

NIST DSE Plant Identification with Remote Sensing Evaluation Plan

Version 0.3, Updated on August 29, 2017

1. Introduction

This document defines the tasks and evaluation plan for the *Data Science for Plant Identification with Remote Sensing* component of the National Institute of Standards and Technology (NIST) Data Science Evaluation (DSE) Series Pilot.

Understanding ecological patterns and processes across geographical scales is crucial to understanding the effects of environmental change on natural systems and human society. However, financial and logistic limitations restrict the scales at which ecological data can be collected by field ecologists on the ground. Recent increases in the availability of remotely sensed imagery from satellites and aircraft provide the potential to observe ecological data at much larger scales than are possible through traditional data collection methods. This DSE focuses on inferring ecological information from remote sensing data.

The DSE focuses on data from the National Science Foundation sponsored National Ecological Observatory Network (NEON). NEON collects continental-scale ecological observations from 81 sites across the United States. Data products range from ground based data on plant and animal locations and traits, to landscape scale airborne remote sensing. Dealing with the volume, velocity, and variety of this data will require interdisciplinary approaches involving ecology, computer science, statistics, and data science. This NIST Data Science Evaluation (DSE) competition is an applied, interdisciplinary, challenge using multiple components of the NEON data stream. Aligning this DSE with NEON has the potential to result in pioneering interdisciplinary approaches and large gains in our understanding of the natural world by engaging the larger Data Science community in this unique data source. In addition, many of the algorithms and techniques are expected to generalize to other problems in data science.

This competition will combine airborne remote sensing data with field-based tree measurements, species identifications, and crown segmentation, to develop and evaluate models to address three common data science tasks (segmentation, alignment, and classification) to estimate tree location, size, and species identity from the remote sensing data.

1. **Segmentation:** Locate objects, or boundaries surrounding objects, in images. Estimate the size, shape, and location of individual tree crowns (the top, sun-exposed portion of the tree visible from above) from remote sensing data.
2. **Alignment:** Associate two or more representations of the same object. Match individual trees measured on the ground with the same tree observed using remotely sensed data.
3. **Classification:** Determine which of a set of categories an object belongs to. Determine the species identity (i.e., the type) of each individual tree from remotely sensed data.

These tasks are designed so that they can be completed independently, but also so that they can be combined to form a data processing pipeline for real world applications.

2. Data

Seven datasets will be used as part of this evaluation:

- Remote Sensing Data -- 43 plots totalling ~1GB of imagery
 - NEON Photographs - 0.0625 m² resolution photographic images (RGB)
 - NEON LiDAR data - 3-dimensional point clouds of the height of trees and ground, which are also provided as 1 m² resolution raster images of the estimated height of vegetation above the ground (Canopy Height Model; CHM).
 - NEON Hyperspectral images - 1 m² resolution raster images with 426 wavelength bands from 380-2510 nm. These data are georectified and atmospherically corrected, and bands that are strongly influenced by the absorbance by atmospheric water (and should be removed in most cases) will be specified.
- Ground Data -- 88 trees
 - NEON Vegetation Structure data - Ground based positions of individual trees, the size of each trunk, and for some trees two measurements of crown size (maximum crown diameter and the diameter of the crown perpendicular to the axis of the maximum diameter)
 - NEON Species Data - the genus and species (type) of each tree
- Individual Tree Crown segmentations (ITC) Data -- 626 stems
 - ITC data - Vector polygon data of tree crown boundaries and vector point data of trunk locations both mapped in the field directly onto remote sensing images. Boundaries and locations are drawn onto remote sensing images from the ground while observing the trees directly and then fine tuned to match clear crown edges in the images. The crown data include confidence scores qualitatively estimating the precision of the tree crown mapping.
 - Correspondence between ITC crowns and ground stems

Spatial data will be provided in the UTM 17N coordinate system. The projection can generally be ignored at this scale of analysis and *x* and *y values* treated as any standard grid.

Figure 1 illustrates these core data types. All data collected by NEON are publicly available with open licenses. The subset of data needed for this DSE will be posted on the DSE website for download.

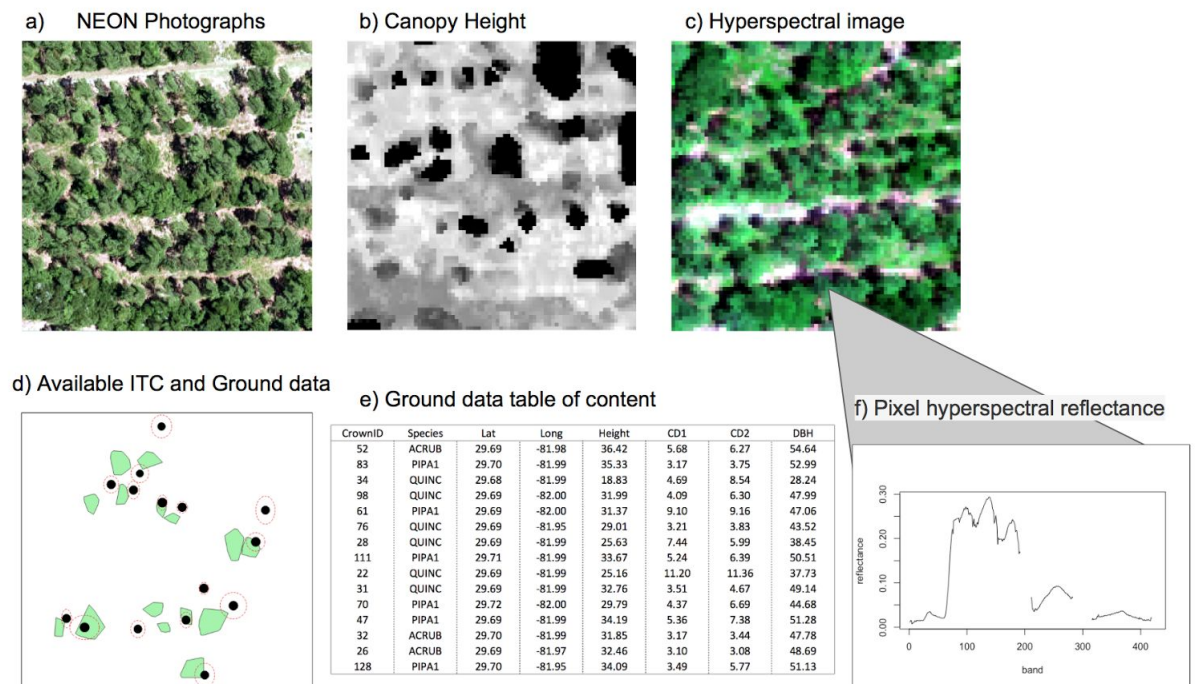
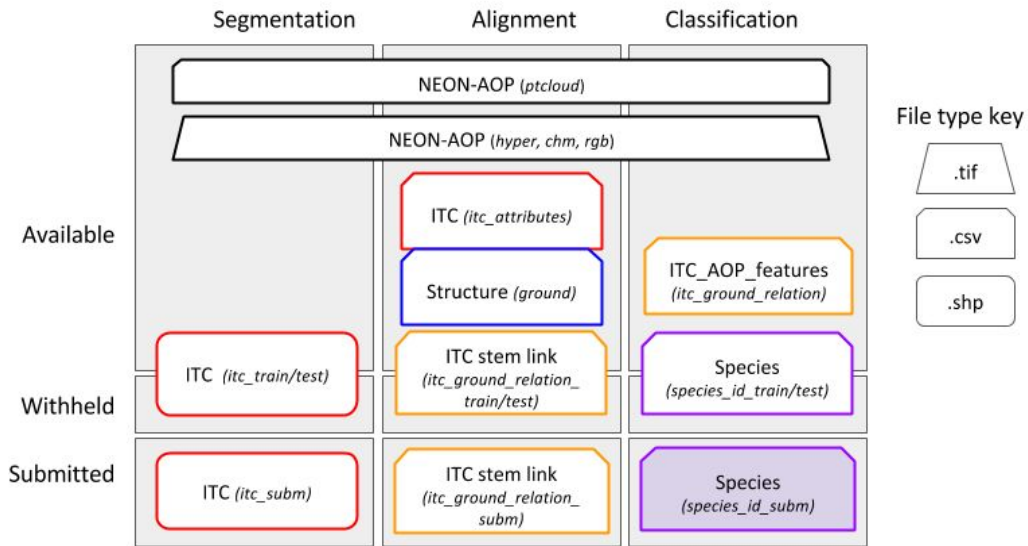


Figure 1. Example data products provided in the contest. Top row represents the remote sensing data: a) is the RGB image of a 80x80m plot, b) a CHM for the same region, c) a representation of the Hyperspectral data in false colors (a mixture of 3 hyperspectral bands; R = band 19, G = band 34, B = band 58). d) represents vectorial spatial data: green polygons ITCs collected for the area, black dots NEON ground data, dots' size is scaled by tree DBH. Red ellipse represent expected crown area, given the widest crown diameter and its orthogonal one. e) is a sample of the attributes associated to ground data. f) is an example of spectral signature for a given hyperspectral pixel.

Location: All data is from the Ordway-Swisher Biological Station (OSBS) NEON site. Field measures and remote sensing data are for 23 OSBS NEON distributed plots (1600 m²) and 20 OSBS NEON tower plots (1600 m²), which include hardwood/woody wetlands and upland pine/sandhill ecosystems.

Input/Output data types for Each tasks:

The segmentation task uses the remote sensing data and 70% of ITC stem data for training and 30% of the ITC stem data for testing. The alignment task uses the remote sensing data, the ITC data, and ground data on stem diameter, height, and location. This task uses only trees whose ITCs are in the training set of the segmentation task with 70% of those stems used for training and 30% for testing. The classification task uses the subset of remote sensing pixels falling within the ITC-stem polygons. Data are in tabular form, with each row representing a pixel of the remote sensing data. Features for each pixel include ID, height, and reflectance for the 426 hyperspectral bands. A different csv file contains species identification for 70% of the ITCs for which species has been identified. The remaining 30% of species identification data is held as testing.



ITC (<i>itc_train/test</i>)	crow_n_id	confidence	plot_ID				
	391	high	OSBS_002				
Structure (<i>ground</i>)	stem_id	tag_E	tag_N	height	diameter	cr_max	cr_perp
	1771	45963	32585	23.3	12.2	3.5	2.1
ITC (<i>itc_attributes</i>)	plot_ID	crow_n_id	ITC_E	ITC_N			
	OSBS_002	391	45960	32595			
ITC stem link (<i>itc_ground_relation_train/test</i>)	stem_id	crow_n_id					
	1771	391					
Species (<i>species_id_train/test</i>)	crow_n_id	species	genus	species_id	genus_id		
	391	<i>Quercus nigra</i>	<i>Quercus</i>	QUNI	QU		
ITC_AOP_features (<i>hyper_bands_train/test</i>)	stem_id	CHM	band_1	...	band_426		
	391	23.3	0.42	0.73	1.00		

Figure 2: Training/testing Input/output data types for each of the three tasks. Data in both the train and test boxes will be split either at the plot level (segmentation and alignment tasks) on the stem level with 70% of the plots/stems used for training and 30% for testing. Products to be submitted are for the test data. Names in italics correspond to the names of the files provided or to be submitted. **NEON-AOP (black)**: Spatial data of the 4 types of NEON-AOP data. **ITC (red)**: Spatial data of Individual Tree Crowns and a unique crown ID number. **ITC stem link (orange)**: Table of unique crown ID number and the NEON ground data stem ID number. **Structure (blue)**: Table of NEON ground data with tree stem attributes; unique stem ID, latitude & longitude coordinates, stem diameter, crown dimensions. **Species (purple)**: Table of NEON ground data with assignment to two class categories (species class and genus class). Species product is a table of the probability of assignment of each ITC for each class in both class categories.

3. Segmentation

Segmentation is a common Data Science problem that involves locating objects or boundaries in images. In addition to being useful on its own, good segmentation is often important for classification of objects within images.

Summary: Segmentation

Inputs: Hyperspectral images, RGB images, and LiDAR point cloud data.

Output: Shapefiles containing estimates of the position and shape of individual tree crowns (ITC data).

Training Files: *ITC_OSBS_plot.shp*, *OSBS_plot_hyper.tiff*, *hyper_bands.csv*, *OSBS_plot_chm.tiff*, *ptcloud_plot.csv*, *rgb_plot.tiff*

Test Files: *hyper_plot.tiff*, *chm_plot.tiff*, *ptcloud_plot.csv*, *rgb_plot.tiff*

Submission File: *itc_subm_plot.shp* (including all 4 sub files; .shp, .shx, .dbf, .prj)

3.1 Segmentation in the ecology domain

Identifying the position and size of individual trees from remote sensing is useful for understanding forest structure and an important first step in species classification. It is also a complex version of the common image segmentation task because trees often overlap each other and look similar, and because the available data is heterogeneous, involving many bands, multiple resolutions, and point cloud height data. Participants will use Hyperspectral, LiDAR and high resolution RGB photographs to locate and segment individual crowns in multiple plots.

3.2 Training Data

Input for the training phase will be data for a subset of plots at Ordway Swisher and will include shapefiles with polygons for each individual tree crown (ITC) along with associated hyperspectral, LiDAR, and RGB data. LiDAR data provides information on the spatial variation in canopy height that may allow partitioning of crowns of neighboring trees with similar spectral signatures. Hyperspectral data allows development of spectral signatures to identify object categories (e.g. by assigning spectrally-similar categories to the same cluster). RGB photographs provide finer resolution information (0.25x0.25 m as opposed to 1.0x1.0 m for the hyperspectral data), which may be helpful to separate trees that are close to one another and to refine boundary placement. Using the remote sensing data, participants will build models to estimate crown boundaries in plots where ITC polygons are not available. The ITC data can be considered as ground truth, since polygons of individual crowns were manually delineated in the field by field experts and drawn directly on remote sensing images.

The input data consists of two types of data: a shape file containing ITC data for 70% of the plots and remote sensing data for all plots, including one of the following for each plot: Hyperspectral (geotiff format), LIDAR (geotiff and csv), and RGB (geotiff). File naming will be of the form *itc_train_plot.shp*, *hyper_plot.tiff*, *chm_plot.tiff*, *ptcloud_plot.csv*, *rgb_plot.tiff* where *plot* indicates the numeric code of the plot of each plot.

- In each ITC shape file, each attribute will be an individual crown and its associated crown ID. Each “shape file” is composed of four files (.shp, .shx, .dbf, .prj), which can be read with widely-available packages (including most GIS software and GDAL).
- Each of the geotiff files contain the values for each band, the dimensions of the 2D or 3D space, and the spatial location information.
 - Hyperspectral - 426 bands on an 80 x 80 pixel lattice at 1 m² resolution
 - CHM (canopy height model) - 1 band on an 80 x 80 pixel lattice at 1 m² resolution
 - RGB - 3 bands on a 320 x 320 pixel lattice at 0.0625 m² resolution
- Information on hyperspectral bands, including which bands should be ignored, will be provided in a 3 column csv file named hyper_bands.csv.
 - *band_name*: the name of the band formatted as *band_number*
 - *nanometer*: the spectrum value
 - *noise*: 0 indicates that the band is good to use, 1 indicates that it is strongly influenced by the absorbance of atmospheric water and should be excluded from most analyses
- The point cloud data will be provided as a 3 column csv file with the x, y, and z coordinates of each LIDAR return/point.
 - *x*: utm coordinates *m*
 - *y*: utm coordinates in *m*
 - *z*: height in *m* above sea level

3.3 Test Data

Data provided for testing will be the hyperspectral, LiDAR, and RGB data described above for the plots not provided in the training dataset. ITC data will not be provided for the test set, as the goal is for participants to predict its contents from the remote sensing data.

3.4 Submission Data

The submission files will include one shape file (including all 4 subfiles) for each test plot named as *itc_subm_plot.shp*.

- In the ITC shape file, each attribute will be an individual crown and a unique crown ID.

3.5 Performance Metrics

Participants’ performance on each plot will be calculated as the mean pairwise Jaccard Coefficient, $J(A,B)$, calculated for each predicted polygon with the corresponding ground truth polygon. When there is ambiguity with regard to which submitted polygon corresponds to which ITC, we will choose the mapping that gives the best score to the participant.

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|}$$

The score for this task is the average of the plot level scores, which are themselves the average scores of the polygons within each plot.

This method is simple, does not require assignment of predicted crowns to specific ITCs by the participants, and has good continuity since the Jaccard coefficient has good continuity.

The output for the crown delineation algorithm is expected to produce polygons with no overlap, but if there are cases of overlapping crowns in the submitted polygons, they will be disambiguated by disregarding the intersection in the numerator of the Jaccard Coefficient. Thus, it is in the participant's best interest to remove areas of intersection by dividing those areas among the overlapping crowns.

4. Alignment

Alignment is a common Data Science problem that involves associating two or more representations of the same object. One common instance of this problem is identifying the same object in data collected from different instruments or data sources.

Summary: Alignment

Inputs: ITC data, Ground data. *Optional:* Hyperspectral images, RGB images, LiDAR data

Output: Labels relating the ITC crown ID to the Ground data stem ID for each tree

Training Files: ground_train.csv, itc_ground_relation_train.csv, *Optional:* hyper_plot.tiff, hyper_bands.csv, chm_plot.tiff, ptcloud_plot.csv, rgb_plot.tiff

Test Files: itc_ground_relation_test.shp, ground_test.csv, *Optional:* hyper_plot.tiff, chm_plot.tiff, ptcloud_plot.csv, rgb_plot.tiff

Submission Files: itc_ground_relation_subm_plot.csv

4.1 Alignment in the ecology domain

To build models relating remote sensing to field data, it is necessary to align the ground data to the remote sensing images. The three remote sensing datasets have been georeferenced and pre-aligned with one another by NEON. The NEON Ground data are georeferenced separately from the remote sensing data using a combination of GPS locations for fixed points and surveying methods to relate tree locations to those fixed points. Both the remote sensing and ground data coordinates involve measurement error, with the error likely being greater in the ground data. In addition, the ground data provides a location for the trunk of the tree, which is not directly visible in the remote sensing data. The trunk is not always located at the center of the crown because crowns are positioned in part to maximize light capture and avoid other tree crowns. As a consequence, apparent positions of trees can differ between the ground and remote sensing data sources with offsets of up to several meters even when no other errors have occurred with geo-referencing. Such offsets make it more difficult to develop models that related the ground and remote sensing data.

The goal of the task is to correctly align each tree crown polygon to a single tree from the ground data, thus allowing data collected on the ground (e.g. species identity, height, stem diameter) to be accurately associated with remote sensing data. This can be accomplished by considering the position and size of stems and crowns. Other information from the remote sensing images may also be helpful.

In real world situations this task will typically be conducted in an unsupervised manner matching segmented crowns to ground data using tree location and size without the use of ITC data for either

training or evaluation. In this DSE we have collected ITC data to allow us to: 1) make this task independent of the segmentation task by providing the ITC segmentation as a near perfect segmentation of the image; 2) evaluate the performance of the algorithms against the known mapping between the ITC crowns and the ground data stems; and 3) provide data that allows for supervised rather than unsupervised algorithm development.

4.2 Training Data

Input data will be for a subset of plots and will include ITC polygon data, NEON ground data (including the stem ID, location of the stem, diameter of the stem, and for some of the plots the maximum crown radius and the radius perpendicular to the axis of the maximum radius). Only ITC polygon data for which ground data is available will be included. The stem diameter of trees is positively related to the crown diameter, which should allow the size of the crown to be estimated from the diameter of the stem in cases where the crown diameter is not measured directly. The remote sensing data (Hyperspectral, RGB, and LiDAR) will also be available for optional use.

The primary input data consists of the ground and ITC data for 70% of the training plots from the segmentation task and will include: 1) one csv file (`ground.csv`) containing information on the location, stem size, crown size, and stem ID of each tree; 2) one csv (`itc_attributes.csv`) with attributes and centroid coordinates for each individual tree crown (ITC); and 3) one csv file per plot containing the stem ID and associated crown ID for each tree (`itc_ground_relation_train.csv`). The hyperspectral, canopy height, and camera data described in 3.2 will also be provided for optional use. The goal is to match stem ID from the ground data to the crown ID for the ITC tree crowns.

- Each ground csv file contains these comma separated fields:
 - *stem_id*: a unique identifier, corresponding to NEON tags on ground data
 - *tag_E*: utm coordinates in *m* collected on the ground
 - *tag_N*: utm coordinates in *m* collected on the ground
 - *diameter*: diameter of the stem in *cm* at a standard height above the ground (~1.3 m)
 - *height*: height of the tree measured from the ground. Empty when missing value
 - *cr_max*: maximum radius of the crown in *m*
 - *cr_perp*: radius of the crown in *m* in the direction perpendicular to the maximum radius
- Each `itc_attributes` csv file contains these comma separated fields:
 - *Plot_id*: name of the plot from which data have been collected
 - *ITC_id*: a unique identifier assigned to each itc polygon, and uncorrelated to *stem_id*
 - *ITC_E*: itc centroid utm coordinates in *m*
 - *ITC_N*: itc centroid utm coordinates in *m*
- Each `itc_ground_relation` relation file contains these comma separated fields
 - *stem_id*: the stem ID from the ground data
 - *crown_id*: the matching crown ID from the ITC data

4.3 Test Data

Data provided for testing will be the ground and ITC attribute data (and optionally the remote sensing data) for the 30% of the training plots from the segmentation task not provided in the training dataset. Relationship data will not be provided for the test set.

4.4 Submission Data

The submission files will include one csv file for each test plot named `itc_ground_relation_subm_plot.csv`. Each file must contain these comma separated fields

- *stem_id*: the stem ID from the ground data
- *crown_id*: the matching crown ID from the ITC data
- *probability*: the probability that the stem_id matches to the crown_id. The probabilities for a given stem_id will be normalized to sum to 1 if the submitted values do not already.

4.5 Performance Metrics

Performance of the pairing of field stems to ITC crowns will be evaluated using the trace of the prediction matrix divided by the sum over the values in that matrix. This method was chosen based on the following reasoning. In the testing stage, suppose we have a set of probe data (ITC) denoted as $\{p_n | n=1, \dots, N\}$, and ground truth data denoted as $\{g_n | n=1, \dots, N\}$. We know in advance that there is a unique one-to-one mapping between P and G sets. Without loss of generality, assume p_n should be mapped to g_n for $n=1, \dots, N$. For each probe data point p_i , a program predicts a non-negative confidence score that should be aligned with ground truth data point g_j , which forms a prediction matrix $M = (m_{ij})$ where $i, j = 1, \dots, N$. Then, the quality of prediction can be measured by the following scoring function:

$$score = \frac{trace(M)}{\sum_{i,j} m_{i,j}}$$

where $trace(\cdot)$ represents trace of a matrix and M represents the prediction matrix which has been aligned in the order which matches the ground truth.

5 Classification

Classification is a common Data Science problem that involves determining which of a set of categories an object belongs to.

Summary: Classification

Inputs: Ground data on species and genus IDs combined with Hyperspectral and CHM values for each pixel within the ITC for the tree

Output: A probability that each tree belongs to each possible species

Training Files: species_id_train.csv, crown_hyper_train.csv, hyper_bands.csv

Test Files: crown_hyper_test.csv

Submission Files: species_id_subm.csv

5.1 Classification in the ecology domain

A large number of ecological, environmental, and conservation oriented questions depend on species identification. This includes efforts to conserve individual species, understand and maintain biodiversity, and incorporate the biosphere into global circulation models. Being able to describe the density and distribution of different species using remote sensing would allow these efforts to occur more rapidly and

at larger scales than field sampling. The goal of this task is to classify trees in remote sensing data into species and genera (higher level groupings of species). In addition to its utility for the domain, this task represents a challenging version of general classification problems because it involves classifying different species with very similar spectral signatures and categorizing data where some categories (species) have only small samples in the training set (i.e. rare species).

This task involves using the Hyperspectral and CHM data from LIDAR to build classifiers for the ground based species identifications. To make this task independent of the segmentation and alignment tasks, and to maximize sample size, pre-aligned ITC data will be used to extract pixel level Hyperspectral and dCHM data for each pixel in all trees to be classified. Participants will determine probability that each tree belongs to each species and genus. Spectral reflectance signatures are the most commonly used data type for building species classifiers. Since classification is at the level of the crown, any pixel-level classification models must be upscaled to the crown.

One important criterion for evaluating a classification model is its ability to deal with species (or genera) from outside the training set. Thus, we also evaluate participants' classification models on their ability to correctly classify crowns as "Other" when they do not belong to any of the previously seen classes. The "Other" label may appear at only the species level (if the new species belongs to an existing genus) or at both the species and genus levels (if the new species belongs to a new genus). Since the "Other" class will include a diverse group of species, it will include a broader range of spectral signatures than the known classes will, providing an additional challenge for classification. Any class (genus or species) with less than 5 individuals in the full dataset is categorized as "Other" in both the test and training sets.

5.2 Training Data

Input data will be for 70% of the trees in the ITC data for which species were identified on the ground. Data will be provided in three files: 1) a csv file named `species_id_train.csv` containing the crown ID and associated genus and species for each tree; 2) a csv file named `crown_hyper_train.csv` containing the crown ID, canopy height (from the CHM data) and 426 hyperspectral bands for each pixel occurring within each crown; and 3) a csv file named `hyper_bands.csv` containing information on the hyperspectral bands. The goal is to determine the species and genus from the hyperspectral and CHM data for each crown.

- The species ID file contains these comma separated fields
 - *crown_id*: the matching crown ID from the ITC data
 - *genus*: the genus of the tree
 - *species*: the species of the tree
 - *genus_id*: a unique two letter code for each genus
 - *species_id*: a unique four letter code for each species
- The hyperspectral data file contains a total of 428 columns and one row per pixel
 - *crown_id*: the matching crown ID from the ITC data
 - *canopy_height*: the height of the tree in *m* above ground
 - 426 columns with hyperspectral band intensities named: `band_1`, `band_2` ... `band_426`
- Information on hyperspectral bands, including which bands should be ignored, will be provided in a 3 column csv file named `hyper_bands.csv`.
 - *band_name*: the name of the band formatted as `band_number`
 - *Nanometer*: the average wavelength of the 5nm wide *i*th spectral band
 - *noise*: 0 indicates that the band is good to use, 1 indicates that it is strongly influenced by the absorbance of atmospheric water and should be excluded from most analyses

5.3 Test Data

Test data will include the non-training hyperspectral data and the data on the hyperspectral bands.

5.4 Submission Data

Submissions will be two csv files named `species_id_subm.csv` and `genus_id_subm.csv` containing information on the crown ID and the species and genus identification probabilities for each crown. These files must contain these comma separated fields

- *crown_id*: the matching crown ID from the ITC data
- *ID*: the predicted genus or species ID (or “Other”)
- *probability*: the probability that the crown belongs to the associated genus or species. The probabilities for a given *crown_id* (including the “Other” category) will be normalized to sum to 1 if the submitted values do not already.

The files should contain one row for each possible *crown_id* and *genus/species* combination(i.e., the number of rows should be equal to the number of *crown_id*'s in the test set times the number of unique *IDs* (either species or genus) present in the training set). The results should be sorted in ascending order by *crown_id* and then *ID*. The dictionary of genus/species classes and their unique IDs will be provided. The performance metrics will be computed separately for genus and species classes.

5.5 Performance Metrics

We will evaluate performance using two metrics. The first one, rank-1 accuracy, is the fraction of crowns in the test set whose ground truth *species_id* (or *genus_id*) was assigned the highest probability by the participant. It is calculated as $\text{mean}(\text{argmax}_k(p_{ik})=g_i)$, where g_i is the ground-truth class of crown i , and p_{ik} is the probability assigned by the participant that crown i belongs to class k . This metric only considers whether the correct class has the highest probability, not whether the probabilities are well-calibrated.

The second metric, average cross-entropy, is defined as:

$$\text{cost} = \frac{-\sum_{n,k} \ln(p_{nk}) * \delta(g_n, k)}{N}$$

given that $p_{nk} \neq 0$, to avoid the singularity. The $\delta(x, y)$ is an indicator function that takes value 1 when $x = y$. This metric rewards participants for submitting well-calibrated probabilities that reflect their uncertainty about which crowns belong to which class.

Performance metrics will be evaluated on the two different class labels (genus and species) separately, resulting in four separate scores. Note that the “Other” label will be included alongside the biological species and genus names in calculating these performance metrics.

6. Submission Descriptions

Each submission should include a brief description for each task that includes:

1. *Task*. Segmentation, alignment, or classification.
2. *Team Affiliation*. The name of the submitting team.
3. *Method*. A brief summary the method used (a few sentences to one paragraph).
4. *Algorithmic Information*. A brief description of additional details about the methods and algorithms used (1-2 paragraphs).
5. *Data Sources*. A brief description of the data used including both the specific subsets of the provided data and any external data sets used.

7. Schedule

The key dates for the NIST evaluation will follow the scheme summarized in Table 1.

Registration Opens	July 3, 2017
Release of NEON AOP data, field data and training ITC set for all three tasks	September 1st, 2017
Registration Deadline	September 8th, 2017
First submission Deadline for all three tasks	October 30th, 2017
Release of first round of results and feedback	November 30th, 2017
Second submission deadline for all three tasks	December 15th, 2017
Release of second round of results and feedback	January 14th, 2018
Publication of Reports from the Participants and Evaluation	March 2018

8. Rules

- Participants may work independently or in teams of any size
- Teams should be determined at the beginning of the competition and are not allowed to merge
- Code must be Free Open Source Software (FOSS) i.e. No Tools
- Only two submissions are allowed for each team for each task
- Decisions of the UF/NIST organizers are final