

A Neuro-Symbolic Framework for Tree Crown Delineation and Tree Species Classification

Ira Harmon, Ben Weinstein, Stephanie Bohlman, Ethan White, and Daisy Zhe Wang

Abstract

Neuro-symbolics combine deep learning and symbolic reasoning models to produce better performing hybrids. Not only do neuro-symbolic models perform better, but they also deal better with data scarcity, allow for the direct incorporation of high level domain knowledge into the model, and are more explainable. However, these benefits come at the cost of increased complexity that may deter the uninitiated from using these models. In this work, we present a framework to simplify the creation of neuro-symbolic models for tree crown delineation and tree species classification through the use of object oriented programming and hyperparameter tuning algorithms. We show that models created using our framework outperform their non-neuro-symbolic counterparts by as much as 2 F1-points for crown delineation and 3 F1-points for species classification. Furthermore, we show that the use of hyperparameter tuning algorithms allows the user to test multiple formulations of domain knowledge without the added overhead of manual tuning.

Index Terms

remote sensing; tree crown delineation; tree species classification; machine learning; neuro-symbolics

I. INTRODUCTION

Remote sensing is central to conducting efficient forest inventories on large spatial scales [1], [2]. To make use of the large volumes of data that remote sensing produces requires automated tools. The fundamental tasks of processing remotely sensed data are designed to answer two questions: where are the trees and what kind of trees are they? These two questions are answered by the tasks of crown detection (or delineation) and species classification. While some forest parameters, such as leaf chemistry [3], [4], can be estimated without identifying tree species,

knowing the species identity of trees allows for improved estimates of size and carbon related measurements as well as the ability to conduct biodiversity based research and conservation [5], [6].

Remote sensing-based tree crown delineation and species classification are inherently difficult problems. Crown boundaries can be hard to determine due to irregularly shaped crowns and overlapping adjacent crowns, especially in dense forests [7]. Depending on the resolution of the image, pixels that comprise the boundaries of trees may receive light from one or more tree canopies as well as light from nearby objects, understory vegetation or the ground [8]. Mixed pixels also make species classification more challenging [9]. Species classification is also made more difficult by features due to variation in crown traits, like leaf chemistry or crown leaf density, and their response to environmental soil nutrient or water availability [10]–[12].

Species may also have high spectral similarity to other species, for example from closely related species with similar shapes and crown densities [13], [14]. The difficulty in separating species increases with the number of species in a forest. Generally, species classification accuracy tends to be inversely related to the number of species considered for classification [15], [16]. Atmospheric contributions to the image, viewing geometry and shadows make both delineation and species classification more challenging [17]–[19].

Machine learning algorithms such as deep learning models are the state of the art in automated forest inventory methods for both crown detection/delineation and species classification [20]–[22]. Deep learning models are statistical models that can have billions of learned parameters that are usually trained on labeled datasets. Deep learning model’s performance on a task tends to improve as the number of learned parameters increase, but the size of the dataset needed for the model to be performant grows superlinearly with the number of model parameters [23], [24].

Convolutional neural networks (CNNs) are one of the most commonly used models for image-based crown delineation and species classification [25]. These models tend to contain on the order of millions of parameters, requiring large datasets to train, usually on the order of thousands to millions of images [23], [26]. Amassing and labeling large datasets is typically one of the highest hurdles to creating a useful model [24].

Another frequently cited drawback of deep learning models is their lack of explainability [27], [28]. These models are considered black boxes, where the reasoning behind their inferences cannot be easily determined by the user. In many instances, explainability is not required for the model to be useful, but there are situations when understanding the why of a model’s

predictions are important, particularly for critical decisions and to gain insight into the problem being researched.

Symbolic models are models that represent real world characteristics as variables linked by a series of operations. Common examples of symbolic models are logical formalism such as propositional logic, first order logic (FOL), or mathematical equations. Contrary to deep learning models, it is easy for a human to understand the reasons behind these models' predictions, however, these models do not typically generalize well. On the other hand, it is easy for humans to represent expert knowledge using logic formalisms, guaranteeing that the model captures high level concepts.

Neuro-symbolic models are deep learning model-symbolic model hybrids that are more robust than either model on its own [29], [30]. Neuro-symbolic models have been shown to still be performant with shrunken datasets, having proven themselves useful in zero-shot and few shot learning scenarios, and function more transparently to the user [31]. Neuro-symbolic models also have the advantage of injecting ecological knowledge into algorithms based purely on image data, analogous to how a field biologist uses their ecological knowledge, such as a habitat in which species are likely to occur, in addition to organismal features, to identify a species. This field of machine learning has been applied to several ecological problems. Xu et al. use a neuro-symbolic approach for fine grained image classification of birds [32]. Sumbul et al. use a neuro-symbolic model for zero-shot learning of tree species [33].

There are several mechanisms for creating neuro-symbolic models from deep learning models. Seo et al., use a regularization technique with a method similar to ensembling [34]. Hu et al. use posterior regularization and knowledge distillation [35]. Dilligenti et al. use a similar technique called semantic based regularization (SBR) to create neuro-symbolic models [36]. In this work we focus on SBR.

Semantic based regularization works by adding a regularization equation in the form of a fuzzy FOL expression to the model's loss function. Fuzzy FOL uses continuous functions as its operators making it differentiable, which is a required property for loss functions in models trained using some form of backpropagation, which is used by most neural models [37]. During training, the model is penalized for predictions that fail to satisfy the FOL encoded rule.

Though neuro-symbolic models are more robust, they are not without shortcomings. Turning domain knowledge into a rule or equation that can be incorporated into a machine learning model can be challenging, particularly for rules written in first order logic. Once a rule is formulated

in FOL, it must be written in a programming language, usually the same language used to program the neural model. Furthermore, after a rule is written, depending on the neuro-symbolic framework, there may be many hyperparameters that need to be tuned. Hyperparameters are model variables that are set by the user rather than learned [38]. The number of hyperparameters varies with the model used, but usually number in the tens. Introduction of FOL statements into the model adds hyperparameters that must also be tuned. Usually each rule has its own hyperparameters. When there are multiple rules, they may interact in unexpected ways that increase the complexity of tuning. Selecting values for rule hyperparameters that work well is a time consuming task.

In this paper we introduce a SBR-based neuro-symbolic framework paired with hyperparameter tuning algorithms as one solution to alleviating some of the difficulties in creating neuro-symbolic models. Our object-oriented approach provides rule templates that allow users to quickly model rules in FOL using object-oriented programming’s inheritance paradigm. Our framework includes 3 hyperparameter tuning algorithms - random search, grid search, and Bayesian optimization - to ease the burden of finding optimum rule parameters [39].

In the remainder of this paper we describe the architecture of our framework and illustrate its use on two ecologically relevant scenarios: crown delineation and species classification. For the crown delineation model we use DeepForest, a popular individual tree crown delineation model based on RetinaNet [40]. DeepForest is designed to detect tree crowns in remote sensing RGB images. To demonstrate a species classification use case we use a model from [41], a well cited paper by Fricker et al. The Fricker model can process both RGB and hyperspectral remote sensing data, but in this paper we focus on RGB. We use data from the National Ecological Observatory Network to train each model [42]. For the crown delineation model we use data from Niwot Ridge, an alpine forest in Colorado [43]. For the species classification model we use data from the Tea Kettle Experimental Forest, TEAK, a mixed-conifer forest in California [44]. Though we use these models and data to demonstrate the use of our framework, the framework is model agnostic and can be applied to any neural model that can be written in Python.

II. FRAMEWORK ARCHITECTURE

Our framework is written in Python version 3.11.6 primarily using PyTorch 2.0.1 and PyTorch Lightning 2.1.0 [45], [46]. Meta’s Bayesian optimization package Ax version 0.3.4 is used for Bayesian optimization [47]. We create our models using PyTorch Lighting, whose predefined

classes are easily extendable to any neural model. The code for our framework is available for download on Github at https://github.com/ihmn02/forest_ecology_neuro_symbolic_framework and is released for general use under MIT licensing. The majority of the code is written using the object oriented programming paradigm. The corresponding data used in each model is available from the NEON at <https://www.neonscience.org/data>, but can also be downloaded from Zenodo at <https://zenodo.org/records/14194555>. A block diagram of the framework's architecture is shown in Figure 1.

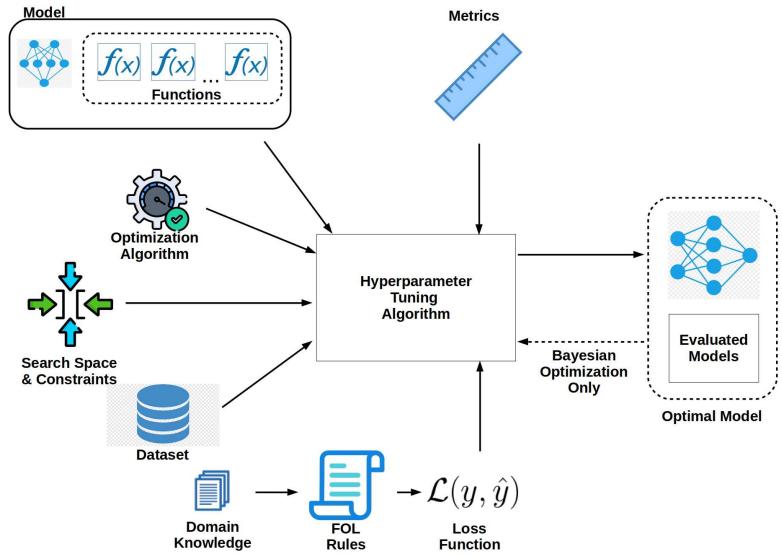


Fig. 1. Framework system overview.

III. USER WORKFLOW

We will touch upon the specifics of the framework as we go through an overview of the user workflow. There are four stages to the workflow: creating rules, modifying the original model and loss function, model tuning, and training and evaluation.

A. Creating Rules

First the user must decide what rules they want to incorporate into the model. Each rule must be converted into a FOL statement. And each statement into an FOL object. In FOL the simplest element of a statement is an atom. Atoms are composed of predicates or functions and their terms. Terms are the arguments of the function or predicate. Predicates return only true or false while functions can return any value in a specified range. We will build our atoms from

functions. The functions contained in the statements are networks from the deep learning model or user defined. Probabilistic FOL operators such as AND, OR, NOT, and IMPLIES connect functions to make a meaningful rule.

Probabilistic FOL operators are continuous and have a range of $[0, 1]$ unlike normal binary operators which are discrete [48]. Because the operators are continuous and differentiable, they don't break the backpropagation algorithm used to train the neural networks [37]. See [48] for an introduction to probabilistic FOL.

Starting from natural language, most concepts can be converted into FOL. The simplest and most common statements are of the form

$$\forall x P(x) \Rightarrow Q(x) \quad (1)$$

which reads for all x , P of x implies Q of x where P and Q are functions of x . For example, suppose the user is training a species classification model on RGB and LiDAR data from a mature forest containing two species of trees, red maples and white pines. The user knows that in this forest, all red maples have a crown height of 40 meters or less, so any trees taller than 40 m have to be white pines. The user can write an FOL statement for this rule as follows,

$$heightGreater Than 40(x) \Rightarrow whitePine(x) \quad (2)$$

where x is an RGB instance from the dataset with its corresponding CHM. $heightGreater Than 40$ and $whitePine$ are functions, in this case neural networks that given an instance from the dataset predict a real number between 0 and 1. A tree much taller than 40 m would produce a prediction near 1 and a tree less than 40 meters would produce a prediction closer to 0. If the instance is a white pine, the function $whitePine$ should predict a value close to 1 otherwise it should predict a value close to 0.

Atoms can be linked by operators. In probabilistic FOL there are several implementations of operators, but in all cases the operators are implemented as functions. In this work we use the Lukaseiwicz t-norm version of the operators given in Table I [49].

The truthfulness of an FOL expression is then calculated by substituting the output from each function into the FOL operators and evaluating the expression. The process of assigning values to the terms of a FOL expression is called grounding. A true expression will evaluate to 1 and a

TABLE I
LUKASEIWICZ OPERATORS AND THEIR IMPLEMENTATION.

Operation	Symbol	Implementation (Lukaseiwicz t-norm)
AND	$x \wedge y$	$\max(0, x + y - 1)$
OR	$x \vee y$	$\min(1, x + y)$
NOT	$\neg x$	$1 - x$
IMPLIES	$x \Rightarrow y$	$\min(1, 1 - x + y)$

false expression to 0. An expression can have any value between 0 and 1. Using the Lukaseiwicz t-norm, the example rule would be evaluated as

$$\min(1, 1 - \text{heightGreater Than}40(x) + \text{whitePine}(x)). \quad (3)$$

Each rule is assigned a real number $\lambda \in [0, \infty)$ that represents the importance of the rule.

Our framework provides templates for basic FOL expressions. These templates are classes that the user can modify to fit their needs using inheritance. There are templates for rules of the following forms: $\forall x P(x) \Rightarrow Q(x)$, $\forall x P(x) \Rightarrow \neg Q(x)$, $\forall x P(x) \Rightarrow Q_1(x) \vee Q_2(x) \vee \dots \vee Q_N(x)$, $\forall x P_1(x) \vee P_2(x) \vee \dots \vee P_N(x)$, $\forall x \neg P(x) \Rightarrow Q(x)$, and $\forall x P(x) \iff Q(x)$. The rule classes have methods for each FOL operation and an eval method to evaluate the rule. The atoms of the expression are the arguments of the eval method. The rule objects have a generic_interface method to simplify linking the model and the rule object. In the generic interface the user should create variables that map the predictions from the model to named variables in the rule object. Table II lists the name of each class and its expression. Classes for new expressions can be added using inheritance.

TABLE II
PREDEFINED EXPRESSION CLASSES AND THEIR CORRESPONDING FOL.

Class Name	FOL
Rule_p_imp_q	$\forall x P(x) \Rightarrow Q(x)$
Rule_p_imp_not_q	$\forall x P(x) \Rightarrow \neg Q(x)$
Rule_p_imp_disj_q	$\forall x P(x) \Rightarrow Q_1(x) \vee Q_2(x) \vee \dots \vee Q_N(x)$
Rule_disj_p	$\forall x P_1(x) \vee P_2(x) \vee \dots \vee P_N(x)$
Rule_not_p_imp_q	$\forall x \neg P(x) \Rightarrow Q(x)$
Rule_p_iff_q	$\forall x P(x) \iff Q(x)$

B. Modifying the Original Model and Loss Function

Assuming the user has an existing deep learning model that they want to make neuro-symbolic, the next step is model modification. Continuing with the example from step 1, let's assume that we have a classifier that can classify tree species from RGB images as our original model. We'll call the RGB image x_i and its corresponding label y_i .

There are many ways to create neuro-symbolic models from vanilla deep neural networks; one of the most straightforward is semantic based regularization. That is the method we use in this paper. As the name implies, SBR adds rules written as statements of FOL to the model's loss function. See [36], [50], [51] for a more comprehensive discussion of SBR.

Functions that make up the atoms of a FOL expression must be added to the model. These functions can be learned or user defined [36]. In this work we use user defined functions, but learned functions can be readily used with our framework. Each new network acts to accomplish a task such as identifying trees that have a crown height of greater than 40 m. The addition of the learned functions effectively turns the model into a multi-task learning deep neural network. Figure 2 shows a diagram of a modified model and the training process. In the diagram, functions $f(x)$ are added to the original network and their predictions are used in the FOL expressions. The user can arbitrarily add many functions. In the case of our toy example, the dataset would need to be augmented with the CHM model aligned with the RGB images. This can be done by appending a CHM raster of the same width and height to the RGB image as an additional layer.

How new networks are added to the model depends on model implementation. In the code provided, we use PyTorch Lightning's `LightningModule` class as the backbone of each model. Then new networks are added using PyTorch's `nn.Module` class. A description of its use can be found in [52]. There are no restrictions on the architecture of the additional networks. However, the output of the final layer must be a real number between 0 and 1 or must be able to be decomposed into a real number between 0 and 1 if the output has more than one dimension.

New loss terms must be added to the original model's loss function. Loss terms must be added for the new networks. A loss term must also be added to complete the SBR implementation. The additional term penalizes the model if it fails to make predictions that make all rules true.

We'll use the notation from [51]. The SBR loss term takes the form

$$\sum_{h=1}^H \lambda_h \cdot (1 - \Phi_h(f(\chi))) \quad (4)$$

where H is the number of rules, λ_h is the user assigned importance of rule Φ_h , Φ_h is the average evaluation of rule h when grounded using the data in the training batch, f is the set of functions comprising the rules,

$$f = \{f_1, f_2, \dots, f_T\}. \quad (5)$$

For our example rule Φ becomes

$$\Phi(f(\chi)) = \frac{1}{|\chi|} \sum_{x \in \chi} \min(1, 1 - \text{heightGreater Than40}(x) + \text{whitePine}(x)) \quad (6)$$

and x is the grounding of the argument for each function for all values in a training batch. The final loss function should have the following form

$$L_s(y, \hat{y}) + k_1 \sum_{h=1}^H \lambda_h \cdot (1 - \Phi_h(f(\chi))) \quad (7)$$

where L_s is the typical loss that includes terms for loss from networks and regularization and the summation term is the loss incurred from the SBR rules.

C. Model Tuning

Hyperparameter tuning is the selection of values for the model's non-learned parameters. The values selected can have a heavy impact on the model's predictive performance, training time, and other performance measurements. Once the model and loss function have been updated, the hyperparameters can be tuned. The process is usually done on a validation set, a set of hold out data separate from the test and training set. To begin the tuning process the user must decide what variables to tune. Neural networks typically have many hyperparameters such as the learning rate, an L2 regularization constant, or the number of layers used in the network. Because the focus of this work is the creation of neuro-symbolic models, we focus on the hyperparameters that relate to the rules, in our case λ_h . After a variable is selected, the user must define the upper and lower bounds of the search space. Selecting bounds is usually done through a combination of intuition and trial and error. Once the user decides on bounds for the search space, one or more metrics must be chosen to gauge the model's performance as the search parameters are varied. Common metrics include F1, accuracy, and the total loss. The user also selects an optimization algorithm

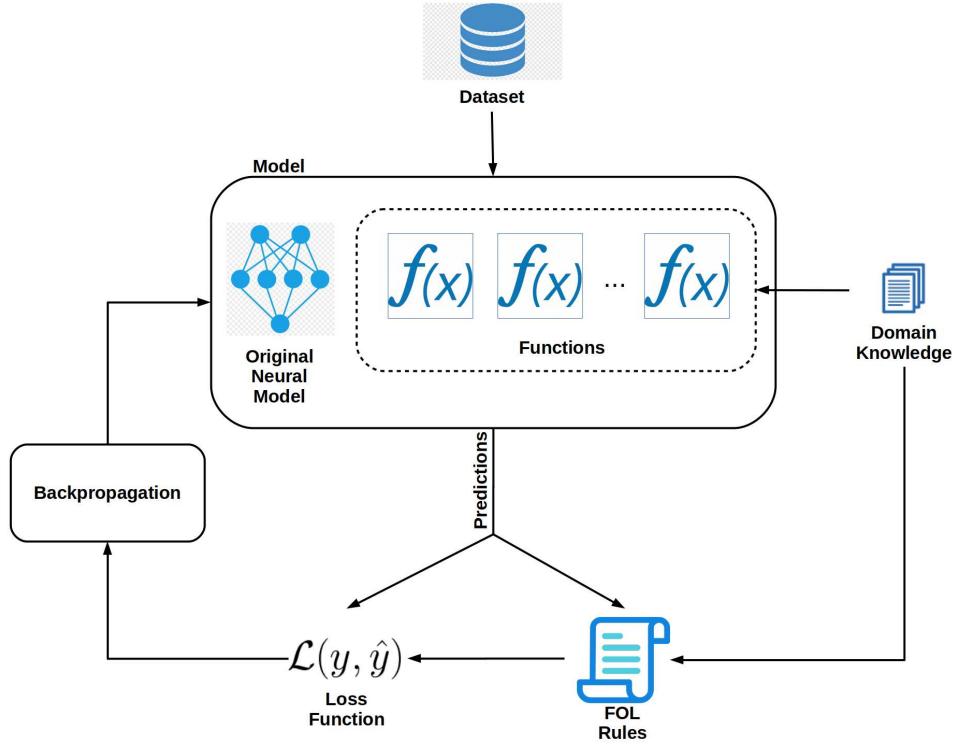


Fig. 2. Model modifications needed to create a neuro-symbolic model.

such as the Adam optimizer or plain stochastic gradient descent. The optimizer selected will also be used during the training process. Searching through the parameter space is an iterative process. Each iteration is known as a trial. The user decides the number of trials to run prior to starting the hyperparameter tuning algorithm. Our framework includes 3 algorithms for hyperparameter tuning, random search, grid search and Bayesian optimization. Random search randomly selects points within the search space with a uniform distribution. Random search is shown to give good results especially for high dimensional search spaces, that require a great deal of time to train. Random search gives near optimum results with as few as 60 trials [53].

Rather than picking random points within the search space, grid search searches the space systematically using an even spacing between search points. The granularity of the search is determined by the number of trials.

Unlike random search and grid search, Bayesian optimization is a closed loop search algorithm; the results of previous trials are used to choose the next points for evaluation. The algorithm attempts to learn which areas of the search space are likely to give optimum results. Bayesian

optimization requires at least 50 trials or more to work well. As the volume of the search space grows, more trials are needed to produce good results. See [54] for a detailed explanation of Bayesian optimization.

After selecting the desired tuning parameters, the model is evaluated for the indicated number of trials. After the last trial, a sorted list of the chosen model parameters and their associated performance is printed in descending order as well as saved to a CSV file.

D. Training and Evaluation

The user can manually enter the best hyperparameters found in the previous step prior to training and testing the model on the test set.

IV. USE CASES

A. Individual Tree Crown Delineation

We demonstrate the use of our framework for crown delineation using DeepForest, a popular open source crown delineation model based on RetinaNet, a CNN [40], [55]. Using our framework we will

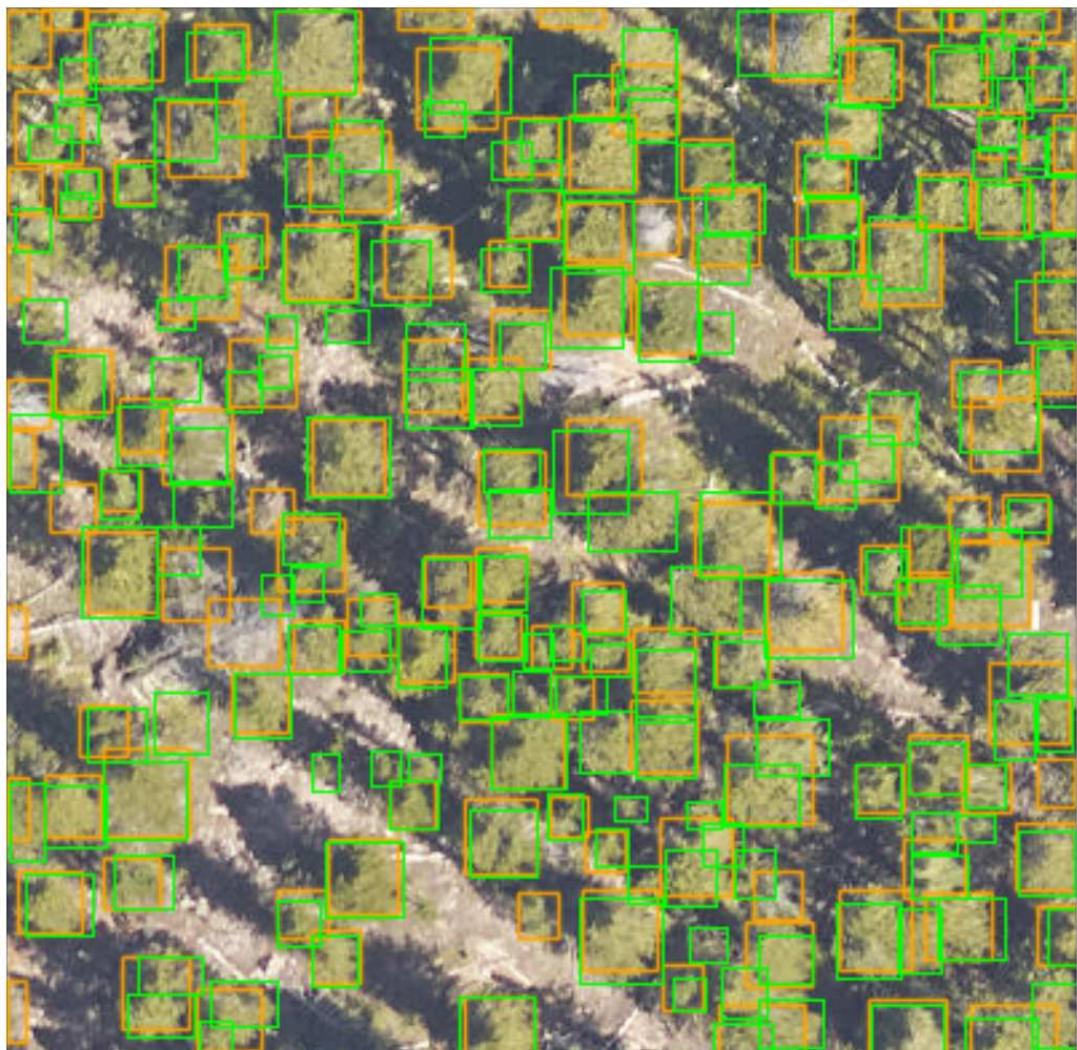
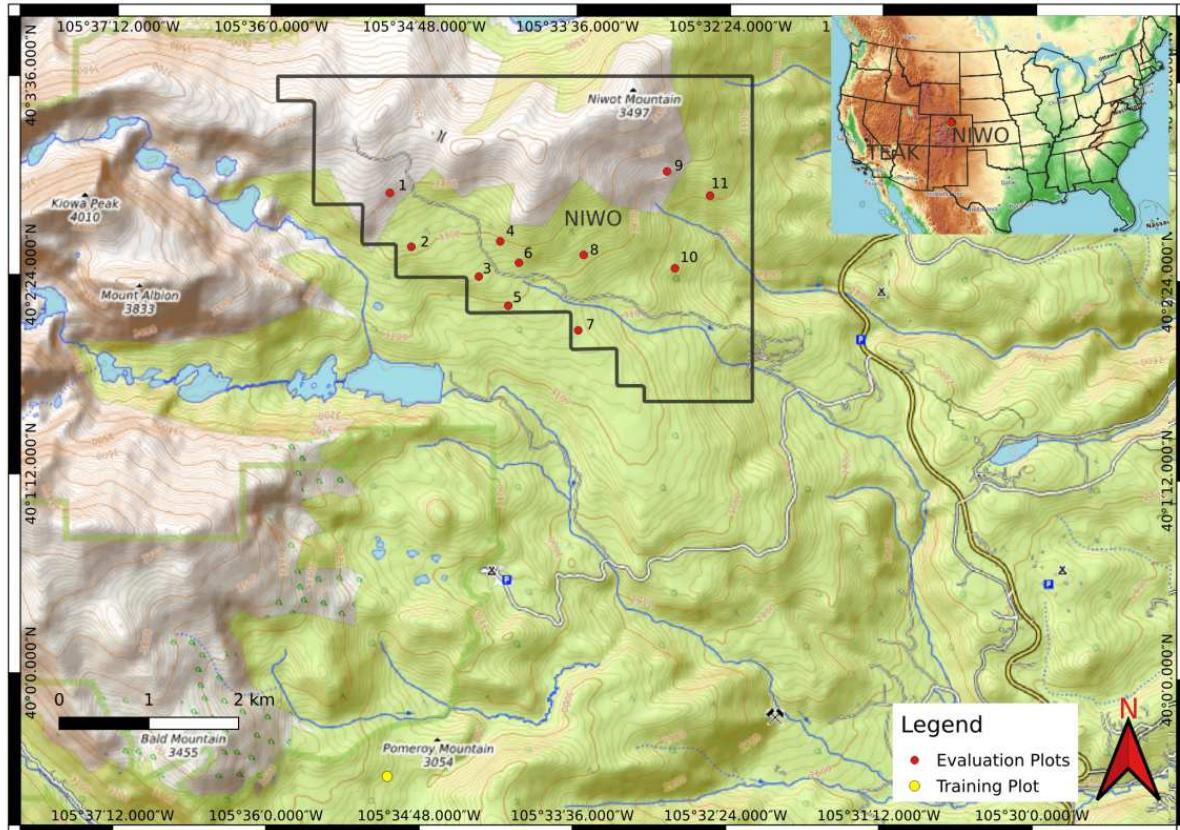
- 1) Create two rules
- 2) Modify the model and loss function to use SBR
- 3) Find optimum values for the rule lambdas
- 4) Evaluate the effectiveness of each rule.

We train the model on data from Niwot Ridge (NIWO), an alpine forest in Colorado USA. Niwot is located at 40.05425° latitude and -105.58237° longitude. Its mean annual precipitation is 1,005 mm and the mean annual temperature is 0.3°C. The site elevation ranges from 2,975 m to 2,583 m. Tree canopy cover varies from continuous at lower elevations and south- and east-facing slopes to open forest with isolated trees at higher elevation. Understory coverage is limited, especially at higher elevations. The mean canopy height is 0.2 m (all canopy height values listed here include bare ground and the prevalence of bare ground at Niwot Ridge is why this value is so small). The dominant tree species are conifers, primarily lodgepole pine, subalpine fir, and Engelmann spruce [43], [56].

1) *Data*: The data was collected in 2018 by the National Ecological Observatory Network (NEON). NEON is a program funded by the US federal government tasked with monitoring environmental health at over 80 sites across the continental United States, Alaska, and Puerto Rico [42]. Part of their mission is to monitor forest health, carbon fluxes, and biodiversity changes through both field surveys and remote sensing. NEON annually overflies forests located at their sites during periods of peak greenness using their airborne observation platform (AOP) [42]. The NEON AOP is an aircraft outfitted with RGB, hyperspectral (HSI), LiDAR, and GPS sensors. The HSI data has a resolution of 1 m and RGB imagery data has a resolution of 0.1 m. The canopy height model (CHM) data produced by the LiDAR sensor has a spatial resolution of $1 m^2$ per pixel. NEON's hyperspectral data is atmospherically corrected and all data is orthorectified and aligned to a uniform spatial grid.

We train our crown delineation model on a combination of RGB and CHM data. The original data comes from NEON's L3 data products which are $1 km^2$ mosaicked tiles. We decomposed rasters larger than 500 pixels in any dimension down to 400 x 400 pixels or less rasters with 5% overlap. A map showing the geolocation of the training and evaluation plots used is shown in Figure 3.

An example evaluation plot is shown in Figure 3 (bottom). Each evaluation plot is 400 x 400 pixels or an area of $1600 m^2$. The number of trees per evaluation plot ranges from 5 to 292. The training plot is 2511 x 4132 pixels with an area of $103,975 m^2$. The plot has 12,412 trees. See [57] and [58] for a detailed description of the data collection and preparation methodology.



The dataset contains a single class - tree. We split the data into training, validation, and test sets. The size of each set is 10,757, 1,655, and 1,624 annotations respectively. The distribution of the crown area (derived from the crown delineations) and the height-crown area allometry (the relationship between each tree’s crown area and its height derived from the canopy height model) are shown in Figure 4. We will use the information contained in these graphs to craft two

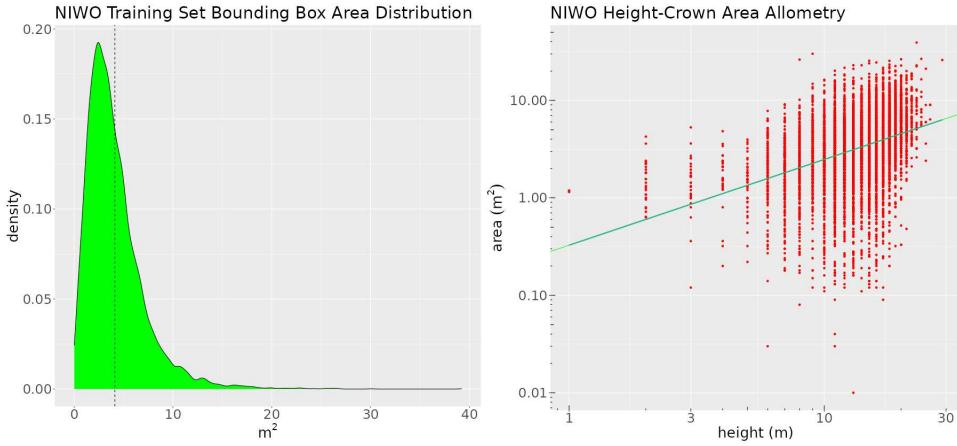


Fig. 4. The left image shows the crown area distribution of the training set at Niwot Ridge. The dashed black line is the mean of the distribution. The right image is a plot of the height-crown area allometry for the Niwot training set. The green line is the fitted log-linear height-crown function.

rules for the site. We append the CHM rasters as a 4th layer to the RGB images. The training data is augmented using geometric transformations.

2) *Creating Rules*: DeepForest delineates crowns using rectangular bounding boxes. These bounding boxes are sometimes larger than the crowns of the trees they detect, which may reduce the performance of the model. We will create rules intended to reduce DeepForest’s tendency to use oversized bounding boxes. We will use two of the rules developed in [58]. As shown in Figure 4, the mean ITC area at Niwot is 400 pixels which corresponds to approximately $4 m^2$. The distribution is right skewed with 60% of the crowns less than the mean. We use this fact to create the following rule: a detected object is a tree if and only if the area of its bounding box is less than $4 m^2$. Despite the use of if-and-only-if, the rule is not as strict as it sounds in natural language. We can control how strictly the rule is enforced by varying the value of the rule’s lambda. This allows us to use a simple rule like this even though we know that there are trees with crowns $\gtrsim 4 m^2$. We also create a second rule, similar to the first, more tightly related to how we think about limits on crown area biologically, by using the height-crown area

allometry model shown in Figure 4. As described in [59] and [60] we fit a power function to map the height to the crown area for the trees at Niwot and use the CHM to predict crown area given the crown height. We use the predicted crown area to create the following rule: a detected object is a tree if and only if its bounding box area is less than or equal to the area predicted by the crown height allometry model. We'll apply both rules at the same time and use the results of our evaluation to choose the better rule. In order to implement each rule, we create two user defined functions. The first function quantifies how much the bounding box area for a detected object deviates from the site mean. If a bounding box is less than 400 pixels the function's output tends towards 1. If the bounding box of a detected object is greater than 400 pixels, its output tends towards 0. We use the sigmoid function to ensure differentiability and write the function as

$$f_1(f_0(x)) = \frac{1}{1 + \exp(-0.5 \cdot (400 - A_{bbox}(f_0(x)))}) \quad (8)$$

where $f_0(x)$ is DeepForest's predictions for x and A_{bbox} is a function that calculates the area of a prediction's bounding box. For Niwot, our fitted power function for the height-crown area allometry is

$$A_{itc}(h) = 0.32658 \cdot h^{0.87992}. \quad (9)$$

To create our second function, we substitute the site mean area with A_{itc} to give

$$f_2(f_0(x)) = \frac{1}{1 + \exp(-0.5 \cdot (A_{itc}(\max_{CHM}(x)) - A_{bbox}(f_0(x)))))}. \quad (10)$$

We use the function \max_{CHM} to extract the maximum value of the CHM for prediction $f_0(x)$. We rename our functions to make the FOL more intuitive. Let f_0 be $isTree(\cdot)$. Let f_1 be $bboxAreaLTMean(\cdot)$. Let f_2 be $bboxAreaLTEHCA(\cdot)$. Using the functions we can write our rules in FOL as follows:

$$\forall x bboxAreaLTMean(isTree(x)) \iff isTree(x) \text{ (rule 1)} \quad (11)$$

and

$$\forall x bboxAreaLTEHCA(isTree(x)) \iff isTree(x) \text{ (rule 2).} \quad (12)$$

3) *Model and Loss Function:* DeepForest is built from RetinaNet [40]. RetinaNet is a deep CNN with a feature pyramid network for scale invariance, a specialized loss function, and residual connections between layers. The network has 32.1 million learned parameters [55]. RetinaNet is well suited for object detection. The network predictions include bounding box coordinates for detected objects as well as the object’s class. To complete the model, we add the additional terms to the loss function as described in section 3.2. As in [58] we delay the application of the rules to the latter epochs of training to give the model time to learn purely from the dataset. We accomplish this by multiplying the rule loss by a function

$$\pi(t) = \begin{cases} 1.0 - \max\{\pi_0, \alpha^{0.029 \cdot t}\}, & 0.029 \cdot t > \pi_s \\ 0, & \text{otherwise} \end{cases} \quad (13)$$

where π_0 and α are constants $\downarrow 1$, t is the training step number, and π_s is a constant that controls at what step the rules are applied. The loss function then becomes

$$L_{tot} = L_s + \pi(t) \cdot L_{rules} \quad (14)$$

where L_{tot} is the summation of all the loss terms, L_s is the total standard loss and L_{rule} is the sum of the loss from all rules. This method was chosen empirically to improve model performance [35].

4) *Hyperparameter Tuning and Training:* Let λ_1 and λ_2 be the hyperparameters associated with rules 1 and 2 respectively. We optimize the values of λ_1 and λ_2 to give an optimum validation set F1 score. We set the bounds of each λ to be between 0.01 and 9.0. Empirically, large values of λ tend to degrade model performance. All model hyperparameters are shown in Table III.

TABLE III

CROWN DELINEATION MODEL HYPERPARAMETERS AND SPECIES CLASSIFICATION MODEL HYPERPARAMETERS FOR TUNING.

Parameter	DeepForest (crown delineation)		Fricker (species classification)
	Value	Value	Value
Variables	$\{\lambda_1, \lambda_2\}$		$\{\lambda_1, \lambda_2\}$
Bounds	$\lambda_1 \in [0.01, 9.0], \lambda_2 \in [0.01, 9.0]$		$\lambda_1 \in [0.01, 20.0], \lambda_2 \in [0.01, 20.0]$
α	0.995		N/A
Batch size	1		32
Epochs	7		5
IoU Threshold	0.4		N/A
Search Algorithms	Bayesian Optimization, Grid Search, Random Search		Bayesian Optimization, Grid Search, Random Search
L2 constant	N/A		1×10^{-3}
Learning Rate	1×10^{-3}		1×10^{-4}
Number of Trials	64		64
Optimization Algorithm	Stochastic Gradient Descent		Adam Optimizer
π_0	0.7		N/A
π_s	9.88		N/A
Search Algo.			
Evaluation Metric	Validation F1		Validation F1
k1	1.0		1.0

To define F1 for an object detection model, first we define the intersection over union score or IoU. IoU is used as a measure of overlap between two bounding boxes. Given bounding box A and bounding box B an IoU of 0 represents no overlap and an IoU of 1 represents full overlap.

IoU is calculated as

$$IoU = \frac{A_{area} \cap B_{area}}{A_{area} \cup B_{area}}. \quad (15)$$

For a DeepForest prediction to be counted as a true positive (TP), the prediction must overlap a ground truth bounding box with an $IoU \geq 0.4$. Prediction bounding boxes that fail to meet this requirement are counted as false positives (FP). If a ground truth bounding does not have a matching prediction, it is counted as false negative (FN). Using these definitions we define precision as

$$Prec = \frac{TP}{TP + FP} \quad (16)$$

and recall as

$$Rec = \frac{TP}{TP + FN}. \quad (17)$$

Then, F1 is

$$F1 = \frac{2 \cdot Prec \cdot Rec}{Prec + Rec}. \quad (18)$$

We tune the model using 3 search algorithms: Bayesian optimization, grid search and random search. We use the default Bayesian optimization settings for the Ax package. We allow 64 trials of each algorithm. Model initialization produces variation in the results so we seeded the random number generator with 15 different seeds and report the average of the results. The model was tuned and trained on one node of a high performance computing cluster using 16 GB of RAM, 1 Nvidia A100 GPU, and 1 CPU.

5) Results: All 3 search algorithms improved the model's performance over the non-neuro-symbolic model. The average change in F1 score for each model in comparison to the non-neuro-symbolic version of DeepForest is shown in Table IV. Bayesian optimization and grid search gave nearly identical results, improving F1 by approximately 2 F1-points. Random search gave the best result, improving model performance by 2.14 F1-points.

The bottom image in Figure 3 shows an example output of the model after training using optimal rule hyperparameters found using the random search algorithm. Ground truth bounding boxes are in green and the model's predicted bounding boxes are in orange.

Figure 5 shows a density map of the points in the search space selected by each search algorithm. The figure highlights the differences in the search algorithms' strategies. Bayesian optimization tries to focus its search to areas of the plane that are likely to give optimum results

TABLE IV
AVERAGE CHANGE IN TEST F1 COMPARED TO THE NON-NEURO-SYMBOLIC MODEL.

Model	Bayesian	Grid	Random
	$\Delta F1$	$\Delta F1$	$\Delta F1$
DeepForest (crown delineation)	+2.03	+2.04	+2.14
Fricker (species classification)	+1.11	+0.8	+3.02

as can be seen by the high point density in the region near the λ_1 axis of its graph. Grid search evenly distributes its points across the plane and random search chooses points at random as reflected by its graph's lack of structure. Despite the differences in the search strategies, the

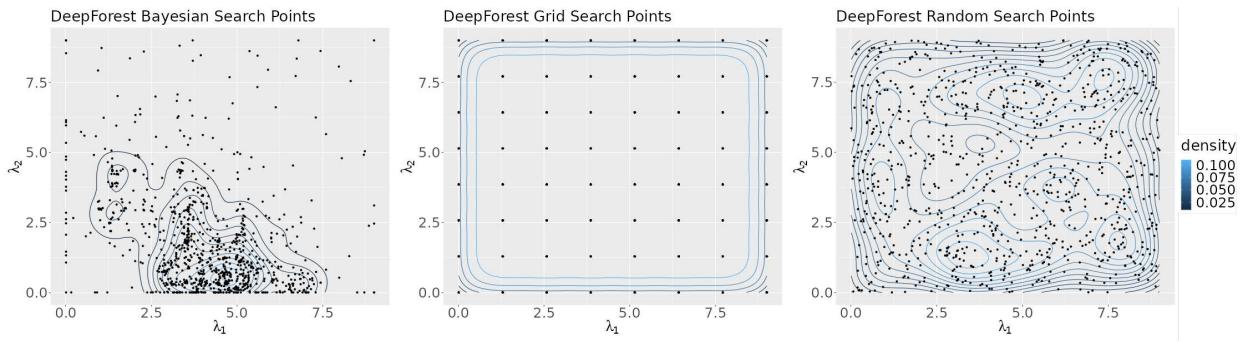


Fig. 5. Contour plots of the points selected by each algorithm for all seeds.

rate of improvement as a function of trial number did not significantly vary between methods. As shown in Figure 6, after 29 trials each method was within 0.5 points of the validation F1's maximum range suggesting that all the search algorithms were able to find optimum parameters with as few as 29 trials.

Figure 7 shows a plot of the validation F1 scores as a function of the rule lambdas. The plot shows optimum results can be obtained with when λ_2 is near zero. This suggests that λ_1 is the more effective of the two tested rules.

V. TREE SPECIES CLASSIFICATION

To demonstrate the use of our framework for a tree species classification use case, we will

- 1) Create two rules
- 2) Modify a tree species classification model and loss function for SBR

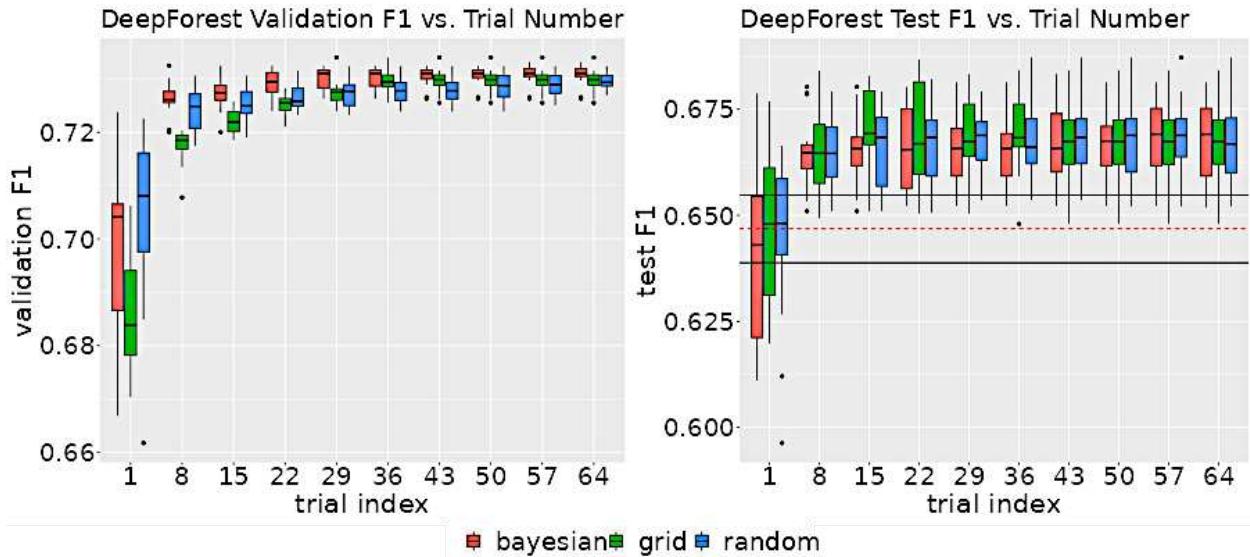


Fig. 6. The left graph shows the distribution of validation scores as a function of the number of trials for each search method. The right graph shows the distribution of test F1 scores as a function of the number of trials for each method. The test F1 score was determined by evaluating the model on the test set using the hyperparameters associated with the highest validation score at the listed trial index. The black horizontal lines on the test F1 graph show the upper and lower 95% confidence intervals for non-neuro-symbolic DeepForest. The dashed red line is the mean test set F1 score of the non-neuro-symbolic model.

- 3) Find optimum values for the rule lambdas
- 4) Evaluate the effectiveness of each rule.

We use the model and data from [41] and two rules developed in [61]. The data is from TEAK, a mixed conifer forest in California USA close to the Nevada border. TEAK is located at 37.000583° of latitude and -119.00602° longitude. It has a mean annual precipitation of 1,223 mm and a mean annual temperature of 8° C. The average canopy height is 35 m. The elevation ranges from 2,086 m - 2,734 m. The dominant tree species are red fir, white fir, Jeffrey pine and lodgepole pine. Sixty-five percent of the study area is mixed conifer forest. The rest of the area is covered by areas dominated by a single species, namely red fir or lodgepole pine. Forest structure ranges from closed canopy forest, to open stands with isolated trees. Understory density ranges from none at higher elevations, to dense understory in some areas at lower elevations. Fricker et al. (2019) gives more information on the site.

1) *Data:* The remote sensing data was collected in 2017 by NEON. NEON surveys its sites annually during periods of peak greenness. Like the data used for crown delineation, the HSI data has a resolution of 1 m and RGB imagery data has a resolution of 0.1 m. The canopy

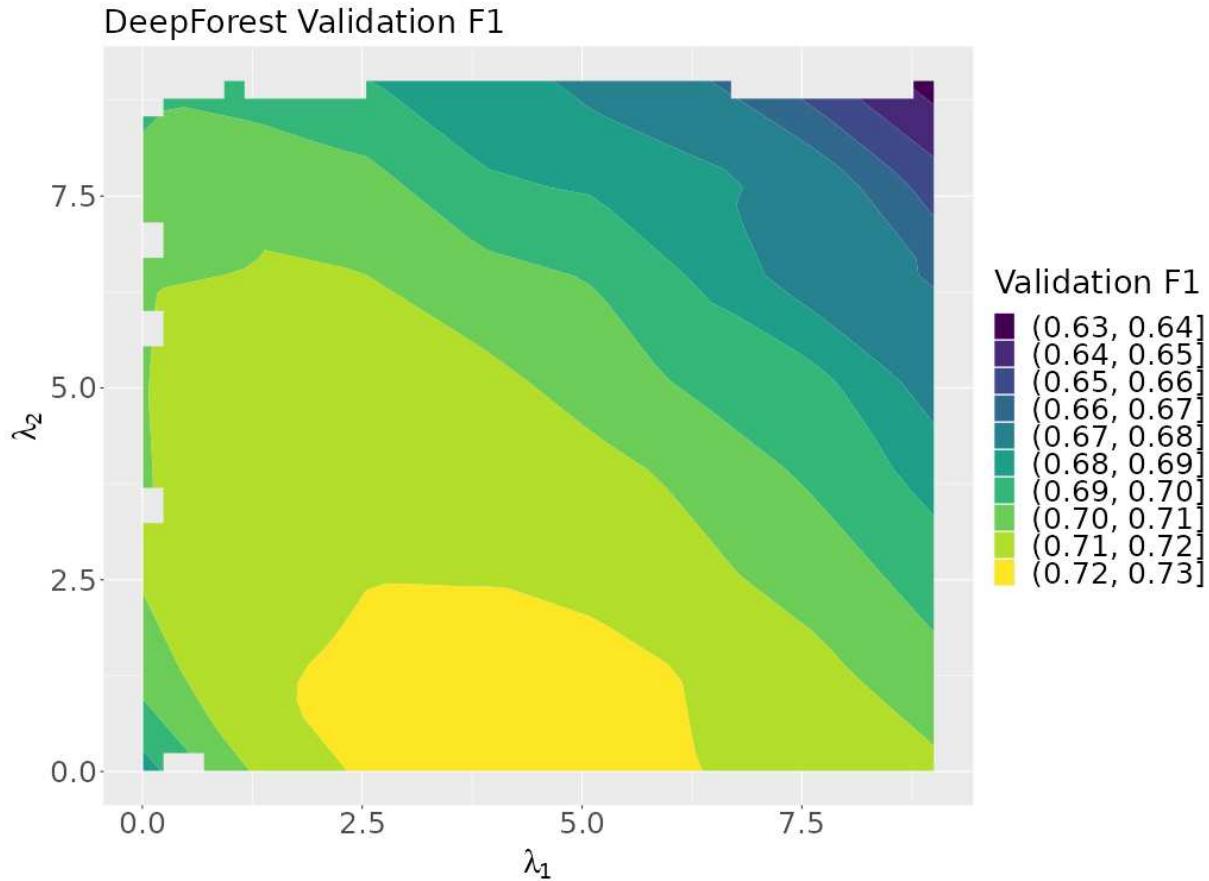


Fig. 7. A contour plot of the validation F1 scores over the search space.

height model produced from the LiDAR data has a spatial resolution of 1 m^2 per pixel. The NEON remote sensing data was augmented with field survey data collected by the author of [41] using sites established by [62]. NEON's HSI data is atmospherically corrected and all data is orthorectified and aligned to a uniform spatial grid. See [41] for a more detailed summary of the data collection and preparation methodology.

In this work, we use RGB and CHM data for species classification, although HSI data is more often used for this task. RGB data is more readily available outside of NEON sites and the RGB based classification is more prone to errors allowing for better demonstration of our methods. Nevertheless, our methods are applicable to HSI data as well. The dataset contains 8 classes: white fir, red fir, incense cedar, Jeffrey pine, sugar pine, black oak, lodgepole pine, and dead. The dead category is composed of standing dead trees of any species. The top image in Figure 8 shows a map of TEAK and the bottom image shows a plot where tree species are identified.

Crowns from each tree were broken into patches of 15x15 pixels. Geometric transformation data augmentation was used to increase the size of the training set using the following transformations: horizontal flips, vertical flips, and 90 degree rotations. Table V shows the number of trees and pixel patches of each class.

TABLE V
DATASET TREE CROWN AND PATCH COUNT BY TREE SPECIES.

Species	Tree Count	Patch Count
white fir	119	2,908
red fir	47	851
incense cedar	66	1,853
Jeffrey pine	164	4,384
sugar pine	68	2,740
black oak	18	111
lodgepole pine	62	895
dead (any species)	169	3,520
Total	713	17,262

We append the CHM raster to RGB images as a 4th layer. The CHM is removed prior to passing the data through the main network. Post data augmentation, we create a training set composed of 111,355 patches, a validation set of 12,373 patches, and a test set of 1,796 patches. The test set patches were not augmented. Figure 9 shows the distribution of tree crown heights in the training set for each species. We will use this information to craft two rules.

2) *Creating Rules*: We will use the observed crown heights to help improve the model’s overall performance. We extend rules 1 and 2 from [61] which are based on the differences in crown height distributions. Rule 1 states that trees taller than 46.0 m are unlikely to be black oak. Rule 2 states that trees taller than 53.2 meters are unlikely to be lodgepole pine. Note from the distributions that the majority of black oaks are less than 46 m height and the majority of lodgepole pines are less than 53.2 m. We implement these rules by first creating our functions

$$f_1(\mathbf{x}) = \frac{1}{1 + \exp(-(1 \times 10^3(-\max_{CHM}(\mathbf{x}) + 46.0)))} \quad (19)$$

and

$$f_2(\mathbf{x}) = \frac{1}{1 + \exp(-(1 \times 10^3(-\max_{CHM}(\mathbf{x}) + 53.2)))}. \quad (20)$$

Both f_1 and f_2 are built on the sigmoid function so both are differentiable. For f_1 when the CHM values of x are less than 46.0 m the output of the function tends toward 1 and when the

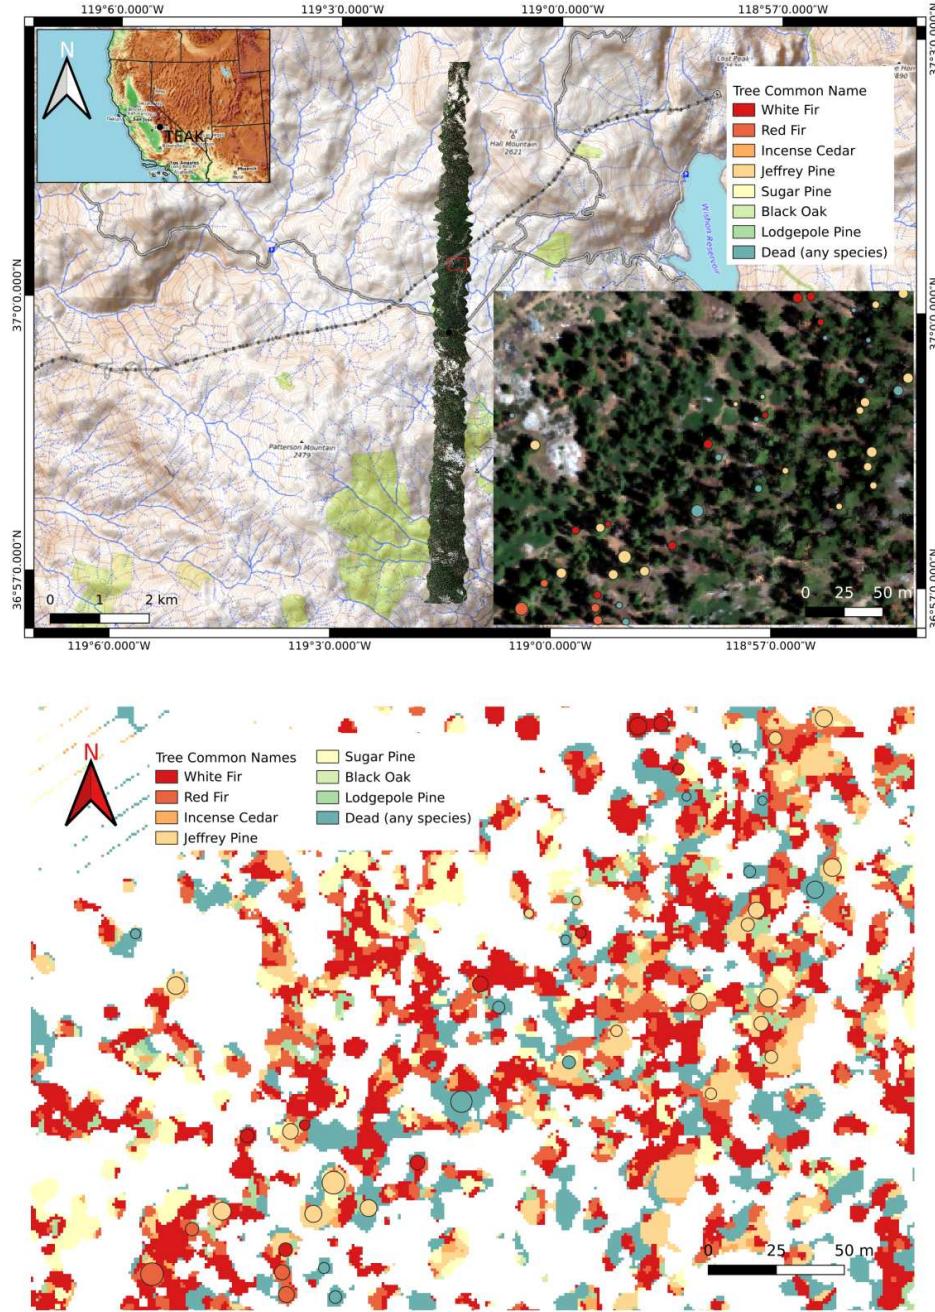


Fig. 8. **(top)** A map of TEAK. The zoomed in region in the lower right corner shows ground truth tree species labels. **(bottom)** Model predictions for the zoomed in region above generated using optimum hyperparameters found using the random search algorithm. Ground truth crown locations (from [41]) are indicated by circles. Colors within the circles are ground truth species identities.

values of x for the CHM are greater than 46 m the output tends toward 0. We rename f_1 and f_2 to make our FOL more readable. Let f_1 be $chmGT46$ and let f_2 be $chmGT53$. We then

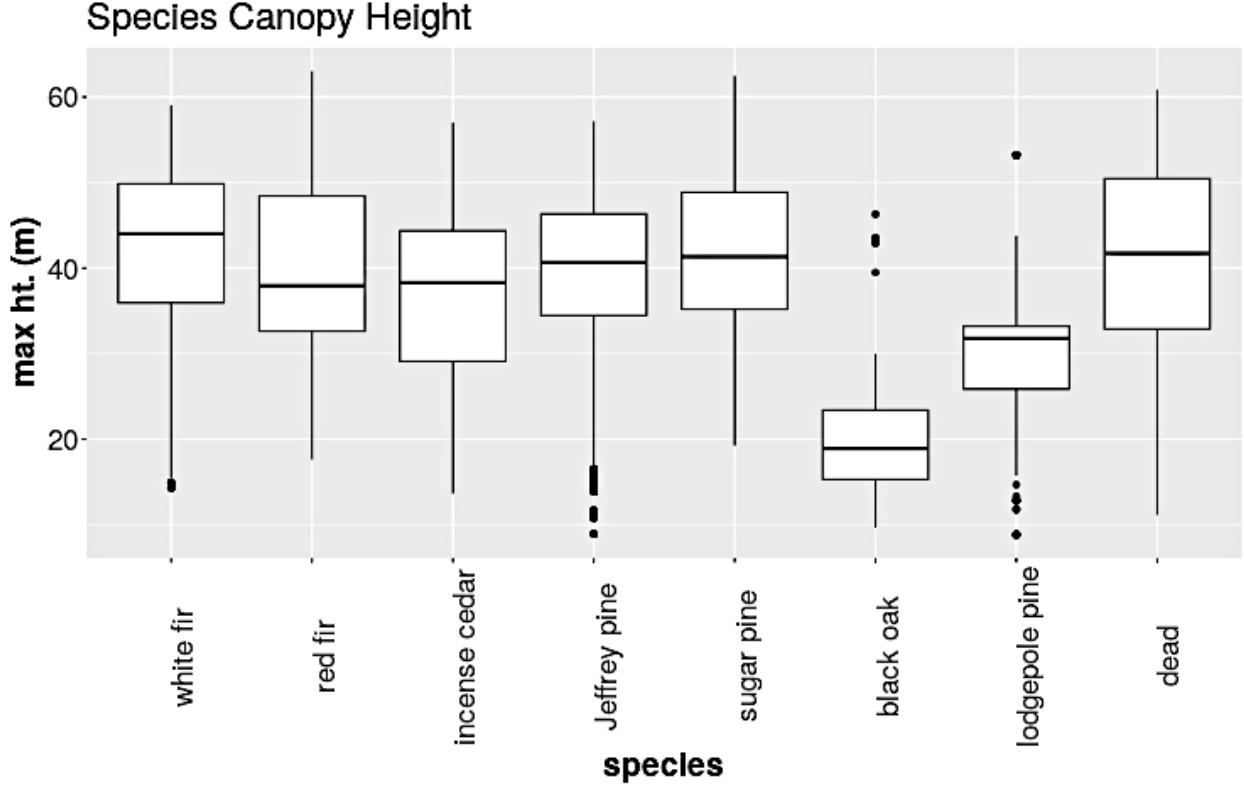


Fig. 9. The box and whisker plot of the crown height distribution for each species. Note that black oak and lodgepole pine are the shortest species.

implement our rules in FOL as

$$\forall x \text{ chmGT46}(x) \Rightarrow \neg \text{isBlackOak}(x) \text{ (rule 1)} \quad (21)$$

and

$$\forall x \text{ chmGT53}(x) \Rightarrow \neg \text{isLodgepolePine}(x) \text{ (rule 2).} \quad (22)$$

The functions *isBlackOak* and *isLodgepolePine* come from the original classifier whose function we'll call f_0 . We use the components of the original classifier's final prediction that correspond to the indices of the black oak and lodgepole pine classes respectively.

3) *Model and Loss Function:* The model comes from [41]. It is an 8 layer fully convolutional neural network with a softmax output layer. It has 684,000 learned parameters. We modify the loss function by adding a term for the rule loss as described in section 2.

4) *Hyperparameter Tuning and Training:* The hyperparameters associated with rule 1 and rule 2, λ_1 and λ_2 respectively, will be the variables we tune the model on. Table III shows all

hyperparameter values used to tune the model. The choice of search space bounds was done through intuition. Large values for rule lambdas tend to worsen model performance. Model performance was found to vary with model initialization, therefore we ran the experiment 10 times, each time initializing the random number generator with a different value and averaged the results. The scale variable was set arbitrarily to 1.0. We ran separate sets of 10 trials of 64 using all 3 search algorithms for comparison. When tuning for a production model, this need not be the case. We chose the validation macro F1 score to compare results across trials. We define F1, precision and recall as in section IV-A4 where TP is the number of true positives, FP is the number of false positives, and FN is the number of false negatives. We report the performance of the tuned model using the test set macro F1 score. The model was tuned and trained on one node of a high performance computing cluster using 16 GB of RAM, 1 Nvidia A100 GPU, and 1 CPU.

5) *Results:* As shown in Table IV all 3 search methods found parameters that improved the model’s performance compared to the non-neuro-symbolic model. Again, the random search algorithm gave the largest improvement with 3.02 F1-points. Bayesian optimization fared second best with 1.11 F1-points, and the grid search algorithm only found parameters that improved the model by 0.8 F1-points. The bottom image in Figure 8 shows the predicted species for the zoomed region at the top of the figure. The circles indicate ground truth labels.

Figure 10 shows density plots of the points selected in the search space by each algorithm. The Bayesian optimization algorithm focused its search heavily on the lower right corner of the search space. Grid search points were evenly distributed throughout the space. The random search algorithms points show no recognizable pattern. Figure 12 shows the distribution of the

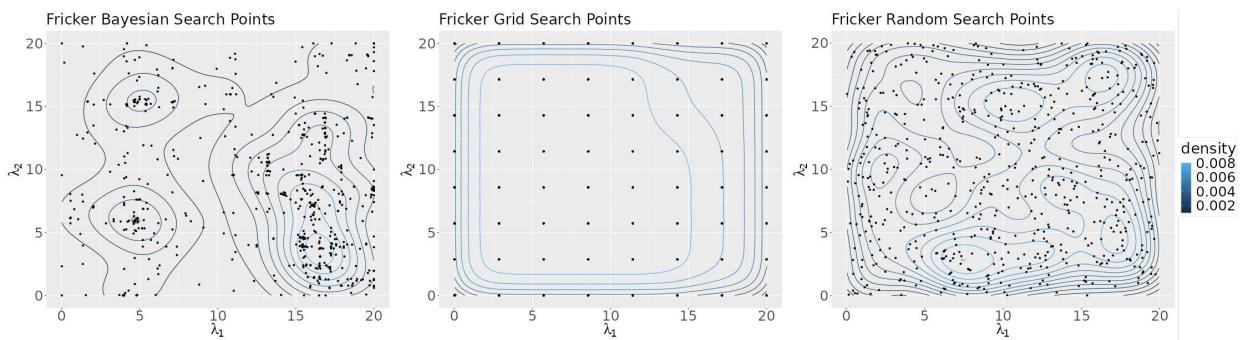


Fig. 10. The density plots of the points selected by each search algorithm.

model validation set and test set F1-scores as function of the number of trials. The median of the validation distribution for all 3 algorithms increases up to 8 trials and remains relatively flat after 43 trials. Figure 11 shows the distribution of the model validation set and test set F1-scores as function of the number of trials. The median of the validation distribution for all 3 algorithms increases up to 8 trials and remains relatively flat after 43 trials. The Bayesian optimization

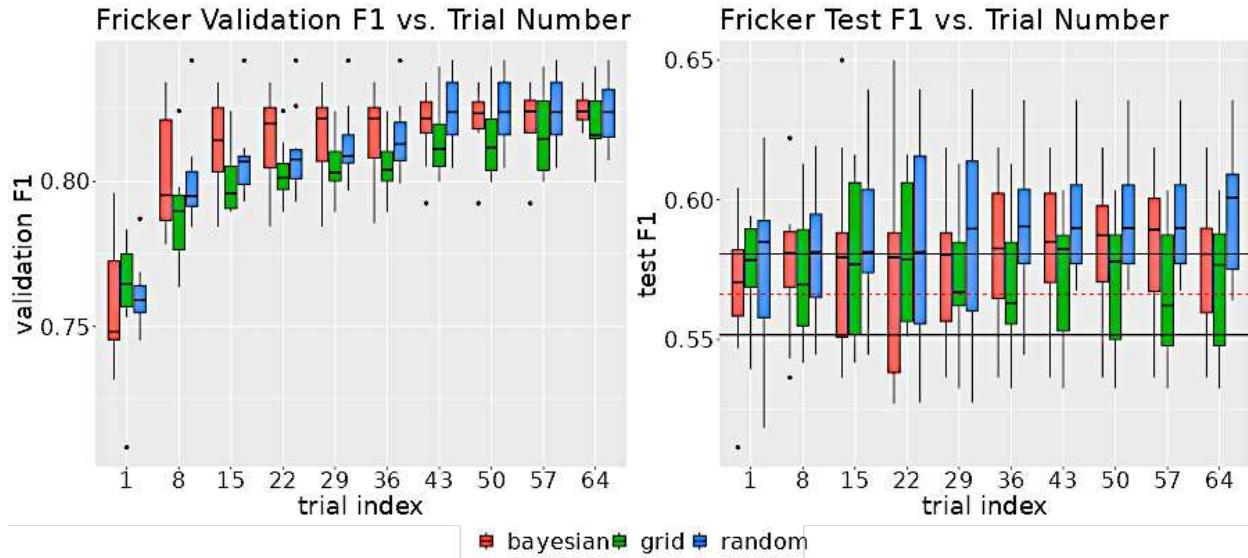


Fig. 11. The left graph shows the distribution of the validation F1 scores for the indicated number of trials for each search algorithm. The right graph shows the distribution of test F1 scores for the indicated number of trials. The hyperparameters selected for the test model were the ones corresponding to the highest validation F1 score up to the indicated trial. In the test graph, the two horizontal black lines indicate the 95% confidence intervals for non-neuro-symbolic test set F1 scores. The horizontal dashed red line indicates the mean of non-neuro-symbolic model test set F1 scores.

model improved at a faster rate than the other two algorithms. There was a great deal of overlap between the test set F1 scores for the non-neuro-symbolic model and the neuro-symbolic model due to the small improvement in performance from the rule implementation that is roughly of the same order as the model variance.

Figure 7 plots the neuro-symbolic model's validation F1 scores over the rule lambda search space. Larger λ_1 values correspond with higher validation scores suggesting that rule 1 is more impactful than rule 2. This is surprising since rule 1 is for the black oak class which has the fewest number of instances in the dataset. The highest validation scores are concentrated in the lower right corner of the search space which corresponds to the area focused on by the Bayesian optimization algorithm.

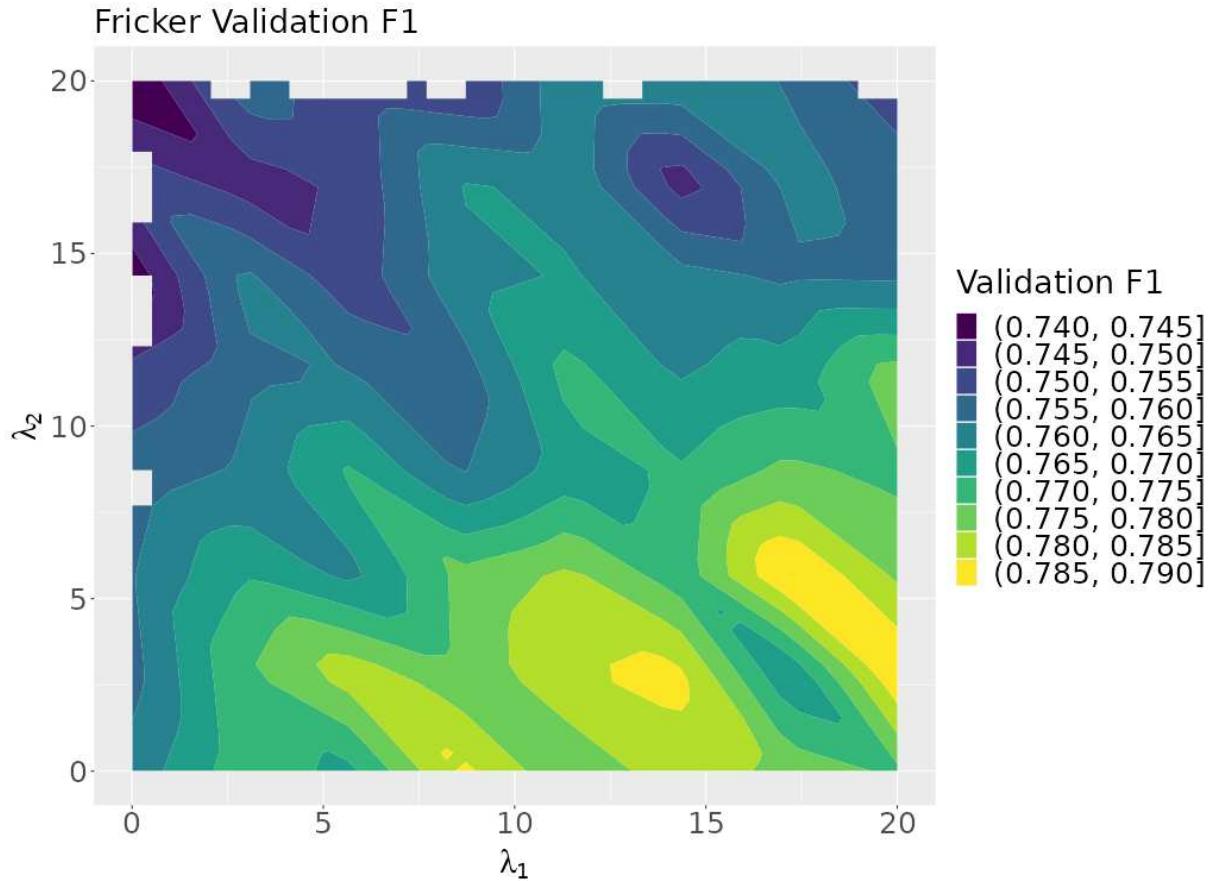


Fig. 12. Validation F1 as a function of λ_1 and λ_2 .

VI. DISCUSSION

We explored the application of our neuro-symbolic framework to crown delineation and species classification. Use of the neuro-symbolic framework increased of the accuracy of both crown delineation and species classification over non-neuro symbolic deep learning algorithms. The increase was modest but found for all 4 rules that were used. This indicates that knowledge about tree crown size and height-crown area allometry, translated into simple rules, can improve crown delineation from image data alone. In addition, simple formulations about heights of co-occurring canopy species can improve species classification. Additional rules developed from ecological knowledge, for example which species are likely to co-occur together or in different elevations, may additionally improve model performance and be tailored to local ecological contexts. Ecological knowledge about natural forests may be translated into FOL rules, although the benefit in classification accuracy of adding these rules needs to be tested more broadly across

more forest types, including closed canopy broadleaf forests. In addition, specific knowledge about management, such as planting spacing, could potentially generate useful rules in plantation forests. The 4 rules that we demonstrated cover the formulation that domain knowledge is likely to take when converted into FOL, using implication and if-and-only-if. However, more complex rules may require more creativity on the part of the user to implement. Nevertheless, the tools needed to express more complex ideas are present within the framework.

Additionally, we demonstrated the use of 3 hyperparameter tuning algorithms to find optimal rule parameters. All 3 algorithms were able to find sets of parameters that improved the neuro-symbolic models performance over the baseline model. Parameters found by the random search algorithm outperformed those found with Bayesian optimization and the grid search algorithm for both models explored. No method was consistently more efficient than any other, which is surprising due to the closed loop nature of the Bayesian optimization algorithm. Our results are partially in line with the findings of [53].

The conventional wisdom is that like machine learning models, there is no best hyperparameter optimization method; each method has its pros and cons. Bayesian optimization is efficient in high dimensional search spaces, but more complex than grid search and random search. Grid search is often only practical for small or low dimensional spaces. One of the greatest advantages of the random search algorithm is that it is easy to parallelize. In our example use cases we only explored search spaces with 2 dimensions which may be an insufficiently powerful test to discriminate search method efficiencies and we did not explore time complexity at all.

Neuro-symbolic models also hold promise in enabling model transferability across sites. Harmon et al. 2022 showed that neurosymbolic models (in this case for crown delineation) can be fine-tuned to improve crown delineation at different sites. Some rules, such as constraining delineation based on mean crown sizes per site, improved delineation scores at all sites. Other rules (constraining delineation based on height to crown allometries as in this study) only improved delineation accuracy at some sites. As in the case of this study, levels of improvement are dependent on many factors, including crown and species characteristics, quality of data at each site and other factors.

VII. CONCLUSIONS

We showed that our framework is applicable to crown delineation and species classification models. We demonstrated the use of our neuro-symbolic framework and tested 3 of the most

common automated hyperparameter tuning algorithms abilities to find optimal rule parameters. We provided a straightforward method for turning domain knowledge into a coded set of rules. The use of the automated tuning allows the user to test multiple rules at once and determine which rule is better suited for a dataset. Manually tuning models is labor intensive and more of an art than a science. Our framework provides an ad hoc method to convert deep learning models into more powerful neuro-symbolic models.

Neuro-symbolic models give users a way to ensure their model learns concepts that may not be easily extracted from the training data due to data scarcity, noise, or other factors. Creative use of rules encoded from domain knowledge is a potential cure for some of the short comings of ML models used for crown delineation and species classification. It also provides a way to utilize ecological knowledge of sites and species characteristics to improve delineation and classification models.

AUTHOR CONTRIBUTIONS

Conceptualization, D.Z.W and I.H.; methodology, I.H., D.Z.W. and E.W.; software, I.H.; validation, I.H.; formal analysis, I.H.; investigation, I.H.; resources, D.Z.W.; data curation, B.W. and I.H.; writing—original draft preparation, I.H.; writing—review and editing, E.W., S.B. and D.Z.W.; visualization, I.H.; supervision, D.Z.W.; project administration, E.W.; funding acquisition, E.W., S.B., D.Z.W. All authors have read and agreed to the published version of the manuscript.

FUNDING

This work was made possible by National Science Foundation (NSF) grant 1926542, USDA National Institute of Food and Agriculture McIntire Stennis project 1024612 to SAB, and USDA National Institute of Food and Agriculture, Hatch project FLA-WEC-005944 to EPW.

DATA AVAILABILITY

The data presented in this study are available on Zenodo at <https://zenodo.org/records/14194555>. These data were derived from the following resources available in the public domain: <https://data.neonscience.org/data-products/explore>.

CONFLICTS OF INTEREST

The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

REFERENCES

- [1] F. E. Fassnacht, J. C. White, M. A. Wulder, and E. Næsset, “Remote sensing in forestry: current challenges, considerations and directions,” *Forestry: An International Journal of Forest Research*, vol. 97, no. 1, pp. 11–37, 2024.
- [2] A. Kangas, R. Astrup, J. Breidenbach, J. Fridman, T. Gobakken, K. T. Korhonen, M. Maltamo, M. Nilsson, T. Nord-Larsen, E. Næsset *et al.*, “Remote sensing and forest inventories in nordic countries—roadmap for the future,” *Scandinavian Journal of Forest Research*, vol. 33, no. 4, pp. 397–412, 2018.
- [3] P. J. Curran, “Remote sensing of foliar chemistry,” *Remote sensing of environment*, vol. 30, no. 3, pp. 271–278, 1989.
- [4] L. Singh, O. Mutanga, P. Mafongoya, K. Peerbhay, and J. Crous, “Hyperspectral remote sensing for foliar nutrient detection in forestry: A near-infrared perspective,” *Remote Sensing Applications: Society and Environment*, vol. 25, p. 100676, 2022.
- [5] Q. Chen, G. V. Laurin, J. J. Battles, and D. Saah, “Integration of airborne lidar and vegetation types derived from aerial photography for mapping aboveground live biomass,” *Remote Sensing of Environment*, vol. 121, pp. 108–117, 2012.
- [6] A. Ghiyamat and H. Z. Shafri, “A review on hyperspectral remote sensing for homogeneous and heterogeneous forest biodiversity assessment,” *International Journal of Remote Sensing*, vol. 31, no. 7, pp. 1837–1856, 2010.
- [7] S. Marconi, S. J. Graves, D. Gong, M. S. Nia, M. Le Bras, B. J. Dorr, P. Fontana, J. Gearhart, C. Greenberg, D. J. Harris *et al.*, “A data science challenge for converting airborne remote sensing data into ecological information,” *PeerJ*, vol. 6, p. e5843, 2019.
- [8] Y. Ke and L. J. Quackenbush, “A review of methods for automatic individual tree-crown detection and delineation from passive remote sensing,” *International Journal of Remote Sensing*, vol. 32, no. 17, pp. 4725–4747, 2011.
- [9] C. Quintano, A. Fernández-Manso, Y. E. Shimabukuro, and G. Pereira, “Spectral unmixing,” *International Journal of Remote Sensing*, vol. 33, no. 17, pp. 5307–5340, 2012.
- [10] J. Lévesque and D. J. King, “Spatial analysis of radiometric fractions from high-resolution multispectral imagery for modelling individual tree crown and forest canopy structure and health,” *Remote Sensing of Environment*, vol. 84, no. 4, pp. 589–602, 2003.
- [11] M. S. Watt, G. D. Pearse, J. P. Dash, N. Melia, and E. M. C. Leonardo, “Application of remote sensing technologies to identify impacts of nutritional deficiencies on forests,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 149, pp. 226–241, 2019.
- [12] R. E. Martin, G. P. Asner, E. Francis, A. Ambrose, W. Baxter, A. J. Das, N. R. Vaughn, T. Paz-Kagan, T. Dawson, K. Nydick *et al.*, “Remote measurement of canopy water content in giant sequoias (*sequoiadendron giganteum*) during drought,” *Forest Ecology and Management*, vol. 419, pp. 279–290, 2018.
- [13] J. Zhang, B. Rivard, A. Sánchez-Azofeifa, and K. Castro-Esau, “Intra-and inter-class spectral variability of tropical tree species at la selva, costa rica: Implications for species identification using hydice imagery,” *Remote Sensing of Environment*, vol. 105, no. 2, pp. 129–141, 2006.
- [14] M. L. Clark, D. A. Roberts, and D. B. Clark, “Hyperspectral discrimination of tropical rain forest tree species at leaf to crown scales,” *Remote sensing of environment*, vol. 96, no. 3-4, pp. 375–398, 2005.
- [15] B. G. Weinstein, S. Marconi, S. J. Graves, A. Zare, A. Singh, S. A. Bohlman, L. Magee, D. J. Johnson, P. A. Townsend, and E. P. White, “Capturing long-tailed individual tree diversity using an airborne imaging and a multi-temporal hierarchical model,” *Remote Sensing in Ecology and Conservation*, vol. 9, no. 5, pp. 656–670, 2023.
- [16] H. Qin, W. Zhou, Y. Yao, and W. Wang, “Individual tree segmentation and tree species classification in subtropical broadleaf forests using uav-based lidar, hyperspectral, and ultrahigh-resolution rgb data,” *Remote Sensing of Environment*, vol. 280, p. 113143, 2022.

- [17] S. K. Alavipanah, M. Karimi Firozjaei, A. Sedighi, S. Fathololoumi, S. Zare Naghadehi, S. Saleh, M. Naghdizadegan, Z. Gomeh, J. J. Arsanjani, M. Makki *et al.*, “The shadow effect on surface biophysical variables derived from remote sensing: a review,” *Land*, vol. 11, no. 11, p. 2025, 2022.
- [18] M. Shahriari Nia, D. Z. Wang, S. A. Bohlman, P. Gader, S. J. Graves, and M. Petrovic, “Impact of atmospheric correction and image filtering on hyperspectral classification of tree species using support vector machine,” *Journal of Applied Remote Sensing*, vol. 9, no. 1, pp. 095 990–095 990, 2015.
- [19] D. G. Leckie, N. Walsworth, and F. A. Gougeon, “Identifying tree crown delineation shapes and need for remediation on high resolution imagery using an evidence based approach,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 114, pp. 206–227, 2016.
- [20] K. Yu, Z. Hao, C. J. Post, E. A. Mikhailova, L. Lin, G. Zhao, S. Tian, and J. Liu, “Comparison of classical methods and mask R-CNN for automatic tree detection and mapping using UAV imagery,” *Remote Sensing*, vol. 14, no. 2, p. 295, 2022, publisher: MDPI.
- [21] C. Zhang, K. Xia, H. Feng, Y. Yang, and X. Du, “Tree species classification using deep learning and rgb optical images obtained by an unmanned aerial vehicle,” *Journal of Forestry Research*, vol. 32, no. 5, pp. 1879–1888, 2021.
- [22] M. Beloiu, L. Heinzmann, N. Rehush, A. Gessler, and V. C. Griess, “Individual tree-crown detection and species identification in heterogeneous forests using aerial RGB imagery and deep learning,” *Remote Sensing*, vol. 15, no. 5, p. 1463, 2023, publisher: MDPI.
- [23] C. Sun, A. Shrivastava, S. Singh, and A. Gupta, “Revisiting unreasonable effectiveness of data in deep learning era,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 843–852.
- [24] L. Alzubaidi, J. Bai, A. Al-Sabaawi, J. Santamaría, A. S. Albahri, B. S. N. Al-dabbagh, M. A. Fadhel, M. Manoufali, J. Zhang, A. H. Al-Timemy, and others, “A survey on deep learning tools dealing with data scarcity: definitions, challenges, solutions, tips, and applications,” *Journal of Big Data*, vol. 10, no. 1, p. 46, 2023, publisher: Springer.
- [25] H. Zhao, J. Morgenroth, G. Pearse, and J. Schindler, “A systematic review of individual tree crown detection and delineation with convolutional neural networks (cnn),” *Current Forestry Reports*, vol. 9, no. 3, pp. 149–170, 2023.
- [26] X. Zhao, L. Wang, Y. Zhang, X. Han, M. Deveci, and M. Parmar, “A review of convolutional neural networks in computer vision,” *Artificial Intelligence Review*, vol. 57, no. 4, p. 99, 2024.
- [27] F. Xu, H. Uszkoreit, Y. Du, W. Fan, D. Zhao, and J. Zhu, “Explainable ai: A brief survey on history, research areas, approaches and challenges,” in *CCF international conference on natural language processing and Chinese computing*. Springer, 2019, pp. 563–574.
- [28] D. Minh, H. X. Wang, Y. F. Li, and T. N. Nguyen, “Explainable artificial intelligence: a comprehensive review,” *Artificial Intelligence Review*, pp. 1–66, 2022.
- [29] A. S. D. Garcez, L. C. Lamb, and D. M. Gabbay, *Neural-symbolic cognitive reasoning*. Springer Science & Business Media, 2008.
- [30] P. Hitzler, A. Eberhart, M. Ebrahimi, M. K. Sarker, and L. Zhou, “Neuro-symbolic approaches in artificial intelligence,” *National Science Review*, vol. 9, no. 6, p. nwac035, 2022.
- [31] E. Giunchiglia, M. C. Stoian, and T. Lukasiewicz, “Deep learning with logical constraints,” *arXiv preprint arXiv:2205.00523*, 2022.
- [32] H. Xu, G. Qi, J. Li, M. Wang, K. Xu, and H. Gao, “Fine-grained Image Classification by Visual-Semantic Embedding.” in *IJCAI*, 2018, pp. 1043–1049.
- [33] G. Sumbul, R. G. Cinbis, and S. Aksoy, “Fine-grained object recognition and zero-shot learning in remote sensing imagery,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 2, pp. 770–779, 2017, publisher: IEEE.

- [34] S. Seo, S. Arik, J. Yoon, X. Zhang, K. Sohn, and T. Pfister, “Controlling Neural Networks with Rule Representations,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [35] Z. Hu, X. Ma, Z. Liu, E. Hovy, and E. Xing, “Harnessing deep neural networks with logic rules,” *arXiv preprint arXiv:1603.06318*, 2016.
- [36] M. Diligenti, M. Gori, and C. Sacca, “Semantic-based regularization for learning and inference,” *Artificial Intelligence*, vol. 244, pp. 143–165, 2017.
- [37] E. van Krieken, E. Acar, and F. van Harmelen, “Analyzing differentiable fuzzy logic operators,” *Artificial Intelligence*, vol. 302, p. 103602, 2022.
- [38] P. Probst, A.-L. Boulesteix, and B. Bischl, “Tunability: Importance of hyperparameters of machine learning algorithms,” *Journal of Machine Learning Research*, vol. 20, no. 53, pp. 1–32, 2019.
- [39] L. Yang and A. Shami, “On hyperparameter optimization of machine learning algorithms: Theory and practice,” *Neurocomputing*, vol. 415, pp. 295–316, 2020.
- [40] B. G. Weinstein, S. Marconi, M. Aubry-Kientz, G. Vincent, H. Senyondo, and E. P. White, “DeepForest: A Python package for RGB deep learning tree crown delineation,” *Methods in Ecology and Evolution*, vol. 11, no. 12, pp. 1743–1751, dec 2020. [Online]. Available: <https://onlinelibrary.wiley.com/doi/10.1111/2041-210X.13472>
- [41] G. A. Fricker, J. D. Ventura, J. A. Wolf, M. P. North, F. W. Davis, and J. Franklin, “A convolutional neural network classifier identifies tree species in mixed-conifer forest from hyperspectral imagery,” *Remote Sensing*, vol. 11, no. 19, p. 2326, 2019.
- [42] T. U. Kampe, B. R. Johnson, M. A. Kuester, and M. Keller, “NEON: the first continental-scale ecological observatory with airborne remote sensing of vegetation canopy biochemistry and structure,” *Journal of Applied Remote Sensing*, vol. 4, no. 1, p. 043510, 2010, publisher: SPIE.
- [43] “Niwot Ridge NEON,” <https://www.neonscience.org/field-sites/niwo>, Jan 2022. [Online]. Available: <https://www.neonscience.org/field-sites/niwo>
- [44] “Teakettle Experimental Forest,” <https://www.fs.fed.us/psw/ef/teakettle/>, Jan 2022. [Online]. Available: <https://www.fs.fed.us/psw/ef/teakettle/>
- [45] J. Ansel, E. Yang, H. He, N. Gimelshein, A. Jain, M. Voznesensky, B. Bao, P. Bell, D. Berard, E. Burovski *et al.*, “Pytorch 2: Faster machine learning through dynamic python bytecode transformation and graph compilation,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, 2024, pp. 929–947.
- [46] W. Falcon, J. Borovcak, A. Wälchli, N. Eggert, J. Schock, J. Jordan, N. Skafte, V. Bereznuk, E. Harris, T. Murrell *et al.*, “Pytorchlightning/pytorch-lightning: 0.7. 6 release,” *Zenodo*, 2020.
- [47] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, “Botorch: Bayesian optimization in pytorch,” Technical report, Tech. Rep., 2019.
- [48] A. Kimmig, S. Bach, M. Broecheler, B. Huang, and L. Getoor, “A short introduction to probabilistic soft logic,” in *Proceedings of the NIPS Workshop on Probabilistic Programming: Foundations and Applications*, 2012, pp. 1–4.
- [49] E. P. Klement, R. Mesiar, and E. Pap, *Triangular norms*. Springer Science & Business Media, 2013, vol. 8.
- [50] M. Diligenti, S. Roychowdhury, and M. Gori, “Integrating prior knowledge into deep learning,” in *2017 16th IEEE international conference on machine learning and applications (ICMLA)*. IEEE, 2017, pp. 920–923.
- [51] S. Roychowdhury, M. Diligenti, and M. Gori, “Regularizing deep networks with prior knowledge: A constraint-based approach,” *Knowledge-Based Systems*, vol. 222, p. 106989, 2021.

- [52] “Pytorch 2.4 documentation,” <https://pytorch.org/docs/stable/nn.html>, Aug 2024. [Online]. Available: <https://pytorch.org/docs/stable/nn.html>
- [53] J. Bergstra and Y. Bengio, “Random search for hyper-parameter optimization.” *Journal of machine learning research*, vol. 13, no. 2, 2012.
- [54] P. I. Frazier, “A tutorial on bayesian optimization,” *arXiv preprint arXiv:1807.02811*, 2018.
- [55] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2980–2988.
- [56] V. M. Scholl, M. E. Cattau, M. B. Joseph, and J. K. Balch, “Integrating national ecological observatory network (neon) airborne remote sensing and in-situ data for optimal tree species classification,” *Remote Sensing*, vol. 12, no. 9, p. 1414, 2020.
- [57] B. G. Weinstein, S. J. Graves, S. Marconi, A. Singh, A. Zare, D. Stewart, S. A. Bohlman, and E. 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*, vol. 17, no. 7, p. e1009180, jul 2021, publisher: Public Library of Science. [Online]. Available: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1009180>
- [58] I. Harmon, S. Marconi, B. Weinstein, S. Graves, D. Z. Wang, A. Zare, S. Bohlman, A. Singh, and E. White, “Injecting domain knowledge into deep neural networks for tree crown delineation,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–19, 2022.
- [59] T. Jucker, J. Caspersen, J. Chave, C. Antin, N. Barbier, F. Bongers, M. Dalponte, K. Y. van Ewijk, D. I. Forrester, M. Haeni, and others, “Allometric equations for integrating remote sensing imagery into forest monitoring programmes,” *Global change biology*, vol. 23, no. 1, pp. 177–190, 2017, publisher: Wiley Online Library.
- [60] C. M. Hulshof, N. G. Swenson, and M. D. Weiser, “Tree height–diameter allometry across the united states,” *Ecology and evolution*, vol. 5, no. 6, pp. 1193–1204, 2015.
- [61] I. Harmon, S. Marconi, B. Weinstein, Y. Bai, D. Z. Wang, E. White, and S. Bohlman, “Improving rare tree species classification using domain knowledge,” *IEEE Geoscience and Remote Sensing Letters*, vol. 20, pp. 1–5, 2023.
- [62] M. P. North, *Vegetation and ecological characteristics of mixed-conifer and red fir forests at the Teakettle Experimental Forest*. US Department of Agriculture, Forest Service, Pacific Southwest Research Station, 2002, vol. 186.