

# Climate-Aware Machine Learning for Above-Ground Biomass Estimation

Aske Meineche, IT University of Copenhagen

---

This study explores the role of data science, machine learning, and artificial intelligence in addressing environmental challenges, specifically focusing on the estimation of Above-Ground Biomass (AGBM) using satellite imagery. The research aims to compare the effectiveness of temporal and spatial modelling techniques in AGBM estimation and to assess the utility of the AI-Climate Alignment Framework proposed by Kaack et al. (2022) in guiding environmentally responsible model development. A tree-based learner and a neural learner are trained on a small dataset, using a temporal and a spatial representation. The results show that the tree-based learner emits less carbon in inference and outperforms the neural learner when training on small samples.

**Keywords:** Machine learning, remote sensing, above-ground biomass estimation, earth observation, geospatial data science, climate-aware machine learning.

---

## 1. Introduction

The environmental crises regarding global climate change and threats to the world's ecosystems have been widely accepted by governmental bodies and international organizations as one of the most important global challenges on the agenda. Important documents, such as the sixth Assessment Report authored by the International Panel on Climate Change (Adler et al., 2022) and the third official accounting of the state of the planetary boundaries (Richardson et al., 2023) have recently underscored the severity of the threats to the fundamental conditions making our planet habitable by human civilization.

Making science-based decisions in the face of these crises requires natural resource management systems (Lister et al., 2020) as field measurements are infeasible at scale (Ghosh and Behra, 2018). Many countries currently mandate national forest inventory programs (NFI) as this tool provides support in both climate and environmental questions (Lister et al., 2020). NFIs contain a variety of measurements regarding forests, among other above ground biomass (AGB). This can support climate modelling as change in biomass strongly influences carbon sequestration (Jevšenak and Skudnik, 2021), as well as ecosystem response to climate change (Novo-Fernández et al., 2019).

While field measurements have traditionally been used in many countries (Ghosh and Behra, 2018; Magnussen et al., 2018), remote sensing is increasingly being utilized as both primary and auxiliary data (Magnussen et al., 2018; Lister et al., 2020). Performing these complex function mappings between remotely sensed data from either UAVs or satellites has in part been aided by

the advances in machine learning, which is playing a significant part in modern AGB-estimation (Li et al., 2020; Ghosh and Behra, 2018; Tamimina et al., 2022).

However, as tree-based machine learning algorithms as used in Li et al. (2020), Ghosh and Behra (2018), and Tamiminia et al. (2022) differ significantly in parameters, performance, complexity and training times from artificial neural nets as used in Vawda et al. (2024). As the data volume of remotely sensed images of terrestrial areas is significant it is relevant to consider complexity and emissions training and inference, rather than solely being guided by accuracy (Strubell et al., 2019).

In this paper two modelling frameworks for Above Ground Biomass are compared with regards to both accuracy and emissions, relying on the novel framework from (Kaack et al., 2022) in guiding emission measurements.

## 2. Literature Review

Machine learning tools applied to remote sensing data is a common combination to predict, classify, estimate or delineate factors of interest, whether it is algal blooms (Hill et al., 2020), land use (Bastini et al., 2023), wind turbines (Zhou et al., 2019), or monitoring of photovoltaic plants (Costa et al., 2021). However, remotely sensed data can be gathered from both satellites, planes or UAVs etc, which can vastly impact the temporal, spatial and spectral resolution, and in turn may require differing model paradigms. Remotely sensed data is often described in both time, space and with a variety of spectral filters. This requires researchers to decide whether to represent the data temporally (Tamiminia et al., 2022), spatially (Vawda et al., 2024), or spatio-temporally (Hill et al., 2020).

The introduction of cloud infrastructure for remote sensing analysis, such as Google Earth Engine, has popularized non-neural machine learning models for these tasks (Tamiminia, 2020). Literature shows inconsistent results as to whether above ground biomass is best modelled neurally (Ma et al., 2023); Dong et al., 2020), or non-neurally (Singh et al., 2022; Li et al., 2022).

As of writing this paper, no literature has been found comparing spatial to temporal modelling of AGB. While a variety of methods exist to extract spatial features from data, one of the most common frameworks is the Convolutional Neural Network (CNN), which was popularized after the performance of AlexNet (Krizhevsky et al., 2012). While temporal modelling can be done in a variety of fashions both neurally, using recurrent neural nets (RNNs) or Long-Short-Term-Memory nets (LSTMs), and non-neurally using gradient boosted trees or similar.

## 3. Climate-Aware Machine Learning

When comparing machine learning models, it is common to solely rely on performance metrics, such as RMSE, R2 or accuracy (Strubell et al., 2019). However, as neural models are increasing in complexity the underlying training time and emissions are also of interest (Strubell et al., 2019). This is especially interesting in an earth observation context as the data volume can become cumbersome to train and process, if the data is represented with a standard spatial resolution of 10m per pixel as is the case with ESA's Sentinel satellites (Tamiminia et al., 2020). Additionally, emissions during training and inference becomes even more semantically relevant when building models for climate and environmental tasks.

A novel framework developed by (Kaack et al., 2022) divides emission analysis of machine learning systems into a three-tier system. The first tier revolves around measuring the compute related impacts through sheer power consumption in training and inference. The second tier revolves around the application area of models. Machine learning can both be used to support renewable energy generation (Perera et al., 2014) or estimate recovery scenarios for oil wells (Mahdaviara et al., 2022). This tier is much more complex to provide estimates of, as it requires intimate sector knowledge and life cycle analysis calculations (Kaack et al., 2022). The third tier attempts to define and estimate system-level impacts. This is defined as the complex societal responses to a machine learning service, such as the effect of self-driving cars on the development of public infrastructure. According to the authors, this tier is even more complex to estimate but contains the highest potential for emissions.

The only concrete tool currently developed to measure either boundary is CarbonTracker (Anthony et al., 2020), which is directly recommended in (Kaack et al., 2020). This framework uses location access to query the current approximate electricity mix on the location of the computer resource. From this it estimates the approximate greenhouse gas emissions from the model.

## 4. Data

### 4.1 Data Acquisition

The feature data has been collected by the Sentinel-1 satellite (S1) and Sentinel-2 (S2) satellite of the Copernicus programme, operated by the European Space Agency. S1 is an active Synthetic Aperture Radar satellite (SAR) and supplies 4 different spectral bands. A SAR satellite, as opposed to an optical satellite, emits radiation and collects the returning signal, making it extremely valuable when observing biomass obscured by a dense tree canopy and makes the satellite unaffected by cloud cover (Zhang et al., 2012). Each monthly observation is, therefore, an average of all 5 images captured of each chip during the month. S2 is an optical satellite that supplies 10 bands, as well as a band estimating the cloud cover in the current image. The optical sensor's inability to penetrate cloud cover, also makes these images more vulnerable to noise. The images from S2 have, therefore, been handpicked by experts, as simply averaging the data out could lead to loss of data quality.

### 4.2 Data Description

While the raw feature data has been collected by the Copernicus program, it has been curated and matched with target data through the cooperation of KU Leuven and University of Liège and made available through the DrivenData platform at the BioMassters competition (DrivenData, 2023). Similarly, the target data has been collected by the Finnish Forest Centre and curated by KU Leuven and University of Liège.

The data modelled in this paper is comprised of images of roughly 8000 2560m by 2560m patches of forest in Finland. All observations were captured between September of 2020 and September 2021. Each patch of forest, referred to as a chip, is captured once a month by two satellites, making up the feature dataset. Each pixel of each image represents a 10m x 10m patch of forest and each tile has been captured using a total of 16 satellite bands. This makes the dataset a multi-spectral raster time-series, as each chip is a raster data point (image) represented by 16 different

spectral channels, for each month during a year. The target data is a single 256x256 matrix corresponding to each chip, where each pixel contains the approximate annual mass of above ground biomass in that 10m x 10m patch, measured in tons. The AGBM target approximations have been measured using a drone mounted Light Detection And Ranging (LiDAR).

#### 4.3 Data Curation

Due to a variety of shortcomings in the technical specifications of S2, only 1251 chips of the original 8000 are complete with all possible measurements. Most incomplete images are missing measurements for 1-2 months, which makes data imputation a feasible, but possibly difficult task. For this study a sample from the complete 1251 chips has been used as the foundation for training and testing.

### 5. Data Representation

The data can be modelled in a strictly spatial fashion, disregarding the temporal dimension. This leaves us with a relatively simple 2D representation, solely relying on information found in neighbouring pixels. This transforms each data point into 12 data points as each month is treated as independent data with the same target values, leading to an increase from 1251 data points to 15012 data points. This makes the model unaware of seasonal changes but enables it to leverage geometric information across the entire image.

Complimentary, temporal representation is also possible, by disregarding all geometric information found in the neighbourhood of each pixel, and instead estimating each pixel based on its historical values (Tamiminia et al., 2022). This pixel-by-pixel assessment was used widely before the introduction of more sophisticated architecture. This representation makes the data volume increase drastically as each data point is effectively turned into 65536 (256 x 256) data points. This results in a total data volume of 81 985 536 data points for the 1251 complete data points. This data structure is easily modelled by most machine learning models as the input is strictly tabular.

### 6. Models

#### 6.1 XGBoost

Extreme Gradient Boosting, or XGBoost, is a tree-based algorithm formally introduced in 2016 (Chen & Guestrin, 2016), and has been the industry standard for simple, tabular machine learning problems since then. The fundamental idea behind gradient boosting is to train a series of weak learners to correct the earlier model's mistakes. XGBoost and other gradient boosted tree models are often used in geoinformatics research, as it is available through the low-code interface of Google Earth Engine (Tamiminia, 2020; Li et al., 2020; Ma et al., 2023; Huang et al., 2023; Wai et al., 2022).

#### 6.2 Convolutional Neural Networks

Convolutional Neural Networks are a type of specialized neural networks adept in representing geometric data, such as images. They were first introduced in the 1980's, however, due to their computational complexity, they were not popularized before the success of AlexNet in 2012 (Krizhevsky et al., 2012). CNNs function by training a predefined number of filters which are moved

across the image, calculating features based on all pixels within the filter. This, essentially, allows information in a given pixel to commute between neighbouring pixels in each convolutional layer. Serialized convolutional layers allows the network to look for more and more complex geometries, and allowing the information in each pixel to "travel" longer distances within the image. The logic of training filters, as opposed to training a traditional, fully connected Multi-Layer Perceptron, is that this offers a form of parameter sharing. This means that the network never has to handle the entire image at once, but only whatever is within the filter - greatly reducing the number of parameters and training data needed to tune them.

## 7. Experiment Design

The experiment was designed to be replicated for multiple types of model architecture. Both models have been trained on 200 randomly selected chips, which results in 2400 separate data points for the CNN and 13 000 000 for the XGBoost. The target chips contain several outlier values, with biomass exceeding 5000 tons in a single 10m x 10m pixel. According to the curators of the dataset all pixels exceeding 500 tons AGBM are probably false measurements. Therefore, all target values in the training data over 500 tons have been set to 500, to keep the model being overly affected by these outliers. All models were tested on the same 50 chips, which were also randomly selected from the 1251 complete chips. Steps were taken to ensure that there is zero overlap between the sampled training data and the test data. A 5-fold cross validation was used to test the models' ability to generalize to new data.

A hyper-parameter tuning was performed for both models using select parameters. The XGBoost was tuned on the max depth parameter and the learning rate parameter which are both very influential on the model's tendency to overfit.

The values tested for the max depth parameter were [3, 4, 5, 6, 7, 8, 9, 10], while the learning rates examined were [0.01, 0.1, 0.3, 0.5].

The values and parameters selected for tuning of the XGBoost were selected from Tamiminia et al. (2022) although several of the parameters tested in that paper have been omitted due to computational limitations.

The networks used in this study all have three layers, a constant filter size of 3x3 and stride 1. The number of filters in each layer is a tunable parameter that has been examined during training. After each convolutional layer, a Relu activation function is applied, followed by a max-pooling layer. The max-pooling layer decreases the size of the image, by letting each pixel be "represented" by the maximum value within the pooling kernel. This is another feature that reduces the number of parameters in the model, and thereby the required amount of training data. All tested CNNs were run using varying numbers of filters. The number of filters for each layer were either 4, 8 or 16 and were all trained for 100 epochs. However, due to computational limitations only the best performing model was essentially run through the entire experiment with all folds, carbon tracking and final test set. Therefore, carbon emissions were only tracked for the best permutation, while 32 permutations of the XGBoost have been tracked.

The model parameters with the best average performance after K-fold cross-validation was selected to be trained on the entire training sample. The fully trained model was then evaluated on the final test set.

The experiments' carbon emissions were tracked in two important phases: training and inference. Training was evaluated for the entire cross-validation as well as parameter tuning. This makes the emissions dependent on the number of parameters selected for tuning and the number of values tested. As only 1 permutation was examined for the CNN while 32 separate configurations of the XGBoost was tested this unfairly favours the convolutional neural network.

Inference emissions were evaluated for the fully trained model on the same 50 test instances, and is, therefore, the most comparable.

To test whether the models can model the problem acceptably they are compared to a baseline set by DrivenData on the same dataset. They use an unknown model on a test set of unknown size with a root mean-squared error of 101 tonnes. While not directly comparable, since the chip selection is different, it provides an approximate benchmark to understand the range of acceptable errors.

## 8. Results

When inspecting the performances of the two models and representations on the dataset we see that both models perform better on the test set than on the train set.

We see that the XGBoost with a temporal representation slightly outperforms the CNN with an average error of approximately 59 tonnes pr. pixel, while the convolutional neural network misses the target by 67 tonnes on average (Figure 1).



Figure 1: Test and Training Performance.

### 8.1 XGBoost

The chip where the XGBoost performs the best is predicted with an average error of 15 tons per pixel (Figure 2). The chip has a very low average biomass, with most of the chip being virtually

devoid of biomass. The model has correctly estimated the shape and general mass of most of the forest, apart from the few very high-density pixels, exceeding 500 tons of biomass per pixel.




Figure 2: Ground truth and XGBoost Prediction on Best Chip.

The pixel where the XGBoost performs the worst is one with significantly higher average biomass. While some of the pixels are somewhat well-approximated, it generalizes an entire error to some average value yielding an average error of 120 tons per pixel (Figure 3). We can see that the model performs very well on a chip with generally low biomass, while completely under-fitting on chips with a wider distribution of values (8, 9). When comparing the ground truth values with the error for each pixel for the two chips, we can see that the model has a generally low error for the low biomass chip, while it clearly has two distinct strategies for the high biomass chip: either estimating around zero or estimating around 220 tons (Figure 4).

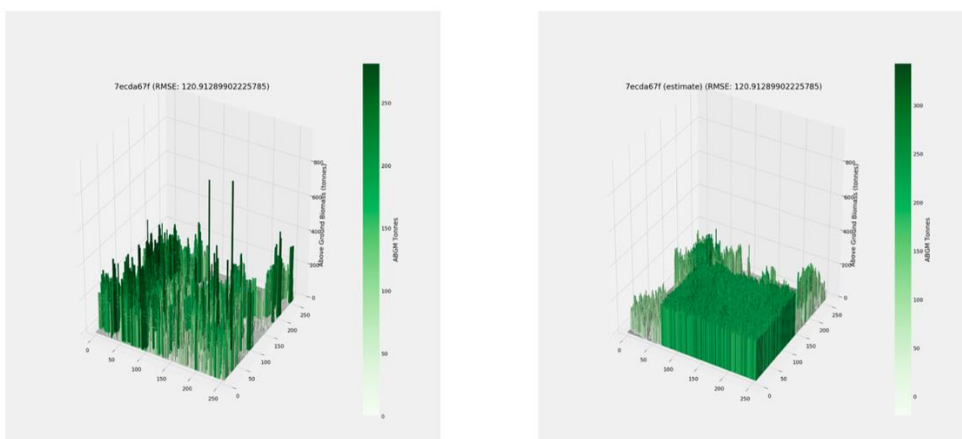


Figure 3: Ground truth and XGBoost Prediction on Worst Chip.

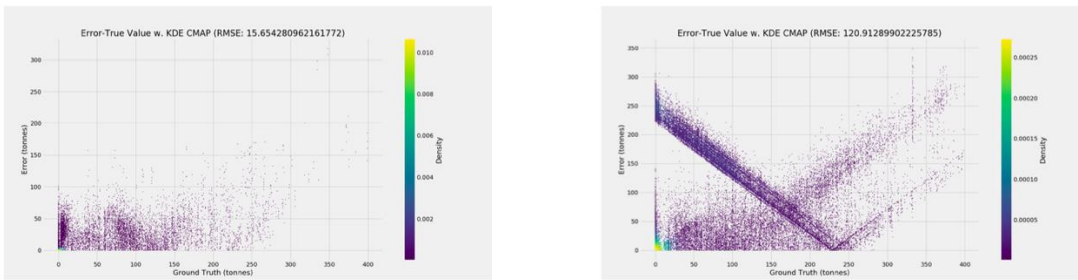


Figure 4: Distribution of Errors on best and worst chip for XGBoost.

### 8.2 CNN

We see that the CNN also performs its best on the same chip as the XGBoost, due to its low general biomass (Figure 5). It performs slightly worse with an average error of 25 tons per pixel.




Figure 5: Ground truth and CNN Prediction on Best Chip.

Just as the XGBoost the CNN also completely excludes the higher AGBM-values, which, in the worst case, probably are also wrong measurements, as biomass exceeding 400 tons in a 10x10m tile. The kernel logic of the CNN is visually apparent as transitions between neighbouring pixels are much smoother than in the, spatially unaware, XGBoost. The worst estimated chip has an average error of 125 tons per pixel and has a relatively high average biomass. The model has almost exclusively estimated values below 150 tons, while the non-outlier tiles range up to 400 tons (Figure 6).



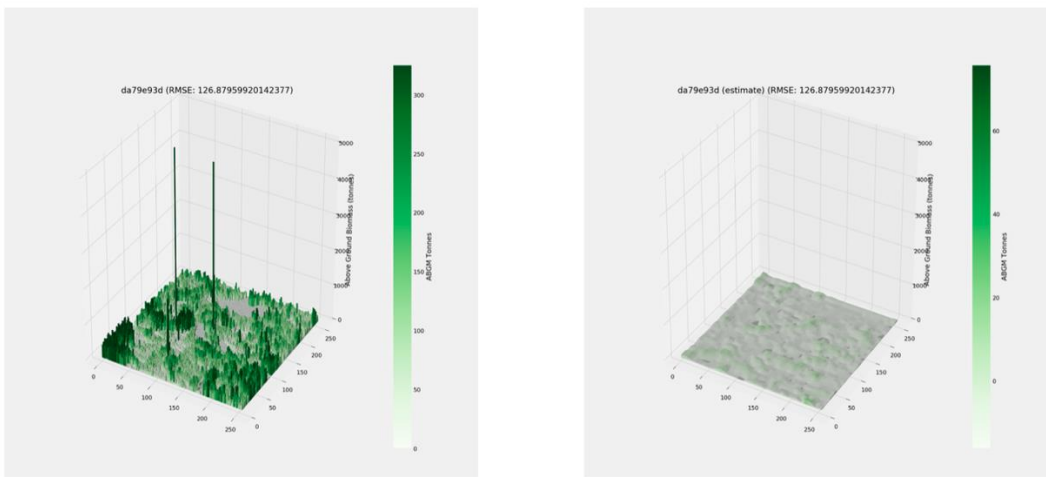


Figure 6: Ground truth and CNN Prediction on Worst Chip.

While the outlier values make it harder to visually inspect the model’s performance on this chip, it has estimated a plateau between 0 and 150 tons, which is also where most of the data is located as seen in Figure 7.

We see that the distribution of estimates is more condensed in the low end of the biomass spectrum, which pays off in the case of the low-biomass chip, while it results in completely wrong estimates in the high-biomass chip.

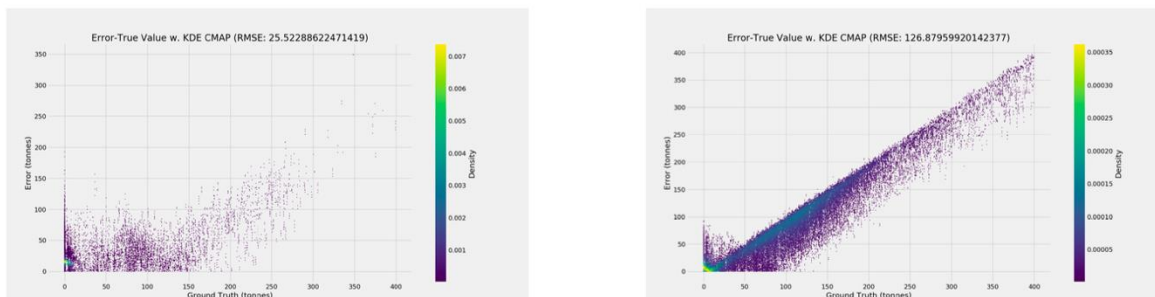


Figure 7: Distribution of Errors on best and worst chip for XGBoost.

### 8.3 Model Comparison

When comparing the models’ average errors on the entire dataset with the similarity between each pixel and its’ neighbours we do see that most pixels are somewhat like their neighbourhood in Figure 8. Similarity has been quantified as the mean difference between all eight neighbours of each pixel. However, while the models are close in similar pixels, the XGBoost seems to be better at estimating the pixels which are very dissimilar to its neighbours.

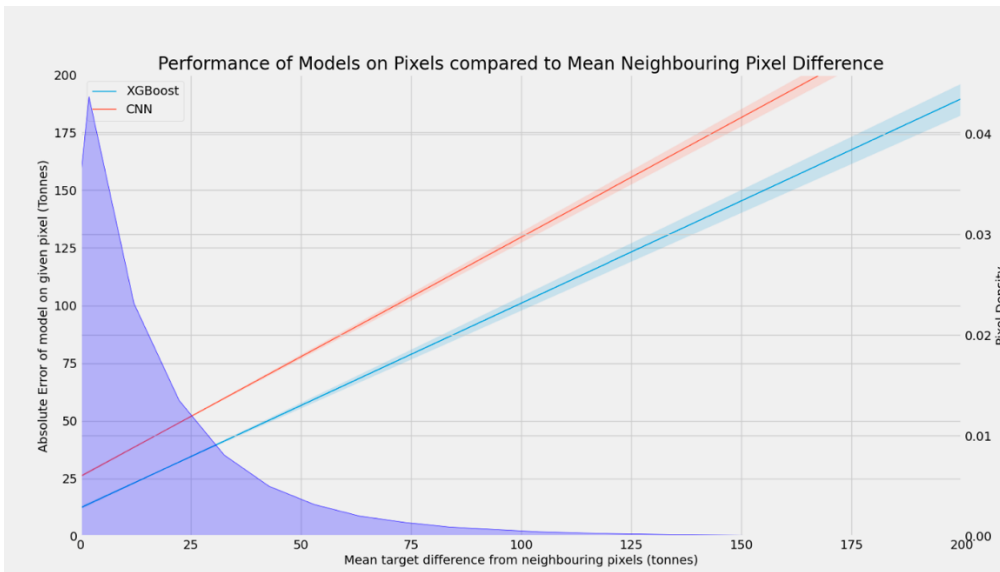


Figure 8: Model Performances compared to pixel neighbourhood similarity.

### 8.4 Carbon Footprints

In **Fejl! Henvisningskilde ikke fundet.** we see that CNN emits 38 times more CO<sub>2</sub> during inference, which is not surprising considering the tree-based structure of the XGBoost compared to the complex neural structure of the CNN. The inference estimation is performed over the entire test set, so the emissions per data point is one 50th of what is displayed in the bar charts. By that estimate the CNN would emit, roughly 2.8 kg of CO<sub>2</sub> when predicting on the entire dataset, while the XGBoost would only emit around 75 grams. Results of emissions testing during training are omitted due to infrastructure problems during training.

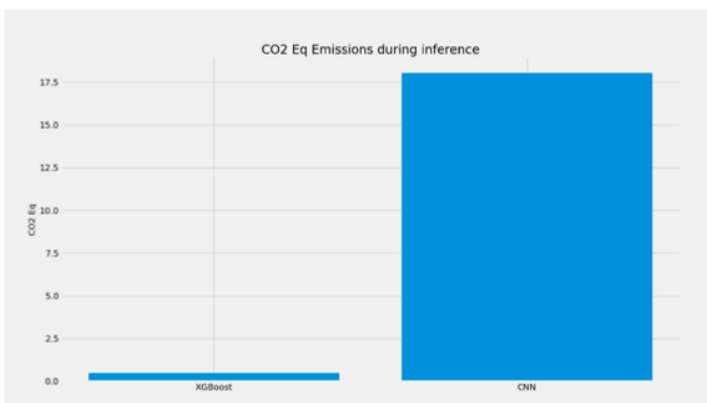


Figure 9: CO<sub>2</sub> Emissions for models during inference.

## 9. Discussion

Both trained models outperform the benchmark model proposed by DrivenData, indicating that there is information to be gained in both the temporal and spatial dimensions. We see that the XGBoost generally outperforms the CNN on the test set in figure 3. The fact that both models perform better during test, compared to training suggests that the randomly selected test set is

comprised of more "trivial" examples than the train set. Models correctly fitted to the global dataset distribution perform better during training than during tests. Both models overfit relatively low biomass data points, due to general over-representation of these pixels in the dataset as seen in Figure 4, Figure 7 and Figure 8. This error has likely occurred due to the problem in performing stratified sampling for the CNN as a single datapoint is represented by 2562 pixels.

One solution is to use a more systematic approach when sampling or re-sampling the training data, possibly using data augmentation to avoid the model overfitting to single high-biomass data points. In figure 8 we see that the XGBoost outperforms the CNN both on pixels that are similar and dissimilar, although the gap widens as the dissimilarity grows. Theoretically, the CNN should have outperformed the XGBoost on high similarity neighbourhoods, as this geometric information is directly used as the primary decision basis for the CNN, while it is unavailable to the XGBoost. It is possible that this means that, either the CNN has underfit the problem, or that the metric used to measure pixel/neighbourhood-similarity is too simple. When examining figures 4 and 7 it seems that the models have two strategies - one to estimate data below 150 tons of biomass, which they perform reasonably well on, and one to estimate data in the 200+ AGBM range, which seems to work sub-optimally. This suggests that both models could implicitly be segmenting the images, in which case explicit image segmentation models such as UNets or SWIN-transformers might help improve the model's performance. While the CNN was subject to a hyperparameter search outside of the experiment design, the experiment and potential performance would still benefit greatly from running the parameter search within the experiment.

The application of the CarbonTracker package worked very well on inference. It is easy to see that the XGBoost emits much less carbon than CNN on inference, which is very useful information as most computation for ML is used during inference rather than training depending on the modelling setup.

A potential issue with using the CarbonTracker framework for ML-tasks is its incompatibility with specific chip manufacturers, such as AMD. AMD provides a significant portion of the computer chips used today, and not supporting this architecture significantly decreases the number of machines which can run CarbonTracker.

Lastly, a concrete practice of comparing emissions to model performance needs to be established. While measuring emissions is a tremendous step towards more sustainable modelling practices, deciding when a model's increase in emissions cannot be justified by its' increase in performance is a complicated question to answer, and probably one that is highly specific to each application.

## 10. Conclusion

Both temporal and spatial data have value in estimating Above Ground Biomass from satellite imagery, both beating a benchmark set by DrivenData. However, the temporal representation paired with an XGBoost outperformed the CNN using the spatial representation when training on a relatively small dataset, which makes it possible to perform this task on smaller machines. As both models under-fitted on high AGBM-values, selective resampling based on average AGBM-values might help the modelling process. As both representations conveyed enough information to beat the benchmark model, combining the representations in a spatio-temporal model might result in a more accurate prediction.

When applying the AI Climate Alignment Framework by Kaack et al. (2022) the CarbonTracker tool (Anthony et al., 2020) clearly shows that the XGBoost emits as much as 38 times less carbon on inference. A standard for performing emission measurements would make comparisons much easier. In conclusion, the first level of the framework is adequately defined and has technological support but lacks community standards for measurements.

## References

- Adler C., Wester P., Bhatt I., Huggel C., Insarov G.E., Morecroft M.D., Muccione V., Prakash A. (2022) Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK and New York, USA: Cambridge University Press. [Chapters 5 & 6]
- Anthony, L.F.W., Kanding, B. and Selvan, R. (2020) 'Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models'. arXiv.
- Bastani, F. et al. (2022) 'SatlasPretrain: A Large-Scale Dataset for Remote Sensing Image Understanding Supplementary Material'.
- Chen, L. et al. (2018) 'Estimation of Forest Above-Ground Biomass by Geographically Weighted Regression and Machine Learning with Sentinel Imagery', *Forests*, 9(10), p. 582.
- Chen, T., & Guestrin, C. (2016, August). Xgboost: A scalable tree boosting system. In *Proceedings of the 22<sup>nd</sup> acm sigkdd international conference on knowledge discovery and data mining* (pp. 785-794).
- Costa, M.V.C.V. da et al. (2021) 'Remote Sensing for Monitoring Photovoltaic Solar Plants in Brazil Using Deep Semantic Segmentation', *Energies*, 14(10), p. 2960.
- Dong, L. et al. (2020) 'Application of Convolutional Neural Network on Lei Bamboo Above Ground Biomass (AGB) Estimation Using Worldview-2', *Remote Sensing*, 12(6), p.958.
- DrivenData (2022) The BioMassters, DrivenData at: <https://www.drivendata.org/competitions/99/biomass-estimation>, [accessed at 19 August 2024]
- Ghosh, S.M. and Behera, M.D. (2018) 'Aboveground biomass estimation using multi-sensor data synergy and machine learning algorithms in a dense tropical forest', *Applied Geography*, 96, pp. 29–40.
- Hill, P.R. et al. (2020) 'HABNet: Machine Learning, Remote Sensing Based Detection and Prediction of Harmful Algal Blooms'. arXiv.
- Heterogeneity Huang, T. et al. (2023) 'Estimating the Aboveground Biomass of Various Forest Types with High at the Provincial Scale Based on Multi-Source Data', *Remote Sensing*, 15(14), p. 3550.
- Jevšenak, J. and Skudnik, M. (2021) 'A random forest model for basal area increment predictions from national forest inventory data', *Forest Ecology and Management*, 479, p. 118601.
- Kaack, L.H. et al. (2022) 'Aligning artificial intelligence with climate change mitigation', *Nature Climate Change*, 12(6), pp. 518–527.
- Li, Y. et al. (2020) 'Forest aboveground biomass estimation using Landsat 8 and Sentinel-1A data with machine learning algorithms', *Scientific Reports*, 10(1), p. 9952.
- Lister, A.J. et al. (2020) 'Use of Remote Sensing Data to Improve the Efficiency of National Forest Inventories: A Case Study from the United States National Forest Inventory', *Forests*, 11(12), p. 1364.
- Ma, Y. et al. (2023) 'Novel Features of Canopy Height Distribution for Aboveground Biomass Estimation Using Machine Learning: A Case Study in Natural Secondary Forests', *Remote Sensing*, 15(18), p. 4364.

- Magnussen, S., Nord-Larsen, T. and Riis-Nielsen, T. (2018) 'Lidar supported estimators of wood volume and aboveground biomass from the Danish national forest inventory (2012–2016)', *Remote Sensing of Environment*, 211, pp. 146–153.
- Mahdaviara, M., Sharifi, M. and Ahmadi, M. (2022) 'Toward evaluation and screening of the enhanced oil recovery scenarios for low permeability reservoirs using statistical and machine learning techniques', *Fuel*, 325, p. 124795.
- Novo-Fernández, A. et al. (2019) 'Integration of National Forest Inventory and Nationwide Airborne Laser Scanning Data to Improve Forest Yield Predictions in North-Western Spain', *Remote Sensing*, 11(14), p. 1693.
- Perera, K.S., Aung, Z. and Woon, W.L. (2014) 'Machine Learning Techniques for Supporting Renewable Energy Generation and Integration: A Survey', in W.L. Woon, Z. Aung, and S. Madnick (eds) *Data Analytics for Renewable Energy Integration*. Cham: Springer International Publishing, pp. 81–96.
- Richardson, K., Steffen, W., Lucht, W., Bendtsen, J., Cornell, S. E., Donges, J. F. & Rockström, J. (2023). Earth beyond six of nine planetary boundaries. *Science advances*, 9(37), eadh2458.
- Singh, C. et al. (2022) 'Remote sensing-based biomass estimation of dry deciduous tropical forest using machine learning and ensemble analysis', *Journal of Environmental Management*, 308, p. 114639.
- Strubell, E., Ganesh, A. and McCallum, A. (2019) 'Energy and Policy Considerations for Deep Learning in NLP'. arXiv.
- Tamiminia, H. et al. (2020) 'Google Earth Engine for geo-big data applications: A meta analysis and systematic review', *ISPRS Journal of Photogrammetry and Remote Sensing*, 164, pp. 152–170.
- Tamiminia, H. et al. (2022) 'Evaluating pixel-based and object-based approaches for forest above-ground biomass estimation using a combination of optical, SAR, and an extreme gradient boosting model', *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, V-3–2022, pp. 485–492.
- Vawda, M.I. et al. (2024) 'Comparing the Utility of Artificial Neural Networks (ANN) and Convolutional Neural Networks (CNN) on Sentinel-2 MSI to Estimate Dry Season Aboveground Grass Biomass', *Sustainability*, 16(3), p. 1051.
- Wai, P., Su, H. and Li, M. (2022) 'Estimating Aboveground Biomass of Two Different Forest Types in Myanmar from Sentinel-2 Data with Machine Learning and Geostatistical Algorithms', *Remote Sensing*, 14(9), p. 2146.
- Zhang, H., Zhang, Y. and Lin, H. (2012) 'A comparison study of impervious surfaces estimation using optical and SAR remote sensing images', *International Journal of Applied Earth Observation and Geoinformation*, 18, pp. 148–156.
- Zhou, S. et al. (2019) 'DeepWind: Weakly Supervised Localization of Wind Turbines in Satellite Imagery'.