

# Contrasting Local and Global Modeling with Machine Learning and Satellite Data: A Case Study Estimating Tree Canopy Height in African Savannas

**Esther Rolf**

*Department of Computer Science  
University of Colorado, Boulder  
Boulder, CO 80309, USA*

ESTHER.ROLF@COLORADO.EDU

**Lucia Gordon**

*John A. Paulson School of Engineering and Applied Sciences  
Harvard University  
Allston, MA 02134, USA*

LUCIAGORDON@G.HARVARD.EDU

**Milind Tambe**

*John A. Paulson School of Engineering and Applied Sciences  
Harvard University  
Allston, MA 02134, USA*

TAMBE@SEAS.HARVARD.EDU

**Andrew Davies**

*Department of Organismic and Evolutionary Biology  
Harvard University  
Cambridge, MA 02138, USA*

ANDREW\_DAVIES@G.HARVARD.EDU

**Editor:** Pradeep Ravikumar

## Abstract

While advances in machine learning with satellite imagery (SatML) are facilitating environmental monitoring at a global scale, developing SatML models that are accurate and useful for local regions remains critical to understanding and acting on an ever-changing planet. As increasing attention and resources are being devoted to training SatML models with global data, it is important to understand when improvements in global models will make it easier to train or fine-tune models that are accurate in specific regions. To explore this question, we design the first study that explicitly contrasts local and global training paradigms for SatML, through a case study of tree canopy height (TCH) mapping in the Karingani Game Reserve, Mozambique. We find that recent advances in global TCH mapping do not necessarily translate to better local modeling abilities in our study region. Specifically, small models trained only with locally-collected data outperform published global TCH maps, and even outperform globally pretrained models that we fine-tune using local data. Analyzing these results further, we identify specific points of conflict and synergy between local and global modeling paradigms that can inform future research toward aligning local and global performance objectives in geospatial machine learning.

**Keywords:** geospatial machine learning, machine learning and remote sensing, machine learning for ecology, tree canopy height, savanna

## 1. Introduction

Paired with a massive amount of publicly available satellite imagery, geospatial machine learning is enabling a new paradigm of environmental monitoring. Machine learning with satellite imagery (SatML) models have been deployed at a global scale to classify patterns as diverse as land cover types (Brown et al., 2022), industrial activity at sea (Paolo et al., 2024), micro-estimates of wealth (Chi et al., 2022), and forest and tree canopy height (Potapov et al., 2021; Lang et al., 2023; Tolan et al., 2024). While efforts like these significantly extend the possibilities for *global-scale* science and policy, understanding *local* trends and variations in key environmental indicators remains paramount. Many important indicators of global change in ecology and environmental sciences manifest at local or regional scales, which global-scale data do not always represent well. Moreover, data-driven policies designed to mitigate global change and its effects (e.g., policies that regulate land use, or distribute resources across populations) are often tailored to particular localities (e.g., administrative regions or protected areas).

Prior work has exposed concerns about the accuracy of geospatial machine learning models developed for global- or national-scale prediction when evaluated in local areas. Kerner et al. (2024) found that global models of land cover maps are sub-optimal at capturing regional land cover variation of agricultural land in Sub-Saharan Africa. Aiken et al. (2023) showed that satellite-based poverty prediction maps can exhibit policy-relevant regional biases across urban and rural lines when trained on entire countries. Nyandwi et al. (2024) highlighted limitations of global building and roads data sets evaluated across diverse neighborhoods in Rwanda. Similarly, Gevaert et al. (2024) found substantial gaps between local and global performance of building segmentation maps from SatML models, with implications for disaster response. While it is tempting to expect that advances in mapping environmental phenomena across the globe will extend to improved model performance in specific local regions, these examples suggest that this may not be the case.

While past studies have exposed accuracy disparities in key local regions, they largely stop short of investigating the root causes and larger implications of these differences. This leaves open several important questions about the suitability of current methods in SatML across local and global scales, as well as how to orient future methodological innovations to address the observed accuracy disparities: Are drops in accuracy in some regions due to those regions being inherently difficult or unique places, or are they symptoms of a deeper discrepancy between the capacities of global and local models? When should we expect a model trained with an abundance of globally representative data will be accurate for a local region? And if local data are available for training, when is fine-tuning a globally pretrained model a better strategy than training a local model from scratch? Towards addressing these critical and practical questions, we seek to develop a deeper fundamental understanding of when discrepancies arise between local and global modeling objectives.

In this work, we devise a set of experiments to better understand if and when global mapping and local mapping use cases constitute somewhat distinct goals for SatML and geospatial machine learning more generally. We conduct a case study of local ecological importance, the mapping of tree canopy height (TCH), which is also representative of a class of local modeling problem settings more generally. Through this case study, we assess the degree to which integrating SatML models and maps designed for global performance

improves the capabilities of local models, compared with a strategy that trains a model “from scratch” based only on local training data. Our findings provide concrete evidence of a divide between local and global modeling performance in practice, and facilitate a nuanced understanding of how data and model design decisions affect this gap.

In our grounding case study, we develop locally-focused models of TCH using publicly available satellite imagery and spatially restricted UAV-LiDAR data from Karingani Game Reserve, Mozambique (Figure 1). Recent global prediction maps of TCH built on globally available satellite data from the Global Ecosystem Dynamics Investigation (GEDI) mission have achieved remarkable performance in mapping variation in tree canopy height across the globe, enabling the monitoring of large-scale change in vegetation extent and structure (Potapov et al., 2021; Lang et al., 2023; Tolan et al., 2024; Pauls et al., 2024). At the same time, local variation in TCH is pivotal for studying vegetation dynamics across landscapes. Local variation in TCH has a strong influence on spatial heterogeneity, which directly affects biodiversity through the creation of diverse ecological niches (Davies and Asner, 2014; Coverdale and Davies, 2023), predator-prey interactions (Loarie et al., 2013; Davies et al., 2016) and spatial patterns of aboveground carbon storage (Aponte et al., 2020; Noulèkoun et al., 2021; Tetemke et al., 2021). As such, there is also a separate line of work in using SatML to expand the spatial range of high-resolution TCH data that has been collected in local regions using Airborne Laser Scanning (ALS) techniques (Wilkes et al., 2015; Astola et al., 2021).

## 1.1 Research Questions and Contributions

Despite the increased use of SatML to map tree canopies at both global and local scales, there is a lack of scholarship connecting and contrasting the efforts of local modeling and global modeling for TCH mapping. This is indicative of a larger gap in research aimed at understanding how to balance (or distinguish between) the goals of local and global usability in geospatial machine learning more generally. Developing such an understanding is crucial to making efficient and sustainable progress in geospatial machine learning for environmental monitoring—from developing global models that can be efficiently fine-tuned for local use cases, to understanding when expensive local data-collection efforts are most needed, and how best to leverage local data among other possible data sources in SatML models.

To better understand the common and distinct facets of local and global mapping paradigms with current methods, our first research question (RQ) is: **(RQ1) How much can we improve the performance of local prediction by using locally collected reference data and/or existing globally (pre)trained models and maps?** The results of our case study indicate that local data are critical for generating ecologically relevant, accurate models in the Karingani Game Reserve region. Our locally-trained models reduce mean absolute error of existing global TCH maps by 40-69%. Similar performance increases extend to measures that are stratified by ecologically-relevant gradients – including true canopy height, distance to rivers, and geology. Surprisingly, using globally-pretrained models to initialize model weights for local fine-tuning did not substantially improve model performance compared to a random initialization. Overall, the best performing model of those we compare for our task is a model trained *only* with local data.

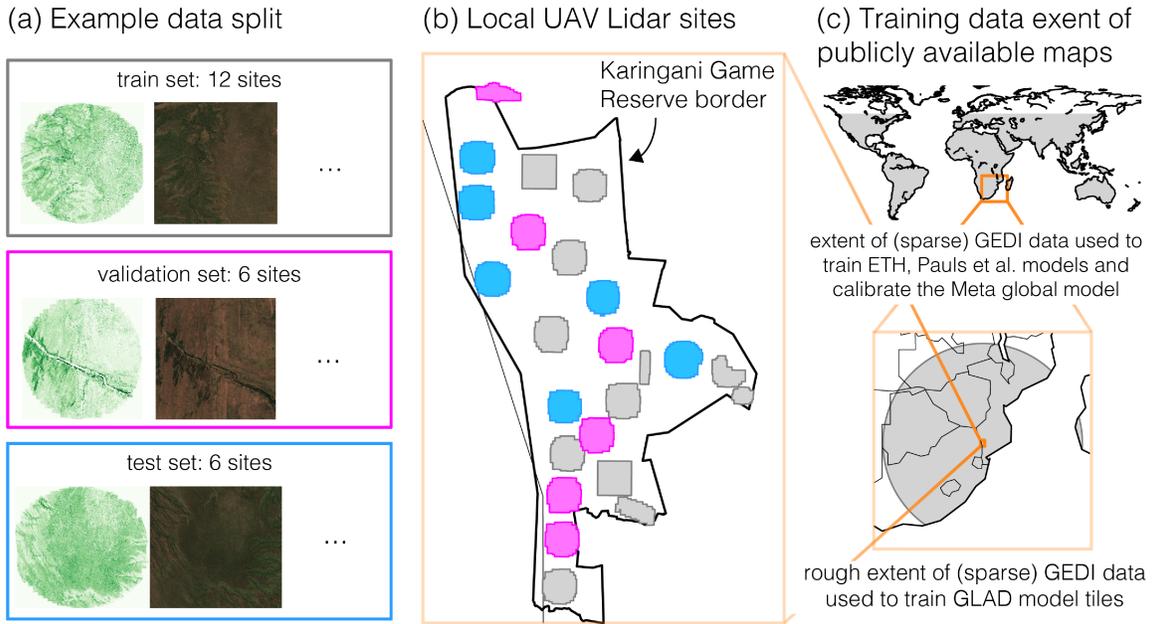


Figure 1: Our case study simulates a common local mapping use case: we use locally-collected tree canopy height (TCH) maps derived from data collected via UAV-mounted LiDAR sensors at sites distributed across Karingani Game Reserve, Mozambique to train a predictive model, then deploy it throughout the study region. (a) The local supervised learning problem consists of high-resolution TCH labels derived from LiDAR measurements (green), paired with 12-channel Sentinel-2 imagery (3-channel visual imagery is shown here) in each site. (b) There are 24 sites where UAVs were flown to collect LiDAR data, which we use to generate the local label data. Colors indicate which sites are allocated to the training (gray), validation (pink), and test (blue) sets for a given split (split 0 pictured here). Splitting the data by site simulates only having data from the training and validation sets, using the spatially disjoint test set sites to estimate the performance that would be achieved in the other parts of the region. (c) Contextualizes our local study area within a global map, highlighting the rough extents of the label data, from the Global Ecosystem Dynamics Investigation (GEDI), that are used as labels to generate the global TCH maps we compare to.

While RQ1 addresses the importance of obtaining local data to build SatML models that perform well locally, the question of how to build accurate local models remains. Developing a locally-focused predictive model involves many design decisions, including: which model architecture to use, which imagery sources and spectral layers to use as model input, and possibly even how much local data should be collected, and where. In order to understand where the gains from new methods and techniques will be greatest, our second research question is: **(RQ2) Which factors of data and model design most influence reliability of local TCH models?** We design experiments to compare the relative importance

of: (i) the quantity and composition of training data sampled across ecological gradients (*geographic representation*), (ii) spectral layers of the satellite imagery (*feature representation*), and (iii) different machine learning architectures (*model representation*), on the performance of our local models. We find that these factors all have a large influence on model performance, with nuanced interactions between factors. Importantly, the amount and distribution of labeled training data has as large an influence on model performance as the model architecture chosen, despite these data-centric design choices receiving relatively little attention compared to model developments in SatML.

Our case study of TCH mapping in the Karingani Game Reserve is emblematic of local mapping settings that arise across many environmental monitoring domains. Towards identifying common challenges across SatML applications, our third research question asks: **(RQ3) What are some specific points of conflict or synergies in local and global modeling efforts that are likely to manifest across application domains in SatML?** Examining the performance of recently proposed global TCH models, we find that recent advances in global performance do not necessarily confer performance improvements in our study domain (in fact, some methods that improve on global performance substantially degrade local performance). While we are careful not to draw absolute conclusions from our case study from a single region, we identify specific design decisions made in past global mapping efforts that may have hindered local usability in favor of global performance, and highlight other decisions that seem to explicitly align the efforts of global and local performance.

The results of our study have implications for current best practices and promising future directions to improve possible synergies between local and global monitoring efforts in geospatial machine learning across prediction domains. Connecting our findings to recent scholarship in other SatML tasks, we find that several indicators of tensions identified in our case study are consistent with the global-local accuracy disparities evidenced in other SatML domains (including crop classification, population density mapping, poverty mapping, and building segmentation). We summarize the key findings of our case study in Section 6, discussing the connections to recent work that suggest the gap between local and global performance objectives may be a pervasive one across prediction tasks in environmental monitoring. We then extract forward-looking takeaways for geospatial machine learning grounded in the results of our case study in Section 7, highlighting the potential of data-centric approaches and modeling paradigms that explicitly align local and global objectives.

## 2. Related Work

We divide related work into machine learning and geospatial machine learning literature on global and local prediction, and work on mapping tree canopy height.

### 2.1 Local vs. Global Prediction in General and Geospatial Machine Learning

The distinction between learning representations that capture local variation and those that best capture global variation has long been established, for example concerning notions of “local” and “global” learning paradigms across general manifolds (Zhou et al., 2003; Saul and Roweis, 2003; Huang et al., 2008). There is a corresponding discrepancy in output (prediction) space as well. A model trained to optimize average performance across an

entire population or distribution can fail to capture meaningful variation within sub-regions or sub-populations due to limited model capacity or finite training data (Rolf et al., 2022), or because the model can mainly predict relative values across sub-groups of a population, but cannot reliably distinguish variation within sub-groups (Aiken et al., 2023).

Past studies have evidenced that the local-global distinction may be especially relevant for *geospatial machine learning*, where the spatial nature of predictions across the globe makes the distinction between local and global prediction particularly clear. Kerner et al. (2024) evaluate existing global and regional land-cover maps in Sub-Saharan Africa, finding that global maps perform significantly worse than maps that were developed for regional use. This phenomenon extends to national-level maps as well; Yeh et al. (2020) produced national models to predict economic well-being from satellite imagery, but found that they could resolve national variation much better than variation within urban or rural areas. This result was replicated across many countries by Aiken et al. (2023). Discrepancies between local and global accuracy of SatML predictions have also been highlighted in the environmental monitoring domains of population estimation (Kuffer et al., 2022), building mapping (Gevaert et al., 2024), road detection (Nyandwi et al., 2024) and tree canopy height mapping (Pauls et al., 2024). With the increasing focus on global modeling efforts and foundation models in SatML, it is important to understand not just where errors in published models are distributed across the globe, but also to more fundamentally understand what aspects of geospatial modeling might differ for a global training and evaluation objective, compared to a more localized one.

## 2.2 Mapping Tree Canopy Height with Satellite Imagery and Machine Learning

Several studies use machine learning with satellite imagery inputs to extend the spatial coverage of LiDAR-based canopy height maps in different local study regions. One thread of methodology uses locally collected LiDAR data, e.g. vegetation height from airborne LiDAR scanning (ALS), as the training labels for these predictive models. For example, Wilkes et al. (2015) used random forest models and publicly available satellite imagery to extend the range of ALS maps in east Victoria, Australia, and Astola et al. (2021) used Sentinel-2 data to generate TCH maps in Finland. Our methodology for “local mapping” (described in Section 4) roughly follows this type of approach.

Four recent studies produce global tree canopy height maps using satellite data and machine learning. While all four maps are aimed at global coverage, each mapping effort has distinct goals that informed their design decisions:

- Potapov et al. (2021) produced a 30m resolution map of canopy height across the globe for the year 2019 using a time series of Landsat satellite images. Features derived from pixelwise time series of Landsat data and digital elevation maps are used as input to bagged regression tree models, using GEDI<sup>1</sup> data for label supervision. To

---

1. While GEDI data are sparse and somewhat noisy compared to ALS or UAV-derived LiDAR data, they have the advantage of global coverage. Relevant to our study region, Li et al. (2023) assess the suitability of the GEDI RH98 data product (designed for dense forests) to African savannas around South Africa, finding that the RH98 data are reliable for canopy heights between 3-15m, with limited ability to detect shrubs below 3m.

capture regional trends, Potapov et al. (2021) train separate random forest models for each Global Land Analysis and Discovery (GLAD) Analysis Ready Data (ARD) tile ( $1^\circ \times 1^\circ$  region), using satellite data from the surrounding  $12^\circ$  radius (see Figure 1 for reference). Following prior convention, we refer to this as the “GLAD” map.

- Lang et al. (2023) produced a 10m resolution global canopy height model for the year 2020 using 12-channel Sentinel-2 images concatenated with 3 channel transformations of locations (lat, lon) as input and GEDI labels as supervision. A key goal of this effort was to improve model performance for tall-height canopy regions, compared to the previous global map. Lang et al. (2023) use a modified XceptionS2 convolutional neural network architecture. Unlike the moving window approach of Potapov et al. (2021), Lang et al. (2023) use a single approach globally, which consists of an ensemble of several models that is used to simultaneously output estimates of TCH and estimates of uncertainty across the globe. Models in the ensemble take as input Sentinel-2 imagery concatenated with a 3-dimensional embedding of latitude and longitude position. Following prior convention, we refer to this as the “ETH” map.
- Tolan et al. (2024) produced a higher-resolution (1m) global canopy height model using 3-channel Maxar imagery spanning the years 2018-2020. To achieve their goal of a high spatial resolution TCH map, they pretrained a vision transformer model with dense ALS data from two regions (California, USA, and São Paulo, Brazil) then used global Maxar satellite data to extend this model across space, using GEDI data to calibrate the pretrained model for regions across the globe. Following prior convention, we refer to this as the “Meta” map.
- Most recently, Pauls et al. (2024) produced a global TCH map at 10m resolution. Noting a gap in performance between regional and global TCH maps, Pauls et al. (2024) focused on global-scale prediction that can be competitive with local regional maps. They devised a shift-resilient loss function to account for specific structures of the GEDI measurements and included elevation data as input (in addition to Sentinel-1 and Sentinel-2 data) to improve performance in mountainous regions. We refer to this map as the Pauls et al. (2024) map.

The design decisions made in producing these existing global TCH maps reflect their common goal: to produce accurate *global* maps of tree canopy height. In contrast, in this study, our aim is to (i) assess the degree to which these globally-focused maps and models can aid in the goal of producing accurate *local* maps, and (ii) compare how the design decisions for global ML models and maps might differ from the design decisions for local models and maps.

Most related to the local modeling setting of our case study, two existing studies use GEDI data to train and calibrate local TCH models. Relevant to our research questions, Healey et al. (2020) studied which spatial ranges are best for calibrating local models. Using data from across the globe, they found that local models trained on data from just a 3km radius outperformed models that were trained on data from larger radii. More recently, Tsao et al. (2023) developed local (country-level) models in Nigeria, focusing on tree height mapping for smallholder palm plantations, where performance in low-height canopy cover is important. Comparing the two globally available maps at the time (GLAD and ETH),

Tsao et al. (2023) found that local calibration significantly outperformed both global maps in their study region. Similarly to Healey et al. (2020), they find that a relatively small radius is best for locally calibrating models for predicting short trees in their region of study. It is worth noting that both of these studies use random forest models. It remains to be seen how these findings would extend to larger machine learning models that incorporate spatial information in imagery.

### 3. Case Study Problem Setting: Mapping Tree Canopy Height in Karingani Game Reserve

Karingani Game Reserve is an ecologically diverse area on the border of Mozambique and South Africa, spanning roughly 145,000 hectares. The Karingani Game Reserve was formally established in 2008 and is comprised primarily of subtropical savanna vegetation varying widely in height and density across gradients of topography, geology, soil type and mean annual rainfall. The reserve also contains a full suite of indigenous herbivore species. In addition to being an area of ecological importance, Karingani is an ideal region for a case study examining how best to map TCH in a local environment. Since the high-resolution LiDAR-based tree canopy height data we use for our case study (described in Section 3.1) has not been used to-date to train or evaluate machine learning models, our case study can simulate different possible avenues and opportunities for producing a new map of TCH in a local region.

Our case study examines how to map tree canopy height across the Karingani Game Reserve using the different resources available (a limited amount of locally collected high-resolution TCH labels (Section 3.1), publicly available satellite imagery (Section 3.2), and existing maps and models (Section 2.2) that have been produced for mapping TCH at a global scale). While we are of course interested in generating accurate maps of this region, the primary focus of our study is to understand more generally which of these current resources are most helpful for creating accurate local maps, and to identify additional resources or techniques that might aid in local mapping efforts.

#### 3.1 Local LiDAR-derived Tree Canopy Height Labels

Our local label data used for model training and evaluation span a study area of 24 sites in the Karingani Game Reserve (Figure 1b), collected in May and early June 2021; 22 sites span roughly 21 km<sup>2</sup> and 2 span roughly 7 km<sup>2</sup>. The 1m resolution reference TCH maps are derived by applying a canopy height model (as detailed in Boucher et al. (2023); Reiner et al. (2023)) to LiDAR sensor data. LiDAR data were collected using the Harvard Animal Landscape Observatory (HALO), a lightweight sensor package assembled by Phoenix LiDAR Systems (Austin, Texas, USA) for integration on an unoccupied aerial vehicle (UAV), which for this study was a FreeFly Alta X. HALO includes a high-resolution LiDAR scanner (a Riegl VUX-1LR), which communicates directly with a centralized navigational computer that utilizes a high-precision MEMS IMU (Sensoror STIM300) and integrated GPS to co-locate each LiDAR data point. The GPS system is capable of post-processing kinematics (PPK) and real-time kinematics (RTK) to maximize absolute accuracy. To improve the accuracy of the UAV’s trajectory in post-processing, we also collected differential GNSS reference data using a StoneX S900 GNSS receiver as a base station. After processing, the

LiDAR data had an absolute horizontal accuracy of  $< 1$  cm and a vertical accuracy of  $< 2$  cm. Comparisons between field- and LiDAR-measured TCH for the HALO system have previously been made in neighboring Kruger National Park, where the two measurements were closely aligned (Singh et al., 2023). The LiDAR data collection via UAV flight took roughly 1 day per site, including travel time.

In contrast to the globally available (but sparse) GEDI data that past studies have used to build and extend TCH models, our local LiDAR-derived TCH labels are highly accurate ground-referenced data. By comparison, while our data exhibit  $< 1$  cm horizontal errors, GEDI data used in past studies exhibit geolocation errors of up to 15-20 m for the GEDI L2B product used in the GLAD model (Potapov et al., 2021), or a few meters for the GEDI L1B product used in the ETH model (Lang et al., 2022). Thus, we consider our local label data to be a valuable resource for evaluating model performance in this region. Moreover, understanding the performance gain achievable when training with highly accurate local TCH data will help inform the value of acquiring local LiDAR data by flying UAVs or other types of aircraft in additional targeted regions in the future.

### 3.2 Extending the Range of Existing Label Data with Machine Learning and Sentinel-2 Satellite Data

While the LiDAR-derived data spans sites across the Karingani Game Reserve, there are still significant gaps in the coverage of the existing TCH data (roughly two thirds of the Karingani Game Reserve are not covered by past UAV surveying efforts). In our case study, the goal is to train machine learning models that take publicly available satellite imagery as input to extend the range of the existing TCH data.

As input to our locally-trained machine learning models, we downloaded three Sentinel-2 satellite data tiles captured at a date near to the collection of the LiDAR data (5/13/21), which span the study area and have low cloud cover percentage over the 24 sites. Since the Sentinel-2 data has roughly 10m spatial resolution, we aligned these images with coarsened 10m-resolution LiDAR-derived TCH using cubic resampling. We chose to restrict our local modeling efforts to mainly 12-band Sentinel-2 satellite imagery for simplicity of comparison. Past global-mapping approaches have used different input imagery and layers: Potapov et al. (2021) use Landsat imagery, Lang et al. (2023) append location embeddings to Sentinel-2 imagery, Tolan et al. (2024) use proprietary Maxar imagery, and Pauls et al. (2024) use Sentinel-1 and Sentinel-2 data and digital elevation maps. We investigate the effect of input layer choices on model performance in Section 5.3.

### 3.3 Site-stratified Training and Evaluation Splits Simulate Conditions of Model Use Outside Training Data Sites

To reflect our case study setting of applying the trained SatML models to fill in gaps where high-fidelity TCH data are not available, we employ a spatially stratified train/validation/test split of our data. Specifically, we split the 24 training sites into three distinct sets: training (12 sites), validation (6 sites) and testing (6 sites), as depicted in Figure 1a,b. For each split, no test sites are seen during the training step or the model validation step in which we pick model parameters and hyperparameters.

Employing these spatially separated train/validation/test splits simulates training our model in regions where we have data, then deploying our model to a separate set of sites that were not used for training or validating the model. An alternative method might assign pixels to train/test/splits within each site, but this would very likely overestimate the performance of our model when applied to new areas within the study region (Rolf, 2023). Thus, blocking our assignment of pixels to train, validation, and test set by data collection site ensures that the performance we evaluate more faithfully represents our intended use case of producing models that generalize across the entire local region in the Karingani Game Reserve. We construct four such spatial data splits, where each LiDAR collection site appears in exactly one test set and exactly one evaluation set. By analyzing performance across the four splits, we can assess the variation among the models across different data settings, characterized by the available training data and target regions for model use.

## 4. Methods

### 4.1 Methods Overview

Our experiments in Section 5 will reflect a range of different possible methodologies for predicting tree canopy height (TCH) in a local region. Since several global TCH maps are available, perhaps the most straightforward way to produce a map of TCH in a local area is to take one of the existing global TCH maps, crop and align it to the extent of our study area, and use that subset of the global map as our predictions. This is our primary methodology to “evaluate existing global maps” in Section 5.1.

An alternative approach would be to use the limited local labels to train a new model specifically adapted to the local region. Here, we distinguish between “local-only” models, which we train from scratch using only labels and satellite imagery from the local region (here Karingani Game Reserve), and “globally-pretrained, locally-adapted” (or locally “fine-tuned”) models, which we fine-tune using local data only, but where the model is initialized from a representation that has been pretrained on global data. We use the term “global pretraining” to encompass both supervised models trained to predict TCH across the globe as well as self-supervised models pretrained on global satellite imagery in a way that is agnostic to possible downstream tasks. In Section 5.2, we use our local data (Section 3.1) to train local-only models and fine-tune existing supervised and self-supervised models that have seen global satellite data, to compare the effectiveness of both approaches. Ultimately, contrasting the performance of using existing global TCH maps “out of the box,” adapting globally trained models using local data, and training models using only local data will enable us to address our first research question regarding the importance of local data for local prediction.

To address our second research question regarding the importance of different design choices for local prediction, we construct a set of experiments in which we systematically vary different data and modeling decisions and measure their effects on local performance. Specifically, in Section 5.3, we train local models using different amounts of training data to assess the importance of the abundance of local data in training or adapting models for local use. We train models with different input layers (Sentinel-2 bands and location encodings) as a way to estimate the relative effect that different data input decisions can

have on model performance. To contextualize the relative effects of these data decisions in practice, we compare the range of performance under these different data conditions to the range of performance due to changing the machine learning model architecture.

Our third research question involves identifying potential points of cohesion or distinction between the goals of local and global modeling in geospatial machine learning. Toward this end, we leverage our experimental results in this case study and pull in context from recent scientific reports across several application domains in geospatial machine learning for environmental monitoring. Contextualizing our in-depth analysis of local modeling from our case study amidst a broader set of trends in geospatial ML, in Section 6.3, we pinpoint specific design decisions as well as general trends in geospatial machine learning that we believe warrant more research attention toward understanding and bridging the divide between global and local modeling.

## 4.2 Implementation Details

The following subsections describe the details of our experimental methodology. Code is available at [https://github.com/estherrolf/tree\\_mapping/](https://github.com/estherrolf/tree_mapping/).

### 4.2.1 MACHINE LEARNING MODEL ARCHITECTURES

We compare multiple neural network model architectures throughout our experiments: (1) an XceptionS2 convolutional neural network architecture used in the ETH global predictions (Lang et al., 2023), which has 8 residual blocks each with two convolutional layers with 256 filters, (2) a U-Net architecture with a ResNet-18 backbone designed for pixel segmentation tasks, which takes the entire input (here a  $64 \times 64$  crop of the satellite image) as context for the predictions, and (3) a small five-layer fully convolutional network (FCN) with 128 filters per layer and leaky ReLUs between layers. See Table 3 for a comparison of the number of parameters for each model. We additionally compare to (4) non-parametric pixel-level forest models, where the input is the multiple channel input for that particular pixel. This pixelwise model does not incorporate any texture in the images, thus it constitutes a simple and informative baseline for machine learning with remotely sensed data.

### 4.2.2 MODEL TRAINING

Our experiments compare the strategies of training models only with local data and using local data to fine-tune models that have been initialized through pretraining with global data. When we train models only with local data (using no labels or imagery outside the study area), we randomly initialize model weights and then optimize the model using supervised learning. For the models with globally pre-trained weights available<sup>2</sup> (the ETH and U-Net models), we use these weights (which were derived using global data), and then fine-tune the models using local data. For the XceptionS2 globally pretrained weights, we use the published weights corresponding to the first model (model 0) in the ensemble from Lang et al. (2023), available at the GitHub repository accompanying their paper. Since this model was already trained in a supervised learning setting for the task of TCH

---

2. We note that while the Meta model is published for use, fine-tuning it would require purchasing proprietary satellite imagery. Pretrained weights from Pauls et al. (2024) were not publicly available at the time we conducted our experiments.

prediction, we use a transfer learning strategy in which we tune only the last few layers of the network. Table 2 shows a small difference in training the last 2 versus the last 3 layers, so we stop there. For the U-Net model, we use the self-supervised learning for earth observation (SSL4EO-S12) weights from Wang et al. (2023), which have been pretrained on global Sentinel-2 imagery using self-supervised learning techniques. We experiment with freezing the pretrained encoder weights and fine-tuning just the decoder versus fine-tuning the entire network for the U-Net.

All of our models are trained to optimize a mean squared error (MSE) objective, defined pixelwise. For each model configuration and each training split, we pick the hyperparameter configuration and model weights that correspond to the best validation set loss (MSE) achieved at any point during training (up to a maximum of 200 training epochs). For each neural network model (trained from scratch or fine-tuned), we perform a hyperparameter validation step over the initial learning rate in  $[0.00001, 0.0001, 0.001, 0.01]$  and weight decay in  $[0.00001, 0.0001, 0.001, 0.01]$ . We use the AdamW (Loshchilov and Hutter, 2019) optimizer and a learning rate scheduler that reduces the learning rate on performance plateaus by a factor of 0.1 with a patience of 10 epochs. Our neural network models are trained using the torchgeo python package (Stewart et al., 2022). For the random forest models, we use the scikit-learn implementation of random forest regression (Pedregosa et al., 2011), and tune the maximum depth of the trees in  $[2, 4, 8, 16]$ , and the number of trees in the ensemble in  $[50, 100, 200]$ .

#### 4.2.3 DATA PREPROCESSING

To account for some anomalously high values in the LiDAR-derived TCH labels, we set any pixels greater than 30m to no-data (NaN) values. Pixels with  $TCH < 0m$  are set to 0. This only occurs in a minimal percentage of the overall data; 99.95% of the non-NaN 1m-TCH pixel values fall within 0-30m. Next, we coarsen the labels to 10m-resolution to be consistent with the resolution of Sentinel-2 imagery. When coarsening, we use the 90th percentile value of the 1m pixel values to maintain consistency with past work (Potapov et al., 2021; Lang et al., 2023).

We preprocess the satellite imagery input to all of our locally trained models by channelwise normalization to zero-mean, unit-variance per channel across the three Sentinel-2 tiles that span the 24 sites. When we initialize the model with weights from a previous training or pretraining procedure, we use the image preprocessing transforms used during that previous procedure, as recommended by Corley et al. (2024).

#### 4.2.4 EVALUATION

We evaluate models pixelwise, at 10m resolution unless otherwise stated. When performance is reported by site, we use the models trained with the split where this site is in the test set. When performance is reported by split, we aggregate all pixels in a split to compute the average. When a single performance metric is given, this represents the average across the four splits.

When evaluating performance of the existing global TCH maps, we exclude from evaluation any pixels that contain no-data values in the global predictions (for the GLAD map, this excludes pixels corresponding to water, snow, ice, or other no-data-valued pixels). Less

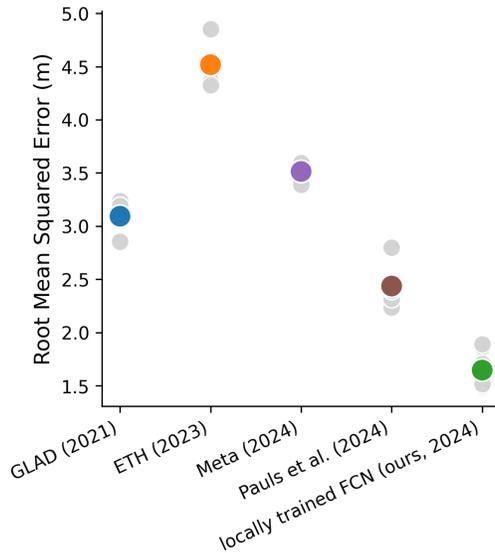


Figure 2: A locally-trained fully convolutional network (FCN) outperforms the three existing global maps in quantitative performance. Average performance is shown across splits (gray dots), and averaged over splits (larger, colored dots). Models on the horizontal axis are ordered by date of publication.

than 0.01% of the non-NaN pixels in the label data have a corresponding NaN in either of the global maps for both 10m and 30m-resolution, so this has a minimal impact on the reported performance metrics.

## 5. Results

Our experimental results evaluate performance in the context of existing global map and using existing models to initialize model training. We then evaluate how data choices affect model performance relative to other model design choices.

### 5.1 Evaluating Existing Global Maps

We first evaluate the performance of existing global TCH maps compared to our best locally trained or fine-tuned model, using both quantitative (Section 5.1.1) and qualitative (Section 5.1.2) assessments. We then assess the degree to which these different maps of TCH capture relative values of aboveground biomass in each site, as a way to gauge the usability of each map for additional ecological assessments in this local area (Section 5.1.3).

#### 5.1.1 QUANTITATIVE PERFORMANCE

Figure 2 summarizes the performance of the existing global models and our best local model (with extended results in Table 1). The existing global maps achieve average root mean squared error (RMSE) between 2.43m and 4.51m, and for all maps the variation in

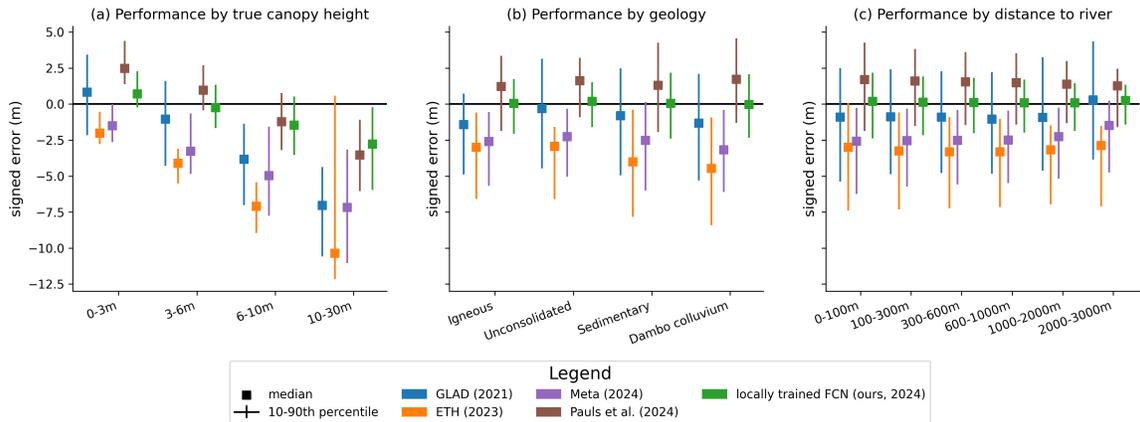


Figure 3: A locally-trained fully convolutional network (FCN) exhibits less prediction bias than the four existing global TCH maps, across ecologically-relevant strata: (a) binned tree canopy height, (b) geology type, and (c) binned distance to nearest river.

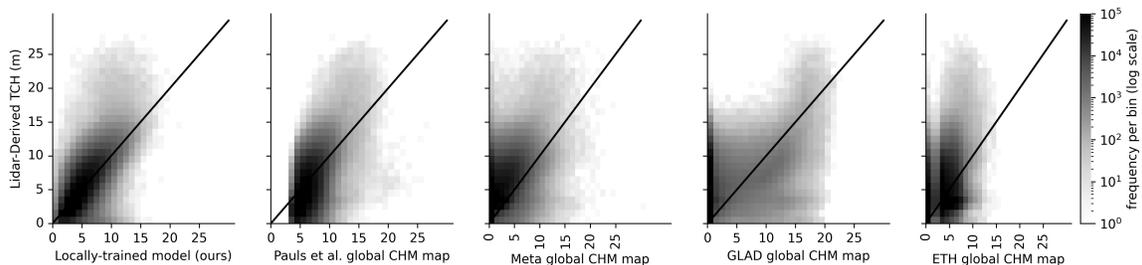


Figure 4: Distribution of LiDAR-derived ground-referenced measures of TCH (vertical axis) and predictions of each model (horizontal axis), binned to the nearest 1m. Distributions span all pixels over the 24 data collection sites (Figure 1) (for our predictions, we combine the four test split predictions). Models are ordered by average root mean squared error.

performance across data splits (grey dots) is generally smaller than the variation across models. The RMSE of these maps in Karingani is significantly lower than the RMSEs of each map evaluated across the globe (3.09m vs. 9.25m for Potapov et al. (2021); 4.51m vs 8.62m for Lang et al. (2023); 2.43m vs. 4.73m for Pauls et al. (2024)).<sup>3</sup> Note that for all models, the RMSE values in Karingani are lower than the global RMSEs. This indicates that the Karingani Game Reserve is not a region of especially high error for the global models, in comparison to other places around the globe.

3. These global metrics are reported through a common evaluation procedure in Pauls et al. (2024), which differ slightly from the original papers. For the purposes of our study, the relative rankings of global performance and local performance is most relevant.

While the global performance generally increased with each subsequent global model that has been published, Figure 2 shows that some of the more recently published maps perform worse in Karingani than older maps. Until the map by Pauls et al. (2024) was released, the best existing TCH map for the Karingani region (by our quantitative metrics) was the 30m resolution map produced by Potapov et al. (2021). We discuss implications of this result for understanding when and how efforts to increase global mapping performance can align with increasing local performance (our RQ3) in Section 6.3.

Our locally-trained model is superior to the existing global maps by a large margin. The predictions of our best-performing model (local FCN) achieve, on average, a 64% lower RSME than the ETH map (4.51m  $\rightarrow$  1.64m), a 47% lower RMSE than the GLAD map (3.09m  $\rightarrow$  1.65m), a 53% lower RMSE than the Meta map (3.51m  $\rightarrow$  1.65m), and a 33% lower RMSE than the Pauls et al. (2024) map (2.43m  $\rightarrow$  1.65m) in our study area. The gray dots in Figure 2 show the performance for each split individually, showing that these differences are consistent across different data splits.

Figure 3 disaggregates performance differences across ecologically relevant strata, including the true TCH values, geology, and distance to the nearest river corresponding to each pixel (see Appendix B for more details of the stratification process for these evaluations). Across these strata, our locally-trained models consistently show less bias (measured with average signed error) compared to the publicly available global models. The existing global models exhibit different biases among strata, with all models tending to under-predict pixels with higher canopy height. All models have similarly consistent prediction bias across geology and distance to rivers, with the model of Pauls et al. (2024) tending to over-predict TCH, and the remaining existing models tending to under-predict TCH. Our locally trained FCN model exhibits a median of around 0 bias, despite the out-of-sample test set (Figure 1a,b) requiring some degree of spatial generalization.

### 5.1.2 QUALITATIVE PERFORMANCE

Figure 5 plots our predictions and the data from the four publicly available maps for three representative sites in our study area. From visual inspection, our local predictions are clearly a better match to the LiDAR-derived TCH values than any of the existing global maps. In contrast, the GLAD map tends to under-predict TCH, and generally has more spatially smooth predictions than the label data (the original resolution of the GLAD map is 30m). The ETH map also tends to predict zeros more frequently, but sometimes captures higher canopy heights (reliably capturing higher canopy heights was a goal of the ETH study (Lang et al., 2023)). For example in the Massingir site, the ETH predictions tend toward more extreme values than the label data convey (Figure 5, first row). The Meta predictions capture more fine-grained textural information better than the GLAD or ETH map, but tend to underpredict TCH overall. The map from Pauls et al. (2024) is the closest existing map to getting the average pixel values correct, but predictions appear to be biased toward the mean value, missing some of the high and low values (note also that Pauls et al. (2024) do not estimate any canopy heights under 3m). Figure 4 plots the distribution of predictions and labels per pixel, aggregated across the 24 sites, confirming that these trends hold broadly across our study area.

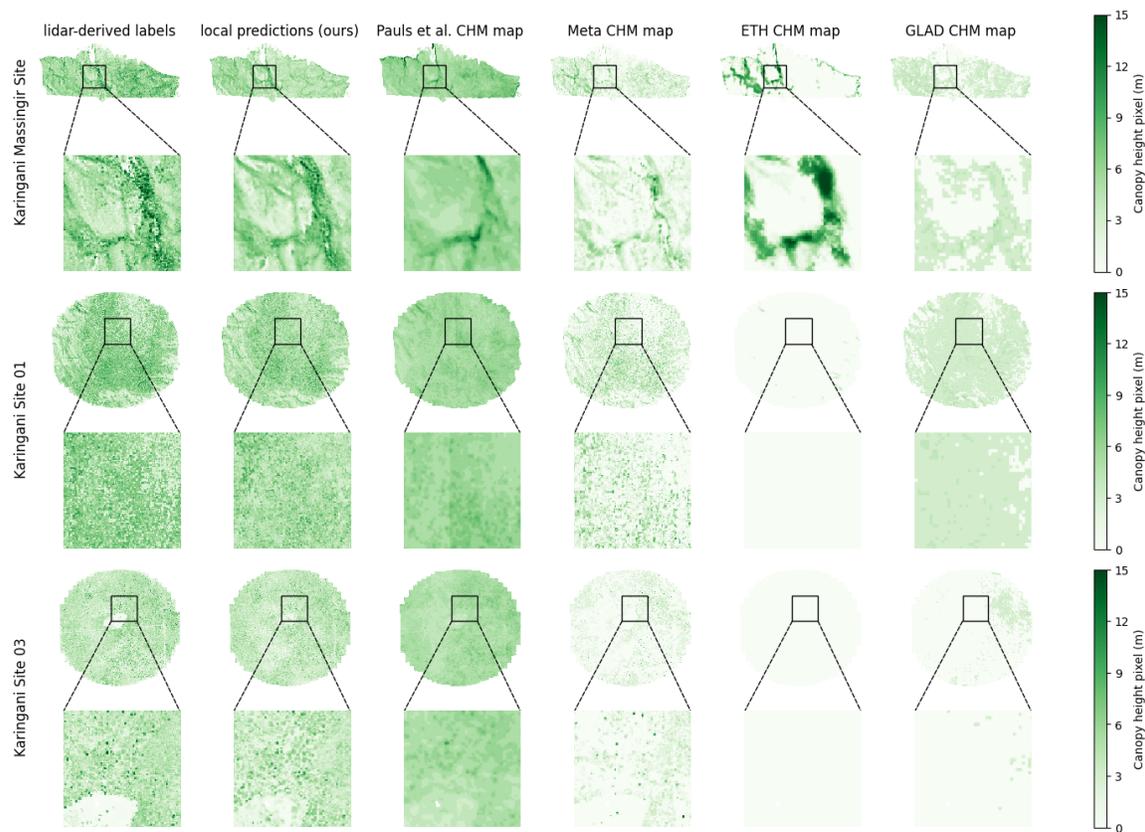


Figure 5: Existing TCH maps exhibit different types of visual error structures, which are largely alleviated by our local model. TCH labels (aggregated to 10m resolution) derived from locally collected LiDAR data, our model trained only with local labeled data (5-layer FCN with 12-channel Sentinel-2 imagery), compared with the Pauls et al. (2024), Meta (Tolan et al., 2024), ETH (Lang et al., 2023), and GLAD (Potapov et al., 2021) global TCH maps at 10m resolution. Existing maps are ordered by publication date, most to least recent.

### 5.1.3 GAUGING USABILITY FOR DOWNSTREAM ECOLOGICAL ANALYSES

To probe the degree to which these TCH maps could be used “out of the box” to assess other ecological spatial patterns related to tree canopy monitoring, we assess the degree to which each map could be locally calibrated to predict aboveground biomass (AGB). Reliable maps of AGB have high relevance to multiple ecological and conservation issues, including estimating carbon stocks and flux. For example, assessment of the impact of fire on savanna carbon, which has wide-ranging policy and management implications, relies on accurate measurement of AGB at fine spatial scales (Zhou et al., 2022).

Using established allometric equations available in the ecology literature (Colgan et al., 2013), we estimate AGB (in kg) for each site using the 1m LiDAR-derived TCH values and the linear-affine formula on the average height ( $H$ ) of pixels designated as canopy cover

( $CC$ ).<sup>4</sup> Canopy cover is defined as a binary variable valued at 1 for pixels with canopy height above a threshold of 1.5m—Colgan et al. (2013) use TCH maps with 30m resolution, while we use 1m resolution data for our reference values, and 10m resolution data for our predicted values, except for those derived from the Meta map, which has 1m resolution). We use the predicted values to estimate the same dependent variable  $X = \text{average}(CC \times H)$  at each site, for each map. We fit a least squares regression to associate these  $X$  values with the calculated AGB values for each site, as a local calibration procedure would do.

The  $R^2$  scores (coefficient of determination) resulting from the linear-affine fits (summarized in Figure 8) are 0.11 (ETH), 0.22 (GLAD), 0.52 (Meta), 0.41 (Pauls et al. (2024)), and 0.52 (our locally trained FCN). Overall, our locally trained FCN predictions with 10m resolution provide a signal of the relative AGB across the 24 study sites that is similar to what we would get from the 1m resolution Meta map, after a linear calibration. The map from Pauls et al. (2024) captures a slightly lower degree of relative variation in AGB across sites, whereas the low  $R^2$  values of the fit from the ETH and GLAD maps suggest that these maps are not as well suited to estimating relative AGB across sites (at least not using this established calibration formula).

## 5.2 Training Local Models From Scratch vs. Initializing With Weights From Global (Pre)Training

We next examine whether existing global SatML models can be a better starting place from which to fine-tune models with local data, versus training from scratch. Table 1 compares the performance of locally fine-tuned models initialized with weights from existing global models, versus initialized randomly.

Table 1 is organized according to two different strategies for adapting globally initialized models for local use: transferring a global supervised model for local use, and fine-tuning a task-agnostic self-supervised SatML model. We use the first strategy to adapt the XceptionS2 models—trained by Lang et al. (2023) for supervised learning of global TCH with GEDI LiDAR data as the original labels—for our local use. When using this pretrained model as a starting point for local training, we deploy a transfer-learning strategy and fine-tune the last three layers of the model. We also train randomly initialized XceptionS2 models for comparison. We introduce a variant without the location embedding input layers, in case these layers hamper performance in our site-stratified setting, which includes some degree of spatial extrapolation. The results in Table 1 show that there are marginal differences in performance when using the pre-initialized weights versus training XceptionS2 models from scratch. Comparing within the set of XceptionS2 models we tested, the difference between the models initialized with the weights from the Lang et al. (2023) model and the best randomly initialized XceptionS2 model is  $<0.05\text{m}$  in RMSE and  $<0.02\text{m}$  in mean absolute error (MAE).

Our second type of local adaption starts with a U-Net model with a ResNet-18 backbone, which has been pretrained by Wang et al. (2023) using self-supervised learning on globally-distributed Sentinel-2 satellite imagery. We experiment with fine-tuning the entire model versus fine-tuning only the decoder layers. Table 1 shows that the largest difference in average performance between a pretrained model after fine-tuning and a model of the same

---

4. Note that the reference AGB data are estimated from ground-truth TCH labels, not direct measurements.

model description	data used		RMSE (m)		MAE (m)		$R^2$	
	local	global	average	std dev.	average	std dev.	average	std dev.
XCEPTIONS2								
global init.	✓	✓	<u>1.98</u>	0.03	<u>1.48</u>	0.02	0.37	0.02
random init. (no latlon)	✓	✗	<u>1.94</u>	0.03	<u>1.47</u>	0.02	<u>0.40</u>	0.02
random init.	✓	✗	2.05	0.04	1.55	0.03	0.33	0.03
U-NET								
SSL init., decoder fine-tuned	✓	✓	<u>2.05</u>	0.04	<u>1.54</u>	0.03	<u>0.32</u>	0.03
SSL init., everything fine-tuned	✓	✓	<u>2.08</u>	0.13	<u>1.55</u>	0.10	<u>0.30</u>	0.09
random init., everything trained	✓	✗	<u>2.27</u>	0.69	<u>1.59</u>	0.20	<u>0.10</u>	0.69
FULLY CONVOLUTIONAL NETWORK								
5-layers, 128 filters	✓	✗	<b><u>1.64</u></b>	0.01	<b><u>1.20</u></b>	0.01	<b><u>0.57</u></b>	<0.01
PIXELWISE RANDOM FOREST								
parameters vary by split	✓	✗	<u>2.28</u>	<0.01	<u>1.80</u>	<0.01	<u>0.16</u>	<0.01
GLOBALLY AVAILABLE TCH MAPS								
GLAD (Potapov et al., 2021)	✗	✓	3.09	-	2.44	-	-0.53	-
ETH (Lang et al., 2023)	✗	✓	4.51	-	3.83	-	-2.26	-
Meta (Tolan et al., 2024)	✗	✓	3.51	-	2.87	-	-0.98	-
Pauls et al. (2024)	✗	✓	<u>2.43</u>	-	<u>2.01</u>	-	<u>0.05</u>	-

Table 1: Performance of fine-tuning globally pretrained models vs. training from scratch, measured by root mean squared error (RMSE), mean absolute error (MAE), and coefficient of determination ( $R^2$ ). Bold denotes the overall best model for each metric, underline indicates the best model of each model architecture type. The standard deviations reported are over 10 training runs with different random seeds. For ties (when average values are within 2 standard errors), we underline all tied models. The XceptionS2 model with “global init.” uses weights from Lang et al. (2023) (model 0). For all three XceptionS2 models, only the last three layers are tuned. The U-Net models use a ResNet-18 backbone, and “SSL init.” denotes that model weights are initialized using the SSL4EO self-supervised learning method for Sentinel-2 satellite imagery (using data distributed across the globe) from Wang et al. (2023). The fully convolutional network (FCN) is described throughout the paper as “ours,” as this is our best performing local-only model.

architecture trained from scratch is  $<0.23\text{m}$  in RMSE, and  $<0.06\text{m}$  in MAE. We note the higher variability (across random seeds during training) in the performance of the U-Net models compared to the other model architectures.<sup>5</sup>

Overall, the most substantial performance difference over the publicly available TCH maps (last four rows of Table 1) comes from using local data in any capacity. Our best locally trained model improves performance compared to the best globally produced data product by  $0.79\text{m}$  (a 33% reduction in error) in average RMSE and  $0.81\text{m}$  (a 40% reduction in error) in average MAE. We note that even our worst locally fine-tuned models from Table 1 improve performance compared to the best globally produced data products. Under conditions where local data are used, the marginal differences in performance due to initializing models randomly or with pretrained weights are very small. In fact, the best overall performing model is still our small FCN model trained only with local data, with an improvement of  $0.34\text{m}$  (a 17% reduction in error) in average RMSE and  $0.28\text{m}$  (a 19% reduction in error) in average MAE compared to the best of our globally initialized, locally fine-tuned models.

### 5.3 Relative Importance of Design Decisions Across Modeling and Data Choices

The results detailed in Sections 5.1 and 5.2 show that models tailored for local prediction can significantly improve the performance of TCH models in local regions. In light of this, we investigate the extent to which several different design decisions a research team faces in designing local models affect the overall performance. One goal of these experiments is to investigate the relative value of future efforts ranging from innovating on the machine learning models, choosing which satellite data layers and sources to use as inputs to the model, and even deciding how much labeled data to collect, and where to collect them.

Figure 6 summarizes an experiment in which we vary one of the following design decisions at a time: which machine learning model architecture is used to define the model (Figure 6a), which data layers are used as input to the model (Figure 6b), and the amount of local labeled data available for training the supervised models (Figure 6c). We find that each design decision has a substantial effect on overall performance. Interestingly, we find that the variation in model performance (RMSE) due to any of these factors is within roughly the same range; note the shared vertical axis range in Figure 6.

When it comes to choice of model architecture, we found that simple convolutional models can perform well at local prediction tasks (Figure 6a). Our best performing model is a small FCN model, with 128 filters per layer, and 5 layers, amounting to a total of 604,417 parameters (Table 3), and a receptive field of only  $11 \times 11$  pixels. This is the model we use throughout (with 12-channel Sentinel-2 inputs) as “locally-trained FCN.” The larger XceptionS2 and U-Net models had slightly worse performance, for both the randomly initialized and the globally pre-initialized variants of each model.

To understand how the availability of local data affects model performance, we take different random samples of the training sites for each of 3, 6, and 9 total sites, and compute the average model performance after training with those subsetted training sets, using our

---

5. This was in part due to model instability toward the outer edges of the image patches – we attempted to mitigate this effect during evaluation by using only the innermost  $4 \times 4$  pixels of each  $64 \times 64$  moving window patch when tiling our predictions as maps, though we found that the FCN and XceptionS2 models were still generally better suited to this task for our training procedure.

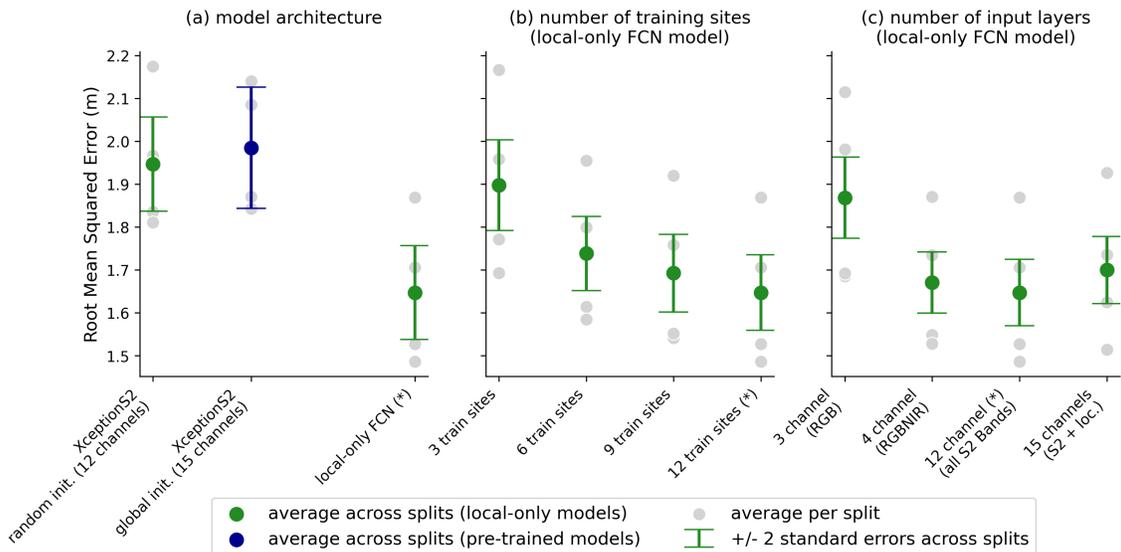


Figure 6: Different design decisions have similar magnitude of effect on local model performance. Performance varies across: (a) different machine learning architectures, (b) amount of training data (number of training sites) available, and (c) different sets of spectral bands used as input to the model. Here S2 stands for Sentinel-2. Note that model conditions with asterisks (\*) are all the same model, for reference across panels (a-c).

best model (locally trained FCN). To reduce the computational cost of this analysis, we fix the model hyperparameters to those used to train the local-only FCN model for each split used in the previous analyses. The total number of instances seen per epoch is kept consistent across conditions. The model performance resulting from using a different number of training sites is shown in Figure 6b. As expected, performance increases on average as the number of sites used for training increases, which is consistent across splits. The rate of returns of the average and per-split performances suggests that having even more training data would likely result in even better model performance.

Choices made about the satellite imagery also impact performance (Figure 6c). For our overall best performing model architecture, performance is worst when using just the red, green, blue (RGB) spectral bands of the image, while average RMSE decreases by 10% with the addition of a fourth near-infrared (NIR) band (Figure 6c). Including the remaining 8 Sentinel-2 bands has a small but positive effect on performance. Adding 3-channel location embeddings as in Lang et al. (2023) also did not change performance greatly for our locally-trained FCN, and if anything decreased average performance. This was also the case for the locally fine-tuned XceptionS2 model (Table 1). The absolute differences in model performance across the panels in Figure 6 contextualize the impact of different design decisions that must be considered in building a SatML model for use in local areas. Overall, the spread in performance due to these different decision factors – which model architecture to use, which data layers to use, and how much training data to

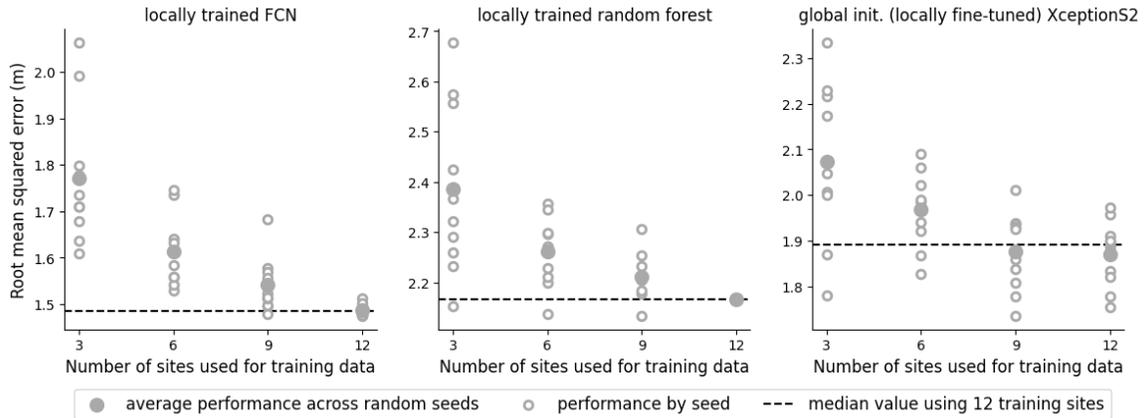


Figure 7: Data composition and model choice interact to affect model performance (train/val/test split 0 shown here; similar trends hold across splits—see Figure 9). For each model, 10 random trials were run with [3,6,9,12] distinct training sites. The same 10 random subsets of training sites are used for each condition to facilitate comparison across models.

use or acquire – is similar in magnitude, as indicated by the similar range of performance within each panel of Figure 6. The performance gap (difference in RMSE averaged across splits) due to using our second-best locally trained model (XceptionS2) versus our best (FCN), is 0.30m, while the performance gap due to using 3-channel (instead of 12-channel) imagery with the FCN is 0.22m. Likewise, the amount of training data impacts performance significantly; the difference between using 3 sites of training data versus 12 results in an average performance gap of 0.25m.

While each of these design decisions impacts local model performance in a distinct way, it is also important to understand if and how these data and modeling choices interact. Figure 7 summarizes the result of ten training runs with different subsets of 3, 6, 9 or all 12 training sites across different model architectures for a single train/val/test split (the split shown corresponds to Figure 1; the overall trends we describe are consistent across all four of the data splits). Figure 7 shows that the average performance generally degrades with fewer sites available for training, extending the results of Figure 6b across different model architectures. The magnitude of the degradation differs across models, as seen by comparing the trends in the filled-in gray dots in Figure 7 across each panel.

For all models, variation in model performance across different subsets (variation in each vertical set of open gray dots) tends to increase as the number of training sites decreases (Figure 7). For the XceptionS2 model, a substantial number of the training runs using fewer than 12 training sites actually outperformed the median performance using 12 sites. Similarly, for the locally trained random forest model, the best of the 10 runs for each number of training sites was better than the best performing model using 12 sites. Given the high variability in performance across each condition on models and number of training sites, it is not possible to conclude from this experiment alone whether the best-performing

runs with subsetted data have better performance because of the data or because of the random instantiation of the model. Examining the random forest models, where variation is low across random seeds for the 12-training-site condition, suggests that some of the performance improvement might be due to the change in training data sites for that model.

## 6. Discussion

Here we return to our key research questions, drawing on our experimental results.

### 6.1 The Value of Local Data for Model Training, Fine-tuning, and Evaluation

In response to our first key research question, we find that **locally collected label data remain pivotal for generating the best possible local maps** of tree canopy height (TCH) in our study region. Our first set of model comparisons (Section 5.1) examined the performance differences of models trained only with locally collected labels derived from UAV-LiDAR, versus using globally-available TCH maps derived from global (but sparse and noisy) GEDI data. While the existing global TCH maps have the benefit of being ready to use “off the shelf” without downloading any satellite imagery, training new models, or collecting local training data, their performance in the Karangani Game Reserve area was significantly worse than what we achieved with a small FCN model trained with *only* a limited amount of local, high-fidelity labeled data. Considering the potential downstream task of using TCH maps to estimate aboveground biomass, the 1m resolution Meta map (Tolan et al., 2024) performed on par with our local model, while the other existing maps performed worse in this use case. This raises the importance of validating global maps for local scientific analyses, even when those maps will be calibrated for local trends.

The intermediate performance of the locally trained pixel-wise random forest (with no spatial context) in Table 1 suggests that some, but not all, of the performance gap between models optimized for global use and models trained or fine-tuned for local use may be due to the calibration to local label distributions. Our second set of experiments (Section 5.2) shows the value of local data for fine-tuning globally pre-initialized models for use in a specific region. While the best overall models we trained only used local data, we still find that models initialized with weights from global training procedures can be substantially improved by using local data for fine-tuning (note, for example, the improvement in performance of the locally fine-tuned XceptionS2 model over the published XceptionS2 global map from Lang et al. (2023) in Table 1). This result underscores the potential for locally collected data to improve the regional performance of globally pretrained models, even as the usefulness of pretrained models for local use cases advances.

Lastly, **we emphasize the importance of the locally collected data for *evaluating* the performance of geospatial ML models**. Despite an increase in high-quality local and regional TCH maps being released recently (mostly in the Global North, e.g., Fogel et al. (2025); National Ecological Observatory Network (NEON) (2024)), understanding and assessing the local performance of global TCH maps across Earth’s diverse continents and biomes remains difficult due to a lack of globally-distributed test sites with high quality, locally collected TCH data. We expect that careful collection of local data will remain paramount to assessing and improving progress in geospatial machine learning for global environmental monitoring.

## 6.2 The Effects of Geographic-, Feature-, and Model-representation on Local Performance

In response to our second research question, we find that **the choices regarding (i) what model architecture to use, (ii) what data layers to use as input, and (iii) where label data are collected can all have a large impact on local performance.** Overall, there are substantial differences in performance across the model architectures we trained and fine-tuned with local data (Table 1). Notably, all of the models we used here have been proposed and used in prior works, and are not designed specifically for our local TCH modeling task. Thus, it is possible that local performance in areas like the Karingani Game Reserve could be further improved with modeling innovations—either by designing models specifically for local training, or by designing global models to be more suitable for downstream adaptation in smaller local areas.

From our experiments in Section 5.3, data choices had a large impact on local performance. For example, the choice of which spectral bands and geographic data layers to use as input also had a significant impact on performance. The inclusion of additional spectral bands (beyond 3-channel RGB) decreased the average error of our best model by 10%. For future researchers, decisions about which imagery to use should be based not only the spectral bands available, but also on the spatial and temporal resolution of the available imagery, which can be pivotal to model success. We expect that the difference in performance due to including all spectral bands of Sentinel-2 imagery is a conservative estimate of the performance differences that would arise from using different satellite products.<sup>6</sup>

We find that **the performance differences due to data design decisions were similar to the performance differences due to model architecture decisions.** The absolute difference in local performance (measured against our best model, which used all training sites and all Sentinel-2 input layers) is strikingly similar across the three different types of variations we tested: (i) using a model architecture from recent TCH-mapping research efforts, (ii) using only RGB spectral bands from the satellite imagery input, or (iii) using only 3 randomly selected training sites instead of 12. From these findings, we conclude that the quantity and composition of the available local training data sites affected model performance with an effect comparable to what we might expect from innovations in model architectures (as roughly reflected by the gap in performance between our best and second-best models in Figure 6a).

## 6.3 Identifying Points of Conflict and Synergy in Local and Global Mapping

Our third research question concerns identifying points of potential disparity or cohesion between the goals of local and global modeling in geospatial machine learning. From the outset, we expect different error rates in different locations (average performance in certain regions will differ from average global performance), as regional data distributions will inherently differ from the global data distribution. Still, we might hope that the existence of global models would make it easier to build good regional models using local data, in a similar way to how Tolan et al. (2024) use global GEDI data to calibrate their TCH models

---

6. The affect of the satellite data source will likely interact with other design choices; when evaluated in Karingani, the regionally trained map that used 30m resolution Landsat data (GLAD) outperformed the global map that used 1m resolution data (Meta).

and address distribution shifts across the globe (see Section 2.2). However, the results of our case study underscore that **currently, globally pretrained models are not necessarily the “right” starting point for training local models, even when high-quality local data is available for fine-tuning these models for specific regions.**

After extensive experimentation, existing globally trained models did not factor in to the best local modeling solution we found in our case study. At the same time, we expect that our local models would perform terribly across the globe, due to domain shifts and lack of global diversity in our training data. This discrepancy suggests that there are at least some points of conflict between local and global mapping endeavors. We caveat that since our case study concerns a single prediction task and study area (which allowed us to focus on the nuances of a specific local downstream use case), our discussion of RQ3—as it pertains to geospatial modeling across general tasks—is more speculative in nature than our discussion of our first two research questions. Interpreting the differences in performance from past global TCH mapping efforts in the context of a broader set of geospatial ML studies, we can still generate insights as to which modeling choices that influenced the differences in local performance in the Karingani Game Reserve are likely to generalize more broadly across different geospatial machine learning tasks.

Our results suggest that **the best decisions about which data to use during training may differ for local and global use cases**, both in terms of which regions to include training data from, and which input layers to use. Regarding the *regions that training data are drawn from*, we hypothesize that the relatively lasting success of the GLAD map (Potapov et al., 2021) is due in part to the fact that it is actually an accumulation of regional models applied across the globe. This is consistent with results of previous work (Healey et al., 2020; Tsao et al., 2023), which showed that there may be an optimal radius for the spatial distribution of training data that is large enough to constitute enough volume and diversity of training data, yet small enough to generate specialized, accurate local models. Regarding *input layers*, Lang et al. (2023) included positional encoding layers (derivatives of latitude and longitude) to achieve good global performance. However, for our local-mapping task, including positional encoding layers did not help with local fine-tuning performance, and possibly were a detriment to local performance (Table 1 and Figure 6). One possible reason for this result is that while direct positional encodings can help with spatial interpolation (e.g. filling in the gaps in sparse but global GEDI data), they may not be necessary over smaller spatial extents, or may even hinder performance in spatial extrapolation settings (e.g. extending the spatial range of local training data). Similar concerns have also been raised in other geospatial machine learning domains (Meyer et al., 2019; Lilja et al., 2024).

The differences in the performance of existing global TCH maps is consistent with the hypothesis that **modeling efforts that account for and correct local error structures in global labels may be an important step to synergizing local and global supervised learning efforts.** Only one of the three most recent global TCH maps (Pauls et al., 2024) improved upon the performance of the original 2021 GLAD global map, when evaluated in Karingani. In their study, Pauls et al. (2024) explicitly aimed to remedy gaps in local performance of global TCH models by accounting for error structures in the global GEDI data they used for training. Specifically, they accounted for localization error due to uncertainty in the position of each GEDI data point, and they removed observations from

high slopes (which make the satellite-based LiDAR measurements especially noisy) during training. The general strategy of accounting for and correcting error structures can be applied in many geospatial ML domains where some form of global labels may be available, but these labels likely come with localization noise, label noise, or statistical biases (Kuffer et al., 2022; Gevaert et al., 2024; Rolf et al., 2024).

## 7. General Takeaways for Geospatial Machine Learning

In addition to addressing our three key research questions as discussed in Section 6, the results from our case study provide three concrete takeaways with implications for future efforts in geospatial machine learning at different spatial scales.

*Takeaway 1: The best global models and the best local models may fundamentally differ in model design and training strategy.* Our results provide evidence that geospatial machine learning models that capture global variation well—as all four existing global maps of tree canopy height do—might not sufficiently capture local variation in the target variable of interest without significant local calibration or fine-tuning. Our results reinforce scholarship that found that SatML models trained in large areas often break down when evaluated in local regions, whereas locally-trained models have been evidenced to do much better. Specifically, the breakdown of global maps in local regions has been shown across domains including crop land classification (Kerner et al., 2024), population density estimation (Kuffer et al., 2022), building detection (Gevaert et al., 2024), and poverty mapping (Aiken et al., 2023).

Our best model for mapping TCH in the Karingani Game Reserve was a relatively small model trained only with local data (LiDAR-derived labels and Sentinel-2 imagery). This suggests that when training local models or pretraining global models for local use, it might be advantageous to use smaller models. At the same time, ultra-small models, such as a pixelwise random forest, obtained reasonable performance but performed worse than models that incorporate spatial context (e.g. texture) in images. Importantly, we found that design decisions aimed to enhance global modeling, such as including direct location embeddings as input to the model, may not be necessary for local models, and can actually hinder performance in local mapping settings.

Local and global mapping efforts are both crucial to informing scientific knowledge about the Earth and its environment. At the same, both efforts have significant costs, and it should not be assumed that global models will be sufficient for capturing local variation, without modification. This brings us to our next major takeaway, concerning the creation of ML models for environmental monitoring for use across both global and local scales.

*Takeaway 2: If geospatial machine learning models need to jointly satisfy global and local monitoring use cases, then novel research innovations may be needed.* It is tempting to expect that working to improve global performance for geospatial ML models will implicitly result in models that are more useful for local modeling tasks as well. However, none of the globally pre-trained methods we considered factored in to our best local modeling approach, even when fine-tuned with local data (Section 5.2). Given the substantial costs of training global models, and the additional substantial cost of collecting local data to calibrate global models (or train entirely local models), it is important to find synergy between the efforts made for global and local modeling with geospatial ML.

Our results have implications for how to build geospatial ML models with global data, especially if the intent is to fine-tune these models for local contexts. We found that smaller models tended to perform better in local areas, suggesting that model distillation methods (Polino et al., 2018; Vemulapalli et al., 2023) may hold promise for improved fine-tuning of large global models for local tasks. Another avenue to design models for local use cases might explicitly design the pretraining strategy for different local regions, similar to how meta-learning has been used to initialize geospatial ML models used for different prediction task domains (Rußwurm et al., 2024).

*Takeaway 3: Data-centric approaches may be a promising avenue for jointly achieving local and global performance with geospatial machine learning.* Several different aspects of our results indicate that data-centric innovations could be key to expanding the joint frontiers of local and global mapping efforts. As discussed in Section 6.3, the best of the existing global models in terms of both global and local performance mostly focuses on curating and preprocessing the training data (Pauls et al., 2024), while the next best model (the GLAD model of Potapov et al. (2021)) is actually an accumulation of regional models. Furthermore, our results in Section 5.3 indicate that even a very coarse and straightforward experimentation with different data conditions and processing strategies could lead to substantial increases in local performance. This suggests that innovation on data-centric approaches—which currently receive less attention than model-centric research—may be an important way to advance progress in SatML.

Many different approaches to data-centric geospatial machine learning exist and are being developed—see the rich discussion in Roscher et al. (2024). Adding to this discussion, our results underscore the importance of studying how best practices regarding data and model decisions in SatML might differ for local and global mapping paradigms. Better guidelines are needed for how best to collect ground-referenced data, which data source(s) to use for which tasks, and how best to preprocess global and local data for training geospatial machine learning models—across a variety of application domains and prediction settings ranging from global to hyperlocal.

## 8. Conclusion

In this paper, we studied perspectives of local and global mapping with machine learning and satellite imagery, rooted in a case study of tree canopy height mapping in the ecologically diverse Karingani Game Reserve in Mozambique. Tree canopy height prediction is a setting in which both global and local maps are pivotal to inform environmental and ecological science and policy. Global maps are crucial for understanding planetary ecosystem and tree composition, while local maps are crucial for scientific studies in landscape ecology and informing locally-enacted policies. The same argument—that both global and local maps are fundamentally needed—can be made for virtually any environmental or ecological variable of interest, from land cover classification to air quality estimation. Consequently, we need geospatial machine learning models that can resolve both large-scale global variation and fine-grained local variation in key environmental variables.

Our study focused on the task of tree canopy height mapping in Karingani Game Reserve, which has not been previously used to train or evaluate geospatial ML models. As such, these data enabled us to ground our study in considerations researchers and modelers

face when applying geospatial machine learning for ecological and environmental modeling to a new area. As discussed in Section 6, we expect that our main findings will generalize to different local areas and to different prediction tasks due to evidence of similar local-global map discrepancies in other geospatial prediction domains. That said, we expect that the exact numbers and nuanced insights may differ from setting to setting.

Our results indicate that, at present, global and local mapping may constitute distinct modeling goals for geospatial machine learning. This hypothesis is supported by previous findings across tasks including cropland mapping (Kerner et al., 2024), poverty prediction (Aiken et al., 2023), population density estimation (Kuffer et al., 2022), and building detection (Gevaert et al., 2024). Even the best global models for tree canopy height mapping, which represent impressive efforts in data processing and model architecture (Potapov et al., 2021; Lang et al., 2023; Tolan et al., 2024; Pauls et al., 2024) are outperformed by simple fully convolutional networks trained on roughly 250 square kilometers of carefully collected highly accurate label data paired with publicly available Sentinel-2 satellite imagery. Moreover, we find that some of the design decisions that past work found important for global modeling were either unnecessary or potentially harmful to local performance.

While we can draw many novel insights from our in-depth analysis in the Karingani Game reserve, the focus on a single region and prediction task is also the main limitation of our study. In order to draw more general insights, we discussed how our findings shed light on trends seen throughout various prediction tasks and regions in applied SatML in recent work. While we used many different model architectures to ensure our results were robust to modeling strategy, it remains to be seen how our findings generalize across prediction tasks (e.g. segmentation tasks), which vary in difficulty. In light of these limitations, it would be informative to repeat our study across many local areas and similar prediction tasks in future work (made possible with a growing number of public data sets representing mapping tasks related to trees (Weinstein et al., 2021; Reiersen et al., 2022; Ouaknine et al., 2025; Ahlswede et al., 2023; Bountos et al., 2023)).

Looking forward, there is a great opportunity to develop models that can flexibly interpolate between different prediction extents, ranging from global to hyperlocal. Our study exposes a need for future research directions that prioritize the performance of geospatial machine learning models in local areas alongside global performance. For example, globally pre-trained SatML models could be designed better for adaptation with high quality local data (labels, or even higher-quality imagery) for settings where those data are available. Distinguishing between the goals of local and global modeling is a first step to establishing best practices for model and design decisions for applications of geospatial machine learning for environmental monitoring. It is also an important step toward developing learning paradigms that enhance the synergy between the two efforts, for example, building models that flexibly leverage both local and global data depending on data availability and the prediction region at hand. In addition to synergistic local-global modeling efforts, an emphasis on the quality of local data for model evaluation and calibration remains critical for local, global, and synergistic model developments.

## Acknowledgments

The authors thank Nico Lang for answering questions that helped us implement the XceptionS2 models from Lang et al. (2023) and Konstantin Klemmer for feedback on an earlier version of this manuscript. Karingani Game Reserve is thanked for permission to collect LiDAR data, as well as for significant logistical support during data collection, with particular thanks to Ellery Worth for on the ground support. Pete Goodman is thanked for providing spatial data on rivers and geology. Tom Lautenbach and Michael Voysey are thanked for helping with data collection, and Jenia Singh and Peter Boucher for LiDAR data processing. Karingani Holdings are acknowledged for funding the LiDAR data collection.

E.R. was supported by postdoctoral fellowships from the Harvard Data Science Initiative and the Harvard Center for Research on Computation and Society. L.G. was supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE2140743.

## Appendix A. Additional Supporting Results

This appendix details supporting figures and tables referenced from the main text.

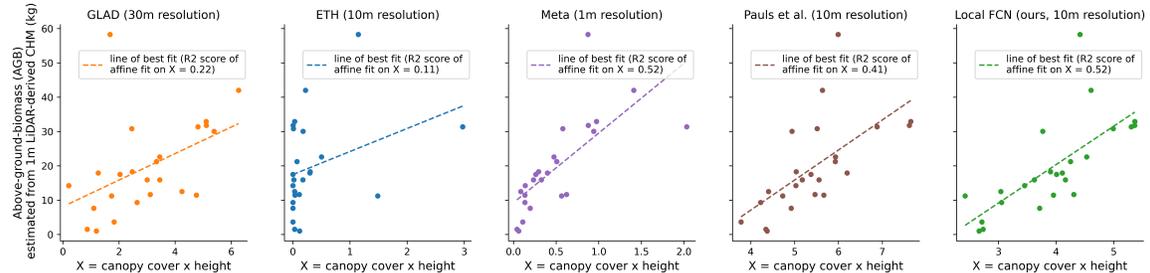


Figure 8: Aboveground biomass (AGB) per site, as estimated using the 1m LiDAR-derived TCH labels and the linear formula  $X = CC \times H$  from Colgan et al. (2013) (where  $CC = \text{canopy cover}$  and  $H = \text{height per pixel}$ , where reference values are computed on label data with 1m resolution). Specifically, equation F “PA harvest” from Colgan et al. (2013) is calculated for each site using our high-fidelity LiDAR-derived local data at 1m resolution, and plotted on the vertical axis. On the horizontal axis of each panel, we plot the same “ $X = CC \times H$ ”, this time estimated from each of the four different data products and our own predictive model. For the GLAD map, the ETH map, the map from Pauls et al. (2024), and the map from our local model, we use 10m resolution. The Meta map has 1m resolution. Dashed lines show the line of best fit describing the least-squares linear fit of that variable to the AGB values at each site calculated from our labels. The goodness-of-fit of each line is measured by the  $R^2$  score of a linear affine projection of the form  $X\beta_1 + \beta_0$ .

model description	layers tuned	RMSE		MAE		R2	
		average	std dev.	average	std dev.	average	std dev.
XceptionS2 (fine-tuned)	last 1	2.09	0.02	1.58	0.02	0.30	0.01
	last 2	1.97	0.03	1.47	0.02	0.38	0.02
	last 3	1.98	0.03	1.48	0.02	0.37	0.02
XceptionS2 (random init.) no latlon	last 1	2.21	0.02	1.72	0.02	0.22	0.02
	last 2	1.97	0.03	1.49	0.03	0.38	0.02
	last 3	1.94	0.03	1.47	0.02	0.40	0.02
XceptionS2 (random init.)	last 1	2.31	0.05	1.82	0.04	0.15	0.03
	last 2	2.08	0.05	1.58	0.04	0.31	0.03
	last 3	2.05	0.04	1.55	0.03	0.33	0.03

Table 2: Performance after fine-tuning the last 1, 2, or 3 layer(s) of an XceptionS2 model from Lang et al. (2023) with our local data. There are small gains to tuning 2 vs 3 layers. This extends the results from Table 1, which shows performance after tuning the last 3 layers in a transfer-learning approach.

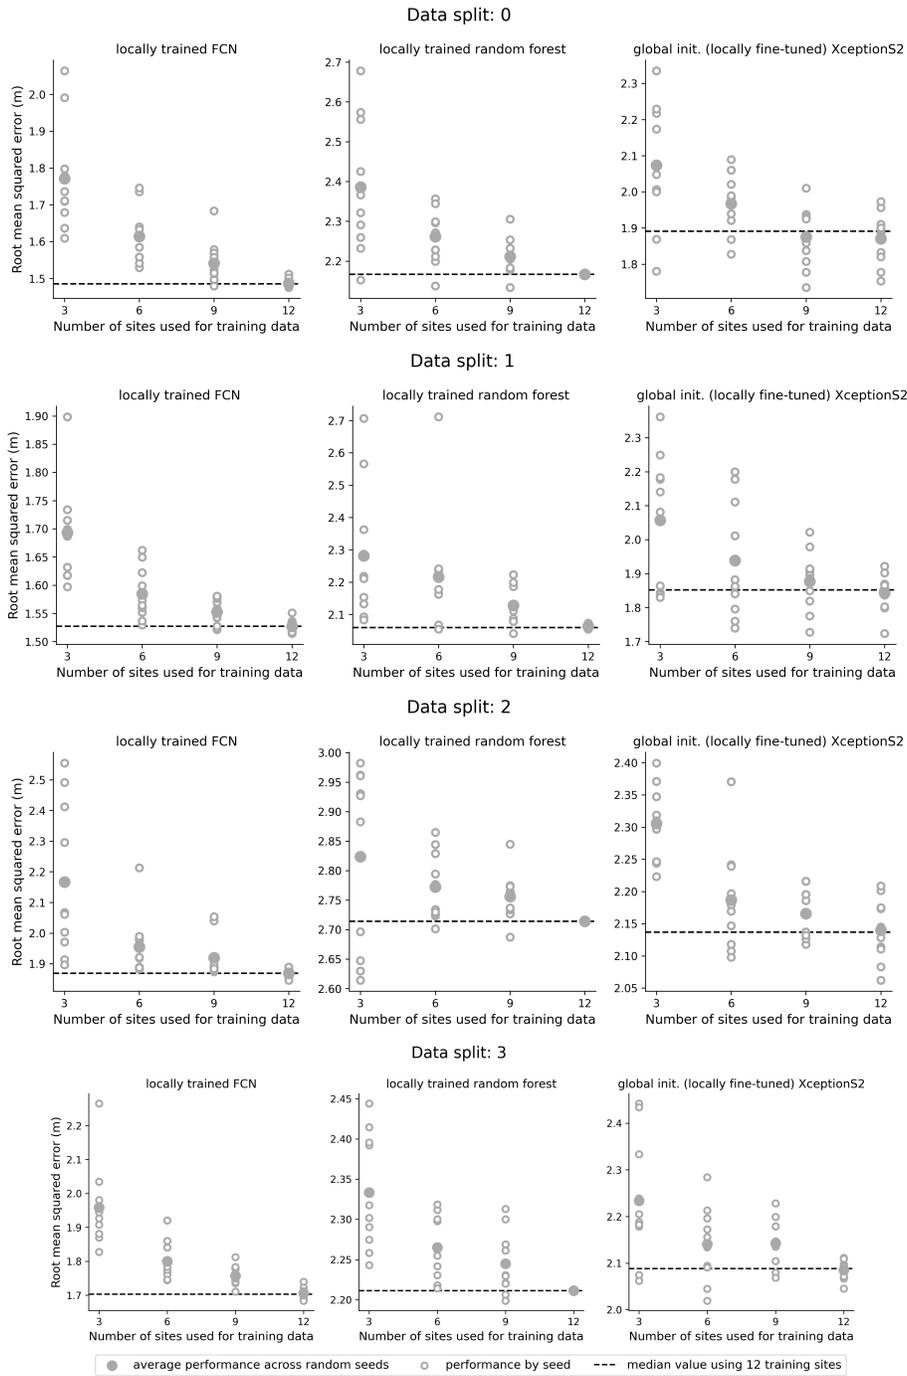


Figure 9: Results extending Figure 7 for all four data splits.

model description	layers tuned	trainable parameters	total parameters
XceptionS2	last 1	257	1,731,587
XceptionS2	last 2	203,265	1,731,587
XceptionS2	last 3	406,273	1,731,587
U-Net (ResNet-18 backbone)	decoder only	3,151,697	14,356,433
U-Net (ResNet-18 backbone)	all	14,356,433	14,356,433
5 layer FCN (“ours”)	all	604,417	604,417

Table 3: Number of total parameters and trainable parameters for the neural network models that we compare in Tables 1 and 2, assuming 12 input channels.

## Appendix B. Evaluation Procedure for Stratifying by Feature

This appendix describes our methodology generating features by which we stratify performance in Figure 3.

### B.1 Ground-truth Height

To stratify performance by the ground-truth canopy height, we calculate the residuals for the pixels whose label falls within the specified range.

### B.2 Distance to River

To stratify performance by distance to river, we start with a data product provided to us by researchers at the Karingani Game Reserve spanning the study area in vector format that has rivers encoded as lines. Note that these “rivers” correspond to paths through which water could flow but are not necessarily active rivers during all parts of the year. We rasterize this vector data and then for each site, we calculate the distance between each pixel in the site and every river pixel. We then take the minimum of those as the “distance to river” for that pixel. During evaluation we then calculate the residuals for the pixels whose minimum distance to any river falls within the specified range.

### B.3 Geology

To stratify performance by geology type, we start with a data product spanning the study area in vector format that has geology type encoded as polygons, provided to us by researchers at the Karingani Game Reserve. We rasterize this vector data and then for each site, we record the geology type for each pixel in the site. During evaluation we then calculate the residuals for the pixels whose geology type matches the specified category.

## References

- S. Ahlswede, C. Schulz, C. Gava, P. Helber, B. Bischke, M. Förster, F. Arias, J. Hees, B. Demir, and B. Kleinschmit. TreeSatAI benchmark archive: a multi-sensor, multi-label dataset for tree species classification in remote sensing. *Earth System Science Data*, 15(2):681–695, 2023. doi: 10.5194/essd-15-681-2023.
- Emily Aiken, Esther Rolf, and Joshua Blumenstock. Fairness and representation in satellite-based poverty maps: Evidence of urban-rural disparities and their impacts on downstream policy. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2023. ISBN 978-1-956792-03-4. doi: 10.24963/ijcai.2023/653.
- Cristina Aponte, Sabine Kasel, Craig R Nitschke, Mihai A Tanase, Helen Vickers, Linda Parker, Melissa Fedrigo, Michele Kohout, Paloma Ruiz-Benito, Miguel A Zavala, et al. Structural diversity underpins carbon storage in Australian temperate forests. *Global Ecology and Biogeography*, 29(5):789–802, 2020.
- Heikki Astola, Lauri Seitsonen, Eelis Halme, Matthieu Molinier, and Anne Lönnqvist. Deep neural networks with transfer learning for forest variable estimation using Sentinel-2 imagery in boreal forest. *Remote Sensing*, 13(12), 2021. ISSN 2072-4292. doi: 10.3390/rs13122392.
- Peter B Boucher, Evan G Hockridge, Jenia Singh, and Andrew B Davies. Flying high: Sampling savanna vegetation with UAV-lidar. *Methods in Ecology and Evolution*, 14(7):1668–1686, 2023.
- Nikolaos Ioannis Bountos, Arthur Ouaknine, and David Rolnick. FoMo-Bench: a multi-modal, multi-scale and multi-task forest monitoring benchmark for remote sensing foundation models. *arXiv preprint arXiv:2312.10114*, 2023.
- Christopher F. Brown, Steven P. Brumby, Brookie Guzder-Williams, Tanya Birch, Samantha Brooks Hyde, Joseph Mazzariello, Wanda Czerwinski, Valerie J. Pasquarella, Robert Haertel, Simon Ilyushchenko, Kurt Schwehr, Mikaela Weisse, Fred Stolle, Craig Hanson, Oliver Guinan, Rebecca Moore, and Alexander M. Tait. Dynamic World, near real-time global 10m land use land cover mapping. *Scientific Data*, 9(1), June 2022. ISSN 2052-4463. doi: 10.1038/s41597-022-01307-4.
- Guanghua Chi, Han Fang, Sourav Chatterjee, and Joshua E. Blumenstock. Microestimates of wealth for all low- and middle-income countries. *Proceedings of the National Academy of Sciences*, 119(3), January 2022. ISSN 1091-6490. doi: 10.1073/pnas.2113658119. URL <http://dx.doi.org/10.1073/pnas.2113658119>.
- Matthew S. Colgan, Gregory P. Asner, and Tony Swemmer. Harvesting tree biomass at the stand level to assess the accuracy of field and airborne biomass estimation in savannas. *Ecological Applications*, 23(5):1170–1184, 2013. doi: <https://doi.org/10.1890/12-0922.1>.
- Isaac Corley, Caleb Robinson, Rahul Dodhia, Juan M Lavista Ferres, and Peyman Najafirad. Revisiting pre-trained remote sensing model benchmarks: resizing and normalization

- matters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3162–3172, 2024.
- Tyler C Coverdale and Andrew B Davies. Unravelling the relationship between plant diversity and vegetation structural complexity: A review and theoretical framework. *Journal of Ecology*, 111(7):1378–1395, 2023.
- Andrew B Davies and Gregory P Asner. Advances in animal ecology from 3D-LiDAR ecosystem mapping. *Trends in Ecology & Evolution*, 29(12):681–691, 2014.
- Andrew B Davies, Craig J Tambling, Graham IH Kerley, and Gregory P Asner. Effects of vegetation structure on the location of lion kill sites in African thicket. *PLoS One*, 11(2):e0149098, 2016.
- Fajwel Fogel, Yohann Perron, Nikola Besic, Laurent Saint-André, Agnès Pellissier-Tanon, Martin Schwartz, Thomas Boudras, Ibrahim Fayad, Alexandre d’Aspremont, Loic Landrieu, et al. Open-Canopy: Towards very high resolution forest monitoring. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1395–1406, 2025.
- Caroline M Gevaert, Thomas Buunk, and Marc JC Van Den Homberg. Auditing geospatial datasets for biases: using global building datasets for disaster risk management. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.
- Sean P. Healey, Zhiqiang Yang, Noel Gorelick, and Simon Ilyushchenko. Highly local model calibration with a new GEDI LiDAR asset on Google Earth Engine reduces Landsat forest height signal saturation. *Remote Sensing*, 12(17), 2020. ISSN 2072-4292. doi: 10.3390/rs12172840.
- Kaizhu Huang, Haiqin Yang, Irwin King, and Michael Lyu. *Machine Learning: Modeling Data Locally and Globally*. Zhejiang University Press and Springer Berlin Heidelberg, Hangzhou and Berlin, 2008.
- Hannah Kerner, Catherine Nakalembe, Adam Yang, Ivan Zvonkov, Ryan McWeeny, Gabriel Tseng, and Inbal Becker-Reshef. How accurate are existing land cover maps for agriculture in Sub-Saharan Africa? *Scientific Data*, 11(486), 2024.
- Monika Kuffer, Maxwell Owusu, Lorraine Oliveira, Richard Sliuzas, and Frank van Rijn. The missing millions in maps: exploring causes of uncertainties in global gridded population datasets. *ISPRS International Journal of Geo-Information*, 11(7):403, 2022.
- Nico Lang, Nikolai Kalischek, John Armston, Konrad Schindler, Ralph Dubayah, and Jan Dirk Wegner. Global canopy height regression and uncertainty estimation from GEDI LIDAR waveforms with deep ensembles. *Remote Sensing of Environment*, 268: 112760, 2022.
- Nico Lang, Walter Jetz, Konrad Schindler, and Jan Dirk Wegner. A high-resolution canopy height model of the Earth. *Nature Ecology & Evolution*, 7, 2023.

- Xiaoxuan Li, Konrad Wessels, John Armston, Steven Hancock, Renaud Mathieu, Russell Main, Laven Naidoo, Barend Erasmus, and Robert Scholes. First validation of GEDI canopy heights in African savannas. *Remote Sensing of Environment*, 285:113402, 2023. ISSN 0034-4257. doi: <https://doi.org/10.1016/j.rse.2022.113402>.
- Adam Lilja, Junsheng Fu, Erik Stenborg, and Lars Hammarstrand. Localization is all you evaluate: Data leakage in online mapping datasets and how to fix it. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22150–22159, 2024.
- Scott R Loarie, Craig J Tambling, and Gregory P Asner. Lion hunting behaviour and vegetation structure in an African savanna. *Animal Behaviour*, 85(5):899–906, 2013.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of the International Conference on Learning Representations*, 2019.
- Hanna Meyer, Christoph Reudenbach, Stephan Wöllauer, and Thomas Nauss. Importance of spatial predictor variable selection in machine learning applications—moving from data reproduction to spatial prediction. *Ecological Modelling*, 411:108815, 2019.
- National Ecological Observatory Network (NEON). Ecosystem structure (dp3.30015.001), 2024. URL <https://data.neonscience.org/data-products/DP3.30015.001/RELEASE-2024>.
- Florent Noulèkoun, Emiru Birhane, Sylvanus Mensah, Habtemariam Kassa, Alemayehu Berhe, Zefere Mulaw Gebremichael, Nuru Mohammed Adem, Yigremachew Seyoum, Tefera Mengistu, Bekele Lemma, et al. Structural diversity consistently mediates species richness effects on aboveground carbon along altitudinal gradients in northern Ethiopian grazing exclosures. *Science of the Total Environment*, 776:145838, 2021.
- Emmanuel Nyandwi, Markus Gerke, and Pedro Achanccaray. Local evaluation of large-scale remote sensing machine learning-generated building and road dataset: The case of Rwanda. *ISPRS International Journal of Geo-Information*, 13(1):1–18, 2024.
- Arthur Ouaknine, Teja Kattenborn, Etienne Laliberté, and David Rolnick. Openforest: a data catalog for machine learning in forest monitoring. *Environmental Data Science*, 4:e15, 2025. doi: [10.1017/eds.2024.53](https://doi.org/10.1017/eds.2024.53).
- Fernando S. Paolo, David Kroodsma, Jennifer Raynor, Tim Hochberg, Pete Davis, Jesse Cleary, Luca Marsaglia, Sara Orofino, Christian Thomas, and Patrick Halpin. Satellite mapping reveals extensive industrial activity at sea. *Nature*, 625(7993):85–91, January 2024. ISSN 1476-4687. doi: [10.1038/s41586-023-06825-8](https://doi.org/10.1038/s41586-023-06825-8).
- Jan Pauls, Max Zimmer, Una M Kelly, Martin Schwartz, Sassan Saatchi, Philippe Ciais, Sebastian Pokutta, Martin Brandt, and Fabian Gieseke. Estimating canopy height at scale. In *Proceedings of the International Conference on Machine Learning*, pages 39972–39988, 2024.

- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. In *Proceedings of the International Conference on Learning Representations*, 2018.
- Peter Potapov, Xinyuan Li, Andres Hernandez-Serna, Alexandra Tyukavina, Matthew C. Hansen, Anil Kommareddy, Amy Pickens, Svetlana Turubanova, Hao Tang, Carlos Ed-ibaldo Silva, John Armston, Ralph Dubayah, J. Bryan Blair, and Michelle Hofton. Mapping global forest canopy height through integration of GEDI and Landsat data. *Remote Sensing of Environment*, 253:112165, 2021. ISSN 0034-4257. doi: <https://doi.org/10.1016/j.rse.2020.112165>.
- Gyri Reiersen, David Dao, Björn Lütjens, Konstantin Klemmer, Kenza Amara, Attila Steinegger, Ce Zhang, and Xiaoxiang Zhu. ReforesTree: A dataset for estimating tropical forest carbon stock with deep learning and aerial imagery. In *Proceedings of the Conference by the Association for the Advancement of Artificial Intelligence*, 2022.
- Florian Reiner, Martin Brandt, Xiaoye Tong, David Skole, Ankit Kariryaa, Philippe Ciais, Andrew Davies, Pierre Hiernaux, Jérôme Chave, Maurice Mugabowindekwe, et al. More than one quarter of Africa’s tree cover is found outside areas previously classified as forest. *Nature Communications*, 14(1):2258, 2023.
- Esther Rolf. Evaluation challenges for geospatial ML. *Machine Learning for Remote Sensing Workshop at the International Conference on Learning Representations*, 2023.
- Esther Rolf, Ben Packer, Alex Beutel, and Fernando Diaz. Striving for data-model efficiency: Identifying data externalities on group performance. *Workshop on Trustworthy and Socially Responsible Machine Learning at Neural Information Processing Symposium*, 2022.
- Esther Rolf, Konstantin Klemmer, Caleb Robinson, and Hannah Kerner. Position: Mission critical – satellite data is a distinct modality in machine learning. In *Proceedings of the International Conference on Machine Learning*, 2024.
- Ribana Roscher, Marc Rußwurm, Caroline Gevaert, Michael Kampffmeyer, Jefersson A Dos Santos, Maria Vakalopoulou, Ronny Hänsch, Stine Hansen, Keiller Nogueira, Jonathan Prexl, et al. Better, not just more: Data-centric machine learning for Earth observation. *IEEE Geoscience and Remote Sensing Magazine*, 2024.
- Marc Rußwurm, Sherrie Wang, Benjamin Kellenberger, Ribana Roscher, and Devis Tuia. Meta-learning to address diverse Earth observation problems across resolutions. *Communications Earth & Environment*, 5(37), 2024.
- Lawrence K. Saul and Sam T. Roweis. Think globally, fit locally: unsupervised learning of low dimensional manifolds. *Journal of Machine Learning Research*, 4:119–155, 2003. ISSN 1532-4435. doi: 10.1162/153244304322972667.

- Jenia Singh, Peter B Boucher, Evan G Hockridge, and Andrew B Davies. Effects of long-term fixed fire regimes on African savanna vegetation biomass, vertical structure and tree stem density. *Journal of Applied Ecology*, 60(7):1223–1238, 2023.
- Adam J Stewart, Caleb Robinson, Isaac A Corley, Anthony Ortiz, Juan M Lavista Ferres, and Arindam Banerjee. Torchgeo: deep learning with geospatial data. In *Proceedings of the International Conference on Advances in Geographic Information Systems*, pages 1–12, 2022.
- Buruh Abebe Tetemke, Emiru Birhane, Meley Mekonen Rannestad, and Tron Eid. Species diversity and stand structural diversity of woody plants predominantly determine above-ground carbon stock of a dry afro-montane forest in northern Ethiopia. *Forest Ecology and Management*, 500:119634, 2021.
- Jamie Tolan, Hung-I Yang, Benjamin Nosarzewski, Guillaume Couairon, Huy V. Vo, John Brandt, Justine Spore, Sayantan Majumdar, Daniel Haziza, Janaki Vamaraju, Theo Moutakanni, Piotr Bojanowski, Tracy Johns, Brian White, Tobias Tiede, and Camille Couprie. Very high resolution canopy height maps from RGB imagery using self-supervised vision transformer and convolutional decoder trained on aerial lidar. *Remote Sensing of Environment*, 300:113888, 2024. ISSN 0034-4257. doi: <https://doi.org/10.1016/j.rse.2023.113888>.
- Angela Tsao, Ikenna Nzewi, Ayodeji Jayeoba, Uzoma Ayogu, and David B. Lobell. Canopy height mapping for plantations in Nigeria using GEDI, Landsat, and Sentinel-2. *Remote Sensing*, 15(21), 2023. ISSN 2072-4292. doi: 10.3390/rs15215162.
- Raviteja Vemulapalli, Hadi Pouransari, Fartash Faghri, Sachin Mehta, Mehrdad Farajtabar, Mohammad Rastegari, and Oncel Tuzel. Label-efficient training of small task-specific models by leveraging vision foundation models. *arXiv preprint arXiv:2311.18237*, 2023.
- Yi Wang, Nassim Ait Ali Braham, Zhitong Xiong, Chenying Liu, Conrad M. Albrecht, and Xiao Xiang Zhu. SSL4EO-S12: A large-scale multimodal, multitemporal dataset for self-supervised learning in Earth observation [software and data sets]. *IEEE Geoscience and Remote Sensing Magazine*, 11(3):98–106, 2023. doi: 10.1109/MGRS.2023.3281651.
- Ben G. Weinstein, Sarah J. Graves, Sergio Marconi, Aditya Singh, Alina Zare, Dylan Stewart, Stephanie A. Bohlman, and Ethan P. White. A benchmark dataset for canopy crown detection and delineation in co-registered airborne RGB, LiDAR and hyperspectral imagery from the National Ecological Observation Network. *PLOS Computational Biology*, 17(7):1–18, 07 2021. doi: 10.1371/journal.pcbi.1009180.
- Phil Wilkes, Simon D. Jones, Lola Suarez, Andrew Mellor, William Woodgate, Mariela Soto-Berelev, Andrew Haywood, and Andrew K. Skidmore. Mapping forest canopy height across large areas by upscaling ALS estimates with freely available satellite data. *Remote Sensing*, 7(9):12563–12587, 2015. ISSN 2072-4292. doi: 10.3390/rs70912563.
- Christopher Yeh, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. Using publicly available satellite imagery and

deep learning to understand economic well-being in Africa. *Nature Communications*, 11 (2583), 2020.

Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *Advances in Neural Information Processing Systems*, volume 16, 2003.

Yong Zhou, Jenia Singh, John R Butnor, Corli Coetsee, Peter B Boucher, Madelon F Case, Evan G Hockridge, Andrew B Davies, and A Carla Staver. Limited increases in savanna carbon stocks over decades of fire suppression. *Nature*, 603(7901):445–449, 2022.