), especially in the winter. Reading into this site reveals it is a pine beetle-attacked forest [51]; disturbances like these can be challenging to model as we discussed in 2.1. This gives future work in DDCFM a vector for potential improvement.







Figure 4.2 NSE (a) and RMSE (b) scores of EcoPerceiver and XGBoost. Each chart represents 10 experiments with different seeds. Whiskers indicate approximately 1.5 x interquartile range.



(a)



(b)

Figure 4.3 Hourly data and model results for GF-Guy (a), an evergreen broadleaf forest station in French Guiana, and CA-LP1 (b), a pine beetle-attacked evergreen needleleaf forest in northern British Columbia.

### 4.2.3 Discussion

While the results are not a sweeping victory, EcoPerceiver performed consistently better than the XGBoost baseline. In most IGBP types it had substantially better NSE, and in the three IGBP types where XGBoost won out, it did not do so by much. The WAT sites are very far out of distribution from the rest of the data; they are typically EC towers mounted above lakes, and both models performed so poorly that we will exclude discussion of them hereafter.

It is ultimately unsurprising that EcoPerceiver performs better. XGBoost requires heavy compression to fit images into its input, which is destroying information. So long as that information has predictive power, it reduces XGBoost’s modelling potential. Conversely, EcoPerceiver takes in a range of input (32 hours of observations in our case). Assuming this information also has predictive power, it will increase EcoPerceiver’s modelling potential. This is the core of why deep MML is beneficial for DDCFM - it simply processes more information in a way tabular methods can not.

One possible explanation for EcoPerceiver’s reduced performance on ENF and WET is data splitting; once hyperparameters were obtained for XGBoost, it was able to train on the entirety of the train split, while EcoPerceiver still had to reserve 20% of the split for validation testing to measure convergence. Both ENF and WET had significant train set prevalence, so these represent IGBP types where XGBoost was most able to take advantage of additional data. Alternatively, we note that ENF and WET are the key ecosystems of the boreal biome, indicating XGBoost retains an advantage in that region. EcoPerceiver uses time-series information (within a strict context window) and better geospatial processing; to speculate, perhaps ENF and WET fluxes are less correlated with these attributes.

EcoPerceiver’s performance may also be dampened by imperfect hyperparameter selection. While XGBoost is lightweight enough for a virtually exhaustive parameter search with cross validation, deep models have significantly higher experimental overhead. Due to time and compute limitations, we were limited in how thoroughly we could explore model configurations for EcoPerceiver. We may be leaving performance on the table in this regard.

Our results also underline the importance of using NSE as the main metric for evaluation. Consider the models’ performance on open savannas (SAV). XGBoost had an RMSE of 1.6514 versus EcoPerceiver’s 1.3070. The magnitude of difference is small, and both values are significantly lower than the RMSE of many other IGBP types. But XGBoost had an NSE of 0.5802 while EcoPerceiver achieved 0.7368 which is a significant improvement. Different ecosystems have different variances in their carbon fluxes, and NSE accounts for this by dividing the performance by the variance of the target.



### 4.3 Ablation Studies

Here we present ablation studies on EcoPerceiver. The boxplots in this section instead represent scores across different IGBP types, since each experiment was run only once. Ranges are truncated to positive values due to the outlier IGBP (WET) making the charts illegible.

#### 4.3.1 Context Window Length


We first conduct an experiment to test the assumption that DDCFM benefits from ingesting data in a temporal context window. The results are shown in Figure 4.4. There is a clear performance advantage going from a context window of 8 hours to 16, and a small advantage going from 16 to 32. This reflects our findings from early hyperparameter tuning experiments. Going from 32 to 64 did not improve performance, but did increase our wall time and memory usage. We therefore used a context window of 32 for our main experiments.

#### 4.3.2 Core Architectural Decisions


Figure 4.5 shows the effects of our core architectural decisions on EcoPerceiver’s performance. We found causal masking to have the smallest impact on performance. Observational dropout and Fourier feature encoding both noticeably improved model performance compared to a model without them. We also tested our model trained on a version of CarbonSense with the low-quality gapfilled values included. We found our model performed better when trained on a smaller amount of high quality data, as the gapfilled model had lower NSE and higher RMSE.

There was a significant improvement of performance with the inclusion of image data (compared with the "No Image Data" model in Figure 4.5). This matched our expectation that providing more information to the model via satellite data would improve the conditioning on the meteorological data and lead to better inference power. Finally, we tested against a vanilla transformer encoder which was fed the same data, but did not use any windowed cross attention operations. This model performed the worst of all our experiments, indicating that windowed cross attention bring significant performance gains for DDCFM.

The quantitative results of these ablation studies are given in table 4.4. Note that we use truncated mean and standard deviations for these tests, since they are robust statistical measures for data with significant outliers (mainly WET ecosystem). Median and interquartile range (iqr) are both given as well.



(a)



(b)

Figure 4.4 Effects of context window length on NSE (a) and RMSE (b) of test site inference for EcoPerceiver.

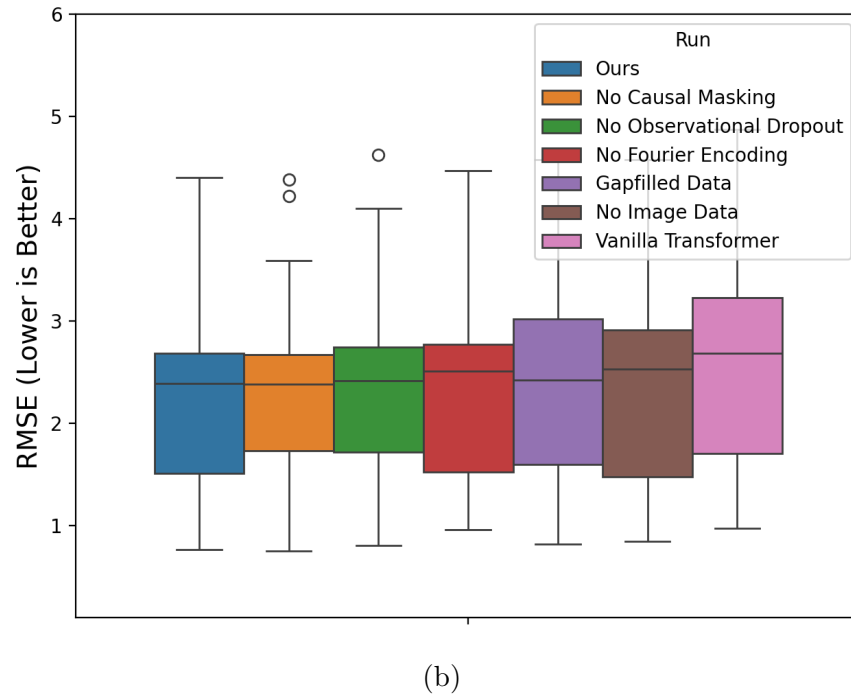
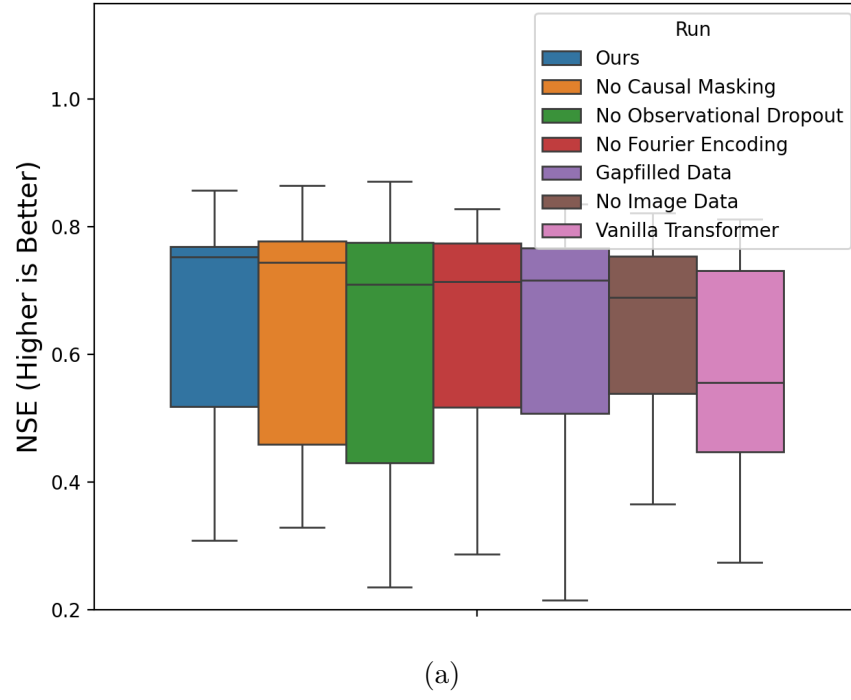


Figure 4.5 Effects of various architectural choices on NSE (a) and RMSE (b) of test site inference for EcoPerceiver.

Table 4.4 NSE and RMSE by model ablation.  $t\_mean$  and  $t\_std$  represent truncated means and standard deviations to account for the outlier biome ("WAT").

Model	NSE				RMSE			
	t_mean	t_std	median	iqr	t_mean	t_std	median	iqr
Vanilla Transformer	0.557	0.249	0.556	0.284	2.709	1.165	2.648	1.603
No Image Data	0.635	0.186	0.689	0.215	2.515	1.134	2.478	1.582
Gapfilled Data	0.635	0.199	0.717	0.259	2.474	1.077	2.376	1.492
No Fourier Encoding	0.624	0.254	0.715	0.257	2.456	1.043	2.481	1.098
No Obs. Dropout	0.634	0.196	0.710	0.345	2.468	1.091	2.339	1.105
No Causal Masking	0.654	0.171	0.744	0.318	2.428	1.077	<b>2.252</b>	0.836
Ours	<b>0.669</b>	0.166	<b>0.753</b>	0.251	<b>2.382</b>	1.055	2.313	1.284

### 4.3.3 Ecosystem-Specific DDCFM

Many research teams study DDCFM in the context of specific regions and IGBP types [7, 8, 11, 20]. This use case is one of the reasons for CarbonSense’s partitioned structure. As a proof of concept, we ran a single experiment with EcoPerceiver where we use the same configuration as our main experiment, but only include the DBF sites. We then compared the test set performance against our main model. Table 4.5 shows the results - our model trained on multiple IGBP types had notably better performance despite a similar convergence time.

Table 4.5 EcoPerceiver DBF performance when trained on DBF data vs trained on all sites

Only DBF		All Sites	
NSE	RMSE	NSE	RMSE
0.7405	3.9782	<b>0.7532</b>	<b>3.8806</b>

## CHAPTER 5 CONCLUSION

### 5.1 Summary of Works

Combating climate change will require coordinated efforts on many fronts. Understanding the carbon dynamics of our biosphere will provide crucial information to decision-makers about the capacity of ecosystems to sequester carbon and how that capacity may change with increasing anthropogenic pressures. In this work we aimed to improve our modelling capabilities by applying recent advances in multimodal deep learning to DDCFM, and provide tools to lower the barrier to entry for deep learning researchers to contribute to this field.

We began by creating CarbonSense, the first openly available multimodal dataset for DDCFM. CarbonSense is an aggregation of EC station data and satellite data from 385 sites around the world, encompassing a wide range of geographic locations and ecosystems. It is meant to act as a common dataset for algorithmic development in this field, as future researchers can directly compare model performance without accounting for differences in training data. CarbonSense is partitioned by EC site so researchers may use it for region- or ecosystem-specific modelling, and we provide the raw data and pipeline to allow others to build upon our work.

We then designed EcoPerceiver, a deep multimodal model for DDCFM. EcoPerceiver leverages the power of attentive neural architectures to ingest meteorological and satellite image data without losing information via compression. We trained our model on the CarbonSense dataset and compared its performance against an XGBoost baseline. Our experiments show that EcoPerceiver achieves SOTA performance across a wide range of ecosystem types despite using less of the overall data for training. We provide EcoPerceiver’s code and weights as a new baseline for researchers looking to build new models for DDCFM.

### 5.2 Limitations

CarbonSense has a data imbalance in not only IGBP types, but geographic location. Africa, Central Asia, and South America are all underrepresented. While these areas contain many EC stations, most do not have readily available data in ONEFlux format, presenting a barrier to their inclusion in CarbonSense. Researchers should be aware of the consequences of developing models with imbalanced data, including poor performance in underrepresented areas.

We also acknowledge limitations with the design and testing of our models. Time and compute constraints limited the amount of hyperparameter tuning we were able to achieve with EcoPerceiver; this may be causing underperformance, or it may not. While we tried to recreate the XGBoost model from recent works, it will never be a truly apples-to-apples comparison; without access to their data pipeline and training code it remains likely that our dataset, model selection, and experiment design had subtle but significant differences. It would be simply impossible to perfectly replicate previous works without extensive collaboration.

Finally, we were unable to organize the necessary data and compute resources for global upscaling with either of our models. This would have been feasible with more time and computational resources. Upscaling is one of the most important tasks in DDCFM; not only does it provide real, usable data, but it also allows domain experts to qualitatively assess the validity of the model output. This is especially true for areas not present in the data such as Sub-Saharan Africa, where expert opinion can act as a stand-in for quantitative metrics.

### 5.3 Future Research

Our work leaves much to be explored in both dataset and model development. CarbonSense can be expanded as more EC station data is incorporated into regional network releases. Additional geospatial data could be added in the form of global soil products or higher resolution satellite imagery. It may also be beneficial to support other fluxes such as methane and nitrous oxide; this would widen the academic uses for CarbonSense and encourage wider adoption.

Future models may work to address the shortcomings of EcoPerceiver in the boreal biome, or incorporate more sophisticated fusion techniques like the use of convolutional layers in image ingestion. We used the simplified regression technique of running a linear layer over the latent space, but more recent MML models like Perceiver IO [52] use cross-attention to dynamically modify the output variables at runtime. This could result in a model which can predict multiple fluxes simultaneously and benefit from weight sharing.

Ultimately, while iterating on EcoPerceiver presents a promising direction, our hope is that researchers will build upon our contributions and develop innovative solutions that go beyond our initial framework. The fight against climate change demands bold and creative approaches. We believe that deep learning has immense potential to help us understand and protect our environment, but this potential will only be realized if the research community commits to exploring and advancing these applications.

## REFERENCES

- [1] A. Vaswani *et al.*, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [2] A. Jaegle *et al.*, “Perceiver: General perception with iterative attention,” in *Proceedings of the 38th International Conference on Machine Learning*, vol. 139, 18–24 Jul 2021, pp. 4651–4664.
- [3] C. Seiler *et al.*, “Are terrestrial biosphere models fit for simulating the global land carbon sink?” *Journal of Advances in Modeling Earth Systems*, vol. 14, 2021.
- [4] D. Baldocchi *et al.*, “Ameriflux: Its impact on our understanding of the ‘breathing of the biosphere’, after 25 years,” *Agricultural and Forest Meteorology*, vol. 348, p. 109929, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168192324000443>
- [5] G. Tramontana *et al.*, “Predicting carbon dioxide and energy fluxes across global fluxnet sites with regression algorithms,” *Biogeosciences*, vol. 13, no. 14, pp. 4291–4313, 2016. [Online]. Available: <https://bg.copernicus.org/articles/13/4291/2016/>
- [6] J. A. Nelson *et al.*, “X-base: the first terrestrial carbon and water flux products from an extended data-driven scaling framework, fluxcom-x,” *EGUsphere*, vol. 2024, pp. 1–51, 2024. [Online]. Available: <https://egusphere.copernicus.org/preprints/2024/egusphere-2024-165/>
- [7] O. Peltola *et al.*, “Monthly gridded data product of northern wetland methane emissions based on upscaling eddy covariance observations,” *Earth System Science Data*, vol. 11, no. 3, pp. 1263–1289, 2019. [Online]. Available: <https://essd.copernicus.org/articles/11/1263/2019/>
- [8] G. McNicol *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*, vol. 4, no. 5, p. e2023AV000956, 2023, e2023AV000956 2023AV000956. [Online]. Available: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2023AV000956>

- [9] A.-M. Virkkala *et al.*, “An increasing arctic-boreal co2 sink despite strong regional sources,” *bioRxiv*, 2024. [Online]. Available: <https://www.biorxiv.org/content/early/2024/02/12/2024.02.09.579581>
- [10] —, “Statistical upscaling of ecosystem co2 fluxes across the terrestrial tundra and boreal domain: Regional patterns and uncertainties,” *Global Change Biology*, vol. 27, no. 17, pp. 4040–4059, 2021. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/gcb.15659>
- [11] C. Zhang *et al.*, “Using an object-based machine learning ensemble approach to upscale evapotranspiration measured from eddy covariance towers in a subtropical wetland,” *Science of The Total Environment*, vol. 831, p. 154969, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0048969722020629>
- [12] P. Xu, X. Zhu, and D. A. Clifton, “Multimodal learning with transformers: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 10, pp. 12 113–12 132, 2023.
- [13] J. Li *et al.*, “Deep learning in multimodal remote sensing data fusion: A comprehensive review,” *International Journal of Applied Earth Observation and Geoinformation*, vol. 112, p. 102926, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1569843222001248>
- [14] G. Michel *et al.*, “Evaluating state-of-the-art process-based and data-driven models in simulating co2 fluxes and their relationship with climate in western european temperate forests,” *EGUsphere*, vol. 2024, pp. 1–42, 2024. [Online]. Available: <https://egusphere.copernicus.org/preprints/2024/egusphere-2024-1758/>
- [15] D. D. Baldocchi, “How eddy covariance flux measurements have contributed to our understanding of global change biology,” *Global Change Biology*, vol. 26, no. 1, pp. 242–260, 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1111/gcb.14807>
- [16] M. Khalil *et al.*, “Machine learning for modeling n2o emissions from wastewater treatment plants: Aligning model performance, complexity, and interpretability,” *Water Research*, vol. 245, p. 120667, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0043135423011077>
- [17] M. Reichstein *et al.*, “On the separation of net ecosystem exchange into assimilation and ecosystem respiration: Review and improved algorithm,” *Global Change Biology*, vol. 11, pp. 1424 – 1439, 07 2005.



- [18] G. Pastorello *et al.*, “The FLUXNET2015 dataset and the ONEFlux processing pipeline for eddy covariance data,” *Scientific Data*, vol. 7, no. 1, p. 225, Jul 2020. [Online]. Available: <https://doi.org/10.1038/s41597-020-0534-3>
- [19] H. Hersbach *et al.*, “The era5 global reanalysis,” *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999–2049, 2020. [Online]. Available: <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803>
- [20] D. PAPAŁE and R. VALENTINI, “A new assessment of european forests carbon exchanges by eddy fluxes and artificial neural network spatialization,” *Global Change Biology*, vol. 9, no. 4, pp. 525–535, 2003. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2486.2003.00609.x>
- [21] M. Jung *et al.*, “Global patterns of land-atmosphere fluxes of carbon dioxide, latent heat, and sensible heat derived from eddy covariance, satellite, and meteorological observations,” *Journal of Geophysical Research: Biogeosciences*, vol. 116, no. G3, 2011. [Online]. Available: <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2010JG001566>
- [22] S. Walther *et al.*, “Technical note: A view from space on global flux towers by modis and landsat: the fluxneteo data set,” *Biogeosciences*, vol. 19, no. 11, pp. 2805–2840, 2022. [Online]. Available: <https://bg.copernicus.org/articles/19/2805/2022/>
- [23] D. Sulla-Menashe and M. A. Friedl, “User Guide to Collection 6 MODIS Land Cover (MCD12Q1 and MCD12C1) Product,” 05 2018. [Online]. Available: <https://www.umb.edu/spectralmass/v006/mcd43a4-nbar-product/>
- [24] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 785–794. [Online]. Available: <https://doi.org/10.1145/2939672.2939785>
- [25] S. Islam *et al.*, “A comprehensive survey on applications of transformers for deep learning tasks,” *Expert Systems with Applications*, vol. 241, p. 122666, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0957417423031688>
- [26] A. Radford *et al.*, “Improving language understanding by generative pre-training,” 2018.
- [27] A. Dosovitskiy *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.

- [28] Y. Zhang *et al.*, “Meta-transformer: A unified framework for multimodal learning,” *arXiv preprint arXiv:2307.10802*, 2023.
- [29] X. Han *et al.*, “A survey of transformer-based multimodal pre-trained modals,” *Neurocomputing*, vol. 515, pp. 89–106, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925231222012346>
- [30] B. P. Yuhas, M. H. Goldstein, and T. J. Sejnowski, “Integration of acoustic and visual speech signals using neural networks,” *IEEE Communications Magazine*, vol. 27, no. 11, pp. 65–71, 1989.
- [31] C. Chen, D. Han, and C.-C. Chang, “Mpcct: Multimodal vision-language learning paradigm with context-based compact transformer,” *Pattern Recognition*, vol. 147, p. 110084, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0031320323007811>
- [32] F. Wang *et al.*, “Tedt: Transformer-based encoding–decoding translation network for multimodal sentiment analysis,” *Cognitive Computation*, vol. 15, no. 1, pp. 289–303, Jan 2023. [Online]. Available: <https://doi.org/10.1007/s12559-022-10073-9>
- [33] H.-Y. Zhou *et al.*, “A transformer-based representation-learning model with unified processing of multimodal input for clinical diagnostics,” *Nature Biomedical Engineering*, vol. 7, no. 6, pp. 743–755, Jun 2023. [Online]. Available: <https://doi.org/10.1038/s41551-023-01045-x>
- [34] J. Yao *et al.*, “Extended vision transformer (exvit) for land use and land cover classification: A multimodal deep learning framework,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–15, 2023.
- [35] V. Benson *et al.*, “Multi-modal Learning for Geospatial Vegetation Forecasting,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 27 788–27 799.
- [36] J. Deng *et al.*, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
- [37] C. Schuhmann *et al.*, “Laion-5b: An open large-scale dataset for training next generation image-text models,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 25 278–25 294, 2022.

- [38] ICOS RI *et al.*, “ICOS Atmosphere Release 2023-1 of Level 2 Greenhouse Gas Mole Fractions of CO<sub>2</sub>, CH<sub>4</sub>, N<sub>2</sub>O, CO, meteorology and 14CO<sub>2</sub>, and flask samples analysed for CO<sub>2</sub>, CH<sub>4</sub>, N<sub>2</sub>O, CO, H<sub>2</sub> and SF<sub>6</sub>.” 2023. [Online]. Available: <https://www.icos-cp.eu/data-products/atmosphere-release>
- [39] Warm Winter 2020 Team and ICOS Ecosystem Thematic Centre, “Warm Winter 2020 ecosystem eddy covariance flux product for 73 stations in FLUXNET-Archive format—release 2022-1 (Version 1.0).” *ICOS Carbon Portal*, 2022. [Online]. Available: <https://www.icos-cp.eu/data-products/2G60-ZHAK>
- [40] H. Chu *et al.*, “Ameriflux base data pipeline to support network growth and data sharing,” *Scientific Data*, vol. 10, no. 1, p. 614, Sep 2023. [Online]. Available: <https://doi.org/10.1038/s41597-023-02531-2>
- [41] C. Schaaf and Z. Wang, “MCD43A4 MODIS/Terra+Aqua BRDF/Albedo Nadir BRDF Adjusted Ref Daily L3 Global - 500m V006,” 2015b. [Online]. Available: <https://www.umb.edu/spectralmass/v006/mcd43a4-nbar-product/>
- [42] —, “MCD43A2 MODIS/Terra+Aqua BRDF/Albedo Quality Daily L3 Global - 500m V006,” 2015a. [Online]. Available: <https://www.umb.edu/spectralmass/v006/mcd43a2-albedo-product/>
- [43] N. Gorelick *et al.*, “Google earth engine: Planetary-scale geospatial analysis for everyone,” *Remote Sensing of Environment*, 2017. [Online]. Available: <https://doi.org/10.1016/j.rse.2017.06.031>
- [44] N. Rahaman *et al.*, “On the spectral bias of neural networks,” in *International conference on machine learning*. PMLR, 2019, pp. 5301–5310.
- [45] M. Tancik *et al.*, “Fourier features let networks learn high frequency functions in low dimensional domains,” *Advances in neural information processing systems*, vol. 33, pp. 7537–7547, 2020.
- [46] B. Mildenhall *et al.*, “Nerf: Representing scenes as neural radiance fields for view synthesis,” in *ECCV*, 2020.
- [47] Z. Liu *et al.*, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 10 012–10 022.

- [48] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *7th International Conference on Learning Representations, (ICLR) 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>
- [49] R. Agarwal *et al.*, “Deep reinforcement learning at the edge of the statistical precipice,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [50] R. McCuen, Z. Knight, and A. Cutter, “Evaluation of the nash–sutcliffe efficiency index,” *Journal of Hydrologic Engineering - J HYDROL ENG*, vol. 11, 11 2006.
- [51] T. A. Black, “AmeriFlux FLUXNET-1F CA-LP1 British Columbia - Mountain pine beetle-attacked lodgepole pine stand , Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [52] A. Jaegle *et al.*, “Perceiver io: A general architecture for structured inputs & outputs,” *arXiv preprint arXiv:2107.14795*, 2021.
- [53] C. Wagner-Riddle, “AmeriFlux FLUXNET-1F CA-ER1 Elora Research Station, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [54] B. Amiro, “AmeriFlux FLUXNET-1F CA-MA1 Manitoba Agricultural Site 1, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [55] —, “AmeriFlux FLUXNET-1F CA-MA2 Manitoba Agricultural Site 2, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [56] D. Billesbach *et al.*, “AmeriFlux FLUXNET-1F US-A74 ARM SGP milo field, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [57] S. Biraud *et al.*, “AmeriFlux FLUXNET-1F US-ARM ARM Southern Great Plains site- Lamont, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [58] C. Rey-Sanchez *et al.*, “AmeriFlux FLUXNET-1F US-Bi1 Bouldin Island Alfalfa, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [59] —, “AmeriFlux FLUXNET-1F US-Bi2 Bouldin Island corn, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [60] C. L. Phillips and D. Huggins, “AmeriFlux FLUXNET-1F US-CF1 CAF-LTAR Cook East, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [61] D. Huggins, “AmeriFlux FLUXNET-1F US-CF2 CAF-LTAR Cook West, Ver. 3-5,” *AmeriFlux AMP*, 2022.

- [62] —, “AmeriFlux FLUXNET-1F US-CF3 CAF-LTAR Boyd North, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [63] —, “AmeriFlux FLUXNET-1F US-CF4 CAF-LTAR Boyd South, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [64] J. Chen and H. Chu, “AmeriFlux FLUXNET-1F US-CRT Curtice Walter-Berger cropland, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [65] A. Desai, “AmeriFlux FLUXNET-1F US-CS1 Central Sands Irrigated Agricultural Field, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [66] —, “AmeriFlux FLUXNET-1F US-CS3 Central Sands Irrigated Agricultural Field, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [67] —, “AmeriFlux FLUXNET-1F US-CS4 Central Sands Irrigated Agricultural Field, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [68] A. Duff and A. Desai, “AmeriFlux FLUXNET-1F US-DFC US Dairy Forage Research Center, Prairie du Sac, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [69] M. R. Schuppenhauer *et al.*, “AmeriFlux FLUXNET-1F US-DS3 Staten Rice 1, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [70] S. Fares, “AmeriFlux FLUXNET-1F US-Lin Lindcove Orange Orchard, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [71] A. Schreiner-McGraw, “AmeriFlux FLUXNET-1F US-Mo1 LTAR CMRB Field 1 (CMRB ASP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [72] —, “AmeriFlux FLUXNET-1F US-Mo3 LTAR CMRB Field 3 (CMRB BAU), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [73] A. Suyker, “AmeriFlux FLUXNET-1F US-Ne1 Mead - irrigated continuous maize site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [74] M. R. Schuppenhauer, S. C. Biraud, and S. Chan, “AmeriFlux FLUXNET-1F US-RGA Arkansas Corn Farm, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [75] M. Schuppenhauer, S. C. Biraud, and S. Chan, “AmeriFlux FLUXNET-1F US-RGB Butte County Rice Farm, Ver. 3-5,” *AmeriFlux AMP*, 2023.

- [76] M. R. Schuppenhauer, S. C. Biraud, and S. Chan, “AmeriFlux FLUXNET-1F US-RGo Glenn County Organic Rice Farm, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [77] J. Baker, T. Griffis, and T. Griffis, “AmeriFlux FLUXNET-1F US-Ro1 Rosemount-G21, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [78] J. Baker and T. Griffis, “AmeriFlux FLUXNET-1F US-Ro2 Rosemount- C7, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [79] —, “AmeriFlux FLUXNET-1F US-Ro5 Rosemount I18 South, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [80] —, “AmeriFlux FLUXNET-1F US-Ro6 Rosemount I18 North, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [81] C. Sturtevant, J. Verfaillie, and D. Baldocchi, “AmeriFlux FLUXNET-1F US-Tw2 Twitchell Corn, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [82] S. D. Chamberlain *et al.*, “AmeriFlux FLUXNET-1F US-Tw3 Twitchell Alfalfa, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [83] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xSL NEON North Sterling, CO (STER), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [84] B. Drake *et al.*, “AmeriFlux FLUXNET-1F US-KS2 Kennedy Space Center (scrub oak), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [85] G. Flerchinger, “AmeriFlux FLUXNET-1F US-Rls RCEW Low Sagebrush, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [86] —, “AmeriFlux FLUXNET-1F US-Rms RCEW Mountain Big Sagebrush, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [87] G. Flerchinger and M. L. Reba, “AmeriFlux FLUXNET-1F US-Rwe RCEW Reynolds Mountain East, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [88] G. Flerchinger, “AmeriFlux FLUXNET-1F US-Rwf RCEW Upper Sheep Prescribed Fire, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [89] S. Goslee, “AmeriFlux FLUXNET-1F US-HWB USDA ARS Pasture Sytems and Watershed Management Research Unit- Hawbecker Site, Ver. 3-5,” *AmeriFlux AMP*, 2022.

- [90] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xDS NEON Disney Wilderness Preserve (DSNY), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [91] R. Staebler, “AmeriFlux FLUXNET-1F CA-Cbo Ontario - Mixed Deciduous, Borden Forest Site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [92] M. A. Arain, “AmeriFlux FLUXNET-1F CA-TPD Ontario - Turkey Point Mature Deciduous, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [93] E. A. Yopez and J. Garatuza, “AmeriFlux FLUXNET-1F MX-Tes Tesopaco, secondary tropical dry forest, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [94] A. Richardson and D. Hollinger, “AmeriFlux FLUXNET-1F US-Bar Bartlett Experimental Forest, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [95] J. W. Munger, “AmeriFlux FLUXNET-1F US-Ha1 Harvard Forest EMS Tower (HFR1), Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [96] K. Novick and R. Phillips, “AmeriFlux FLUXNET-1F US-MMS Morgan Monroe State Forest, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [97] J. Wood and L. Gu, “AmeriFlux FLUXNET-1F US-MOz Missouri Ozark Site, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [98] J. Chen, H. Chu, and A. Noormets, “AmeriFlux FLUXNET-1F US-Oho Oak Openings, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [99] M. Ueyama, H. Iwata, and Y. Harazono, “AmeriFlux FLUXNET-1F US-Rpf Poker Flat Research Range: Succession from fire scar to deciduous forest, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [100] C. Gough, G. Bohrer, and P. Curtis, “AmeriFlux FLUXNET-1F US-UMB Univ. of Mich. Biological Station, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [101] —, “AmeriFlux FLUXNET-1F US-UMd UMBS Disturbance, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [102] J. Chen, “AmeriFlux FLUXNET-1F US-Wi1 Intermediate hardwood (IHW), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [103] —, “AmeriFlux FLUXNET-1F US-Wi3 Mature hardwood (MHW), Ver. 3-5,” *AmeriFlux AMP*, 2023.

- [104] —, “AmeriFlux FLUXNET-1F US-Wi8 Young hardwood clearcut (YHW), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [105] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xBL NEON Blandy Experimental Farm (BLAN), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [106] —, “AmeriFlux FLUXNET-1F US-xBR NEON Bartlett Experimental Forest (BART), Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [107] —, “AmeriFlux FLUXNET-1F US-xGR NEON Great Smoky Mountains National Park, Twin Creeks (GRSM), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [108] —, “AmeriFlux FLUXNET-1F US-xHA NEON Harvard Forest (HARV), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [109] —, “AmeriFlux FLUXNET-1F US-xML NEON Mountain Lake Biological Station (MLBS), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [110] —, “AmeriFlux FLUXNET-1F US-xSC NEON Smithsonian Conservation Biology Institute (SCBI), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [111] —, “AmeriFlux FLUXNET-1F US-xSE NEON Smithsonian Environmental Research Center (SERC), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [112] —, “AmeriFlux FLUXNET-1F US-xST NEON Steigerwaldt Land Services (STEI), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [113] —, “AmeriFlux FLUXNET-1F US-xTR NEON Treehaven (TREE), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [114] —, “AmeriFlux FLUXNET-1F US-xUK NEON The University of Kansas Field Station (UKFS), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [115] A. Antonino, “AmeriFlux FLUXNET-1F BR-CST Caatinga Serra Talhada, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [116] T. A. Black, “AmeriFlux FLUXNET-1F CA-Ca1 British Columbia - 1949 Douglas-fir stand, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [117] —, “AmeriFlux FLUXNET-1F CA-Ca2 British Columbia - Clearcut Douglas-fir stand (harvested winter 1999/2000), Ver. 3-5,” *AmeriFlux AMP*, 2023.



- [118] M. Goulden, “AmeriFlux FLUXNET-1F CA-NS1 UCI-1850 burn site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [119] —, “AmeriFlux FLUXNET-1F CA-NS2 UCI-1930 burn site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [120] —, “AmeriFlux FLUXNET-1F CA-NS3 UCI-1964 burn site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [121] —, “AmeriFlux FLUXNET-1F CA-NS4 UCI-1964 burn site wet, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [122] —, “AmeriFlux FLUXNET-1F CA-NS5 UCI-1981 burn site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [123] H. A. Margolis, “AmeriFlux FLUXNET-1F CA-Qfo Quebec - Eastern Boreal, Mature Black Spruce, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [124] B. Amiro, “AmeriFlux FLUXNET-1F CA-SF1 Saskatchewan - Western Boreal, forest burned in 1977, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [125] —, “AmeriFlux FLUXNET-1F CA-SF2 Saskatchewan - Western Boreal, forest burned in 1989, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [126] M. A. Arain, “AmeriFlux FLUXNET-1F CA-TP1 Ontario - Turkey Point 2002 Plantation White Pine, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [127] —, “AmeriFlux FLUXNET-1F CA-TP3 Ontario - Turkey Point 1974 Plantation White Pine, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [128] E. Euskirchen, “AmeriFlux FLUXNET-1F US-BZS Bonanza Creek Black Spruce, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [129] A. Desai, “AmeriFlux FLUXNET-1F US-CS2 Tri county school Pine Forest, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [130] S. Dore and T. Kolb, “AmeriFlux FLUXNET-1F US-Fmf Flagstaff - Managed Forest, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [131] —, “AmeriFlux FLUXNET-1F US-Fuf Flagstaff - Unmanaged Forest, Ver. 3-5,” *AmeriFlux AMP*, 2023.

- [132] J. Frank and B. Massman, “AmeriFlux FLUXNET-1F US-GLE GLEES, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [133] J. D. Forsythe, M. A. Kline, and T. L. O’Halloran, “AmeriFlux FLUXNET-1F US-HB2 Hobcaw Barony Mature Longleaf Pine, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [134] —, “AmeriFlux FLUXNET-1F US-HB3 Hobcaw Barony Longleaf Pine Restoration, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [135] D. Hollinger, “AmeriFlux FLUXNET-1F US-Ho2 Howland Forest (west tower), Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [136] B. Drake *et al.*, “AmeriFlux FLUXNET-1F US-KS1 Kennedy Space Center (slash pine), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [137] B. Law, “AmeriFlux FLUXNET-1F US-Me1 Metolius - Eyerly burn, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [138] —, “AmeriFlux FLUXNET-1F US-Me2 Metolius mature ponderosa pine, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [139] —, “AmeriFlux FLUXNET-1F US-Me3 Metolius-second young aged pine, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [140] —, “AmeriFlux FLUXNET-1F US-Me6 Metolius Young Pine Burn, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [141] A. Noormets *et al.*, “AmeriFlux FLUXNET-1F US-NC1 NC Clearcut, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [142] —, “AmeriFlux FLUXNET-1F US-NC3 NC Clearcut 3, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [143] P. D. Blanken *et al.*, “AmeriFlux FLUXNET-1F US-NR1 Niwot Ridge Forest (LTER NWT1), Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [144] M. Litvak, “AmeriFlux FLUXNET-1F US-Vcm Valles Caldera Mixed Conifer, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [145] —, “AmeriFlux FLUXNET-1F US-Vcp Valles Caldera Ponderosa Pine, Ver. 3-5,” *AmeriFlux AMP*, 2023.

- [146] J. Chen, “AmeriFlux FLUXNET-1F US-Wi0 Young red pine (YRP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [147] —, “AmeriFlux FLUXNET-1F US-Wi4 Mature red pine (MRP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [148] —, “AmeriFlux FLUXNET-1F US-Wi5 Mixed young jack pine (MYJP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [149] —, “AmeriFlux FLUXNET-1F US-Wi9 Young Jack pine (YJP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [150] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xAB NEON Abby Road (ABBY), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [151] —, “AmeriFlux FLUXNET-1F US-xBN NEON Caribou Creek - Poker Flats Watershed (BONA), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [152] —, “AmeriFlux FLUXNET-1F US-xDJ NEON Delta Junction (DEJU), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [153] —, “AmeriFlux FLUXNET-1F US-xJE NEON Jones Ecological Research Center (JERC), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [154] —, “AmeriFlux FLUXNET-1F US-xRM NEON Rocky Mountain National Park, CASTNET (RMNP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [155] —, “AmeriFlux FLUXNET-1F US-xSB NEON Ordway-Swisher Biological Station (OSBS), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [156] —, “AmeriFlux FLUXNET-1F US-xTA NEON Talladega National Forest (TALL), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [157] —, “AmeriFlux FLUXNET-1F US-xYE NEON Yellowstone Northern Range (Frog Rock) (YELL), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [158] B. Amiro, “AmeriFlux FLUXNET-1F CA-MA3 Manitoba Agricultural Site 3, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [159] D. Billesbach *et al.*, “AmeriFlux FLUXNET-1F US-A32 ARM-SGP Medford hay pasture, Ver. 3-5,” *AmeriFlux AMP*, 2022.

- [160] D. Billesbach, J. Bradford, and M. Torn, “AmeriFlux FLUXNET-1F US-AR1 ARM USDA UNL OSU Woodward Switchgrass 1, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [161] —, “AmeriFlux FLUXNET-1F US-AR2 ARM USDA UNL OSU Woodward Switchgrass 2, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [162] M. Torn, “AmeriFlux FLUXNET-1F US-ARb ARM Southern Great Plains burn site-Lamont, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [163] —, “AmeriFlux FLUXNET-1F US-ARc ARM Southern Great Plains control site-Lamont, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [164] K. Novick, “AmeriFlux FLUXNET-1F US-BRG Bayles Road Grassland Tower, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [165] D. Bowling, “AmeriFlux FLUXNET-1F US-Cop Corral Pocket, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [166] H. Liu, M. Huang, and X. Chen, “AmeriFlux FLUXNET-1F US-Hn2 Hanford 100H grassland, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [167] N. Brunsell, “AmeriFlux FLUXNET-1F US-KFS Kansas Field Station, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [168] —, “AmeriFlux FLUXNET-1F US-KLS Kansas Land Institute, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [169] —, “AmeriFlux FLUXNET-1F US-Kon Konza Prairie LTER (KNZ), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [170] A. Schreiner-McGraw, “AmeriFlux FLUXNET-1F US-Mo2 LTAR CMRB Tucker Prairie (CMRB TP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [171] M. Torn and S. Dengel, “AmeriFlux FLUXNET-1F US-NGC NGEE Arctic Council, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [172] M. L. Silveira and R. Bracho, “AmeriFlux FLUXNET-1F US-ONA Florida pine flatwoods, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [173] J. Baker and T. Griffis, “AmeriFlux FLUXNET-1F US-Ro4 Rosemount Prairie, Ver. 3-5,” *AmeriFlux AMP*, 2022.

- [174] R. Scott, “AmeriFlux FLUXNET-1F US-SRG Santa Rita Grassland, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [175] M. Litvak, “AmeriFlux FLUXNET-1F US-Seg Sevilleta grassland, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [176] R. Shortt *et al.*, “AmeriFlux FLUXNET-1F US-Sne Sherman Island Restored Wetland, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [177] K. Kusak *et al.*, “AmeriFlux FLUXNET-1F US-Snf Sherman Barn, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [178] S. Ma *et al.*, “AmeriFlux FLUXNET-1F US-Var Vaira Ranch- Ione, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [179] R. Scott, “AmeriFlux FLUXNET-1F US-Wkg Walnut Gulch Kendall Grasslands, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [180] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xAE NEON Klemme Range Research Station (OAES), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [181] —, “AmeriFlux FLUXNET-1F US-xCL NEON LBJ National Grassland (CLBJ), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [182] —, “AmeriFlux FLUXNET-1F US-xCP NEON Central Plains Experimental Range (CPER), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [183] —, “AmeriFlux FLUXNET-1F US-xDC NEON Dakota Coteau Field School (DCFS), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [184] —, “AmeriFlux FLUXNET-1F US-xKA NEON Konza Prairie Biological Station - Relocatable (KONA), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [185] —, “AmeriFlux FLUXNET-1F US-xKZ NEON Konza Prairie Biological Station (KONZ), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [186] —, “AmeriFlux FLUXNET-1F US-xNG NEON Northern Great Plains Research Laboratory (NOGP), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [187] —, “AmeriFlux FLUXNET-1F US-xWD NEON Woodworth (WOOD), Ver. 3-5,” *AmeriFlux AMP*, 2023.

- [188] H. McCaughey, “AmeriFlux FLUXNET-1F CA-Gro Ontario - Groundhog River, Boreal Mixedwood Forest, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [189] A. Desai, “AmeriFlux FLUXNET-1F US-Syv Sylvania Wilderness Area, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [190] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xDL NEON Dead Lake (DELA), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [191] —, “AmeriFlux FLUXNET-1F US-xUN NEON University of Notre Dame Environmental Research Center (UNDE), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [192] M. Goulden, “AmeriFlux FLUXNET-1F CA-NS6 UCI-1989 burn site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [193] R. Bracho *et al.*, “AmeriFlux FLUXNET-1F US-EML Eight Mile Lake Permafrost thaw gradient, Healy Alaska., Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [194] M. Ueyama, H. Iwata, and Y. Harazono, “AmeriFlux FLUXNET-1F US-Fcr Cascaden Ridge Fire Scar, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [195] H. Liu, M. Huang, and X. Chen, “AmeriFlux FLUXNET-1F US-Hn3 Hanford 100H sagebrush, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [196] E. Euskirchen, G. Shaver, and S. Bret-Harte, “AmeriFlux FLUXNET-1F US-IC<sub>h</sub> Imnavait Creek Watershed Heath Tundra, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [197] —, “AmeriFlux FLUXNET-1F US-IC<sub>t</sub> Imnavait Creek Watershed Tussock Tundra, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [198] C. Tweedie, “AmeriFlux FLUXNET-1F US-Jo1 Jornada Experimental Range Bajada Site, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [199] E. R. Vivoni and E. R. Perez-Ruiz, “AmeriFlux FLUXNET-1F US-Jo2 Jornada Experimental Range Mixed Shrubland, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [200] G. Flerchinger, “AmeriFlux FLUXNET-1F US-Rws Reynolds Creek Wyoming big sagebrush, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [201] S. Kurc, “AmeriFlux FLUXNET-1F US-SRC Santa Rita Creosote, Ver. 3-5,” *AmeriFlux AMP*, 2022.

- [202] M. Litvak, “AmeriFlux FLUXNET-1F US-Ses Sevilleta shrubland, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [203] R. Scott, “AmeriFlux FLUXNET-1F US-Whs Walnut Gulch Lucky Hills Shrub, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [204] J. Chen, “AmeriFlux FLUXNET-1F US-Wi6 Pine barrens 1 (PB1), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [205] —, “AmeriFlux FLUXNET-1F US-Wi7 Red pine clearcut (RPCC), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [206] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xHE NEON Healy (HEAL), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [207] —, “AmeriFlux FLUXNET-1F US-xJR NEON Jornada LTER (JORN), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [208] —, “AmeriFlux FLUXNET-1F US-xMB NEON Moab (MOAB), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [209] —, “AmeriFlux FLUXNET-1F US-xNQ NEON Onaqui-Ault (ONAQ), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [210] —, “AmeriFlux FLUXNET-1F US-xSR NEON Santa Rita Experimental Range (SRER), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [211] R. Scott, “AmeriFlux FLUXNET-1F US-LS2 San Pedro River Lewis Springs Savanna, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [212] M. Litvak, “AmeriFlux FLUXNET-1F US-Wjs Willard Juniper Savannah, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [213] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xSJ NEON San Joaquin Experimental Range (SJER), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [214] M. Torn and S. Dengel, “AmeriFlux FLUXNET-1F US-NGB NGEE Arctic Barrow, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [215] A. Desai, “AmeriFlux FLUXNET-1F US-Pnp Lake Mendota, Picnic Point Site, Ver. 3-5,” *AmeriFlux AMP*, 2023.

- [216] G. Bohrer, “AmeriFlux FLUXNET-1F US-UM3 Douglas Lake, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [217] L. Kutzbach, “AmeriFlux FLUXNET-1F AR-TF1 Rio Moat bog, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [218] A. Todd and E. Humphreys, “AmeriFlux FLUXNET-1F CA-ARB Attawapiskat River Bog, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [219] —, “AmeriFlux FLUXNET-1F CA-ARF Attawapiskat River Fen, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [220] T. Papakyriakou, “AmeriFlux FLUXNET-1F CA-CF1 Churchill Fen Site 1, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [221] S. Knox, “AmeriFlux FLUXNET-1F CA-DB2 Delta Burns Bog 2, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [222] A. Christen and S. Knox, “AmeriFlux FLUXNET-1F CA-DBB Delta Burns Bog, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [223] T. Roman *et al.*, “AmeriFlux FLUXNET-1F PE-QFR Quistococha Forest Reserve, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [224] B. Olson, “AmeriFlux FLUXNET-1F US-ALQ Allequash Creek Site, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [225] E. Euskirchen, “AmeriFlux FLUXNET-1F US-BZB Bonanza Creek Thermokarst Bog, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [226] —, “AmeriFlux FLUXNET-1F US-BZF Bonanza Creek Rich Fen, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [227] —, “AmeriFlux FLUXNET-1F US-BZo Bonanza Creek Old Thermokarst Bog, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [228] P. Oikawa, “AmeriFlux FLUXNET-1F US-EDN Eden Landing Ecological Reserve, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [229] J. D. Forsythe, M. A. Kline, and T. L. O’Halloran, “AmeriFlux FLUXNET-1F US-HB1 North Inlet Crab Haul Creek, Ver. 3-5,” *AmeriFlux AMP*, 2021.



- [230] E. Euskirchen, G. Shaver, and S. Bret-Harte, “AmeriFlux FLUXNET-1F US-ICs Imnavait Creek Watershed Wet Sedge Tundra, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [231] R. Bracho and C. R. Hinkle, “AmeriFlux FLUXNET-1F US-KS3 Kennedy Space Center (salt marsh), Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [232] J. H. Matthes *et al.*, “AmeriFlux FLUXNET-1F US-Myb Mayberry Wetland, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [233] A. Noormets *et al.*, “AmeriFlux FLUXNET-1F US-NC4 NC AlligatorRiver, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [234] G. Bohrer, “AmeriFlux FLUXNET-1F US-ORv Olentangy River Wetland Research Park, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [235] G. Bohrer and J. Kerns, “AmeriFlux FLUXNET-1F US-OWC Old Woman Creek, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [236] B. Bergamaschi and L. Windham-Myers, “AmeriFlux FLUXNET-1F US-Srr Suisun marsh - Rush Ranch, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [237] R. Vargas, “AmeriFlux FLUXNET-1F US-StJ St Jones Reserve, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [238] A. Valach *et al.*, “AmeriFlux FLUXNET-1F US-Tw1 Twitchell Wetland West Pond, Ver. 3-5,” *AmeriFlux AMP*, 2021.
- [239] E. Eichelmann *et al.*, “AmeriFlux FLUXNET-1F US-Tw4 Twitchell East End Wetland, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [240] A. Valach *et al.*, “AmeriFlux FLUXNET-1F US-Tw5 East Pond Wetland, Ver. 3-5,” *AmeriFlux AMP*, 2022.
- [241] J. Chen and H. Chu, “AmeriFlux FLUXNET-1F US-WPT Winous Point North Marsh, Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [242] National Ecological Observatory Network (NEON), “AmeriFlux FLUXNET-1F US-xBA NEON Barrow Environmental Observatory (BARR), Ver. 3-5,” *AmeriFlux AMP*, 2023.
- [243] G. Vourlitis *et al.*, “AmeriFlux FLUXNET-1F BR-Npw Northern Pantanal Wetland, Ver. 3-5,” *AmeriFlux AMP*, 2022.

## APPENDIX A EDDY COVARIANCE SITE DETAILS

Here we provide an exhaustive list of EC sites used in CarbonSense along with their most recent publication. As per Ameriflux’s data policy, each site has an individual citation with DOI; other networks simply required citation of the unified release. It would be impractical to have each site’s full description in these tables, but the first two letters of each code represent the country where the site is located (ex "DE" for Germany).

Table A.1 EC Sites

Croplands (CRO)					
BE-Lon [38]	CA-ER1 [53]	CA-MA1 [54]	CA-MA2 [55]	CH-Oe2 [39]	CZ-KrP [39]
DE-Geb [38]	DE-Kli [38]	DE-RuS [38]	DE-Seh [18]	DK-Fou [18]	DK-Vng [38]
FI-Jok [18]	FI-Qvd [39]	FR-Aur [38]	FR-EM2 [38]	FR-Gri [38]	FR-Lam [38]
IT-BCi [39]	IT-CA2 [18]	US-A74 [56]	US-ARM [57]	US-Bi1 [58]	US-Bi2 [59]
US-CF1 [60]	US-CF2 [61]	US-CF3 [62]	US-CF4 [63]	US-CRT [64]	US-CS1 [65]
US-CS3 [66]	US-CS4 [67]	US-DFC [68]	US-DS3 [69]	US-Lin [70]	US-Mo1 [71]
US-Mo3 [72]	US-Ne1 [73]	US-RGA [74]	US-RGB [75]	US-RGo [76]	US-Ro1 [77]
US-Ro2 [78]	US-Ro5 [79]	US-Ro6 [80]	US-Tw2 [81]	US-Tw3 [82]	US-Twt [18]
US-xSL [83]					
Closed Shrublands (CSH)					
BE-Maa [38]	IT-Noe [18]	US-KS2 [84]	US-Rls [85]	US-Rms [86]	US-Rwe [87]
US-Rwf [88]					
Cropland/Natural Vegetation Mosaics (CVM)					
US-HWB [89]	US-xDS [90]				
Deciduous Broadleaf Forests (DBF)					
AU-Lox [18]	BE-Lcr [38]	CA-Cbo [91]	CA-Oas [18]	CA-TPD [92]	CZ-Lnz [38]
CZ-Stn [39]	DE-Hai [38]	DE-Hzd [39]	DE-Lnf [18]	DK-Sor [38]	FR-Fon [38]
FR-Hes [38]	IT-BFt [38]	IT-CA1 [18]	IT-CA3 [18]	IT-Col [18]	IT-Isp [18]
IT-PT1 [18]	IT-Ro1 [18]	IT-Ro2 [18]	JP-MBF [18]	MX-Tes [93]	PA-SPn [18]
US-Bar [94]	US-Ha1 [95]	US-MMS [96]	US-MOz [97]	US-Oho [98]	US-Rpf [99]
US-UMB [100]	US-UMd [101]	US-WCr [18]	US-Wi1 [102]	US-Wi3 [103]	US-Wi8 [104]
US-xBL [105]	US-xBR [106]	US-xGR [107]	US-xHA [108]	US-xML [109]	US-xSC [110]
US-xSE [111]	US-xST [112]	US-xTR [113]	US-xUK [114]	ZM-Mon [18]	

Table A.2 EC Sites (cont'd)

Deciduous Needleleaf Forests (DNF)					
BR-CST [115]					
Evergreen Broadleaf Forests (EBF)					
AU-Cum [18]	AU-Rob [18]	AU-Wac [18]	AU-Whr [18]	AU-Wom [18]	BR-Sa3 [18]
CN-Din [18]	FR-Pue [38]	GF-Guy [38]	GH-Ank [18]	IT-Cp2 [38]	IT-Cpz [18]
MY-PSO [18]					
Evergreen Needleleaf Forests (ENF)					
AR-Vir [18]	CA-Ca1 [116]	CA-Ca2 [117]	CA-LP1 [51]	CA-Man [18]	CA-NS1 [118]
CA-NS2 [119]	CA-NS3 [120]	CA-NS4 [121]	CA-NS5 [122]	CA-Obs [18]	CA-Qfo [123]
CA-SF1 [124]	CA-SF2 [125]	CA-TP1 [126]	CA-TP2 [18]	CA-TP3 [127]	CA-TP4 [18]
CH-Dav [38]	CN-Qia [18]	CZ-BK1 [38]	CZ-RAJ [39]	DE-Lkb [18]	DE-Msr [38]
DE-Obe [39]	DE-RuW [38]	DE-Tha [38]	DK-Gds [38]	FI-Hyy [38]	FI-Ken [38]
FI-Let [38]	FI-Sod [18]	FI-Var [38]	FR-Bil [38]	FR-FBn [39]	FR-LBr [18]
IL-Yat [39]	IT-La2 [18]	IT-Lav [39]	IT-Ren [38]	IT-SR2 [38]	IT-SRo [18]
NL-Loo [18]	RU-Fy2 [39]	RU-Fyo [39]	SE-Htm [38]	SE-Nor [38]	SE-Ros [39]
SE-Svb [38]	US-BZS [128]	US-Blo [18]	US-CS2 [129]	US-Fmf [130]	US-Fuf [131]
US-GBT [18]	US-GLE [132]	US-HB2 [133]	US-HB3 [134]	US-Ho2 [135]	US-KS1 [136]
US-Me1 [137]	US-Me2 [138]	US-Me3 [139]	US-Me4 [18]	US-Me5 [18]	US-Me6 [140]
US-NC1 [141]	US-NC3 [142]	US-NR1 [143]	US-Prr [18]	US-Vcm [144]	US-Vcp [145]
US-Wi0 [146]	US-Wi2 [18]	US-Wi4 [147]	US-Wi5 [148]	US-Wi9 [149]	US-xAB [150]
US-xBN [151]	US-xDJ [152]	US-xJE [153]	US-xRM [154]	US-xSB [155]	US-xTA [156]
US-xYE [157]					
Grasslands (GRA)					
AT-Neu [18]	AU-DaP [18]	AU-Emr [18]	AU-Rig [18]	AU-Stp [18]	AU-TTE [18]
AU-Ync [18]	BE-Dor [39]	CA-MA3 [158]	CH-Aws [39]	CH-Cha [39]	CH-Fru [39]
CH-Oe1 [18]	CN-Cng [18]	CN-Dan [18]	CN-Du2 [18]	CN-Du3 [18]	CN-HaM [18]
CN-Sw2 [18]	CZ-BK2 [18]	DE-Gri [38]	DE-RuR [38]	DK-Eng [18]	FR-Mej [38]
FR-Tou [38]	GL-ZaH [38]	IT-MBo [39]	IT-Niv [38]	IT-Tor [38]	NL-Hor [18]
PA-SPs [18]	RU-Ha1 [18]	SE-Deg [38]	US-A32 [159]	US-AR1 [160]	US-AR2 [161]
US-ARb [162]	US-ARc [163]	US-BRG [164]	US-Cop [165]	US-Goo [18]	US-Hn2 [166]
US-IB2 [18]	US-KFS [167]	US-KLS [168]	US-Kon [169]	US-Mo2 [170]	US-NGC [171]
US-ONA [172]	US-Ro4 [173]	US-SRG [174]	US-Seg [175]	US-Sne [176]	US-Snf [177]
US-Var [178]	US-Wkg [179]	US-xAE [180]	US-xCL [181]	US-xCP [182]	US-xDC [183]
US-xKA [184]	US-xKZ [185]	US-xNG [186]	US-xWD [187]		
Mixed Forests (MF)					
AR-SLu [18]	BE-Bra [38]	BE-Vie [38]	CA-Gro [188]	CD-Ygb [38]	CH-Lae [39]
CN-Cha [18]	DE-Har [38]	DE-HoH [38]	JP-SMF [18]	US-Syv [189]	US-xDL [190]
US-xUN [191]					

Table A.3 EC Sites (cont'd)

Open Shrublands (OSH)					
CA-NS6 [192]	CA-NS7 [40]	CA-SF3 [18]	ES-Agu [39]	ES-Amo [18]	ES-LJu [39]
ES-LgS [18]	ES-Ln2 [18]	GL-Dsk [38]	IT-Lsn [38]	RU-Cok [18]	US-EML [193]
US-Fcr [194]	US-Hn3 [195]	US-ICH [196]	US-ICt [197]	US-Jo1 [198]	US-Jo2 [199]
US-Rws [200]	US-SRC [201]	US-Ses [202]	US-Sta [18]	US-Whs [203]	US-Wi6 [204]
US-Wi7 [205]	US-xHE [206]	US-xJR [207]	US-xMB [208]	US-xNQ [209]	US-xSR [210]
Savannas (SAV)					
AU-ASM [18]	AU-Cpr [18]	AU-DaS [18]	AU-Dry [18]	AU-GWW [18]	CG-Tch [18]
ES-Abr [39]	ES-LM1 [39]	ES-LM2 [39]	SD-Dem [18]	SN-Dhr [18]	US-LS2 [211]
US-Wjs [212]	US-xSJ [213]				
Snow and Ice (SNO)					
US-NGB [214]					
Water Bodies (WAT)					
US-Pnp [215]	US-UM3 [216]				
Permanent Wetlands (WET)					
AR-TF1 [217]	AU-Fog [18]	CA-ARB [218]	CA-ARF [219]	CA-CF1 [220]	CA-DB2 [221]
CA-DBB [222]	CN-Ha2 [18]	CZ-wet [38]	DE-Akm [39]	DE-SfN [18]	DE-Spw [18]
DE-Zrk [18]	DK-Skj [38]	FI-Lom [18]	FI-Sii [38]	FR-LGt [38]	GL-NuF [38]
GL-ZaF [18]	IE-Cra [39]	PE-QFR [223]	RU-Che [18]	SE-Sto [38]	SJ-Adv [18]
UK-AMo [38]	US-ALQ [224]	US-Atq [18]	US-BZB [225]	US-BZF [226]	US-BZo [227]
US-EDN [228]	US-HB1 [229]	US-ICs [230]	US-Ivo [18]	US-KS3 [231]	US-Los [18]
US-Myb [232]	US-NC4 [233]	US-ORv [234]	US-OWC [235]	US-Srr [236]	US-StJ [237]
US-Tw1 [238]	US-Tw4 [239]	US-Tw5 [240]	US-WPT [241]	US-xBA [242]	
Woody Savannas (WSA)					
AU-Ade [18]	AU-Gin [18]	AU-How [18]	AU-RDF [18]	BR-Npw [243]	ES-Cnd [39]