TETRA TECH

Enabling Crop Analytics At Scale

ENABLING CROP ANALYTICS AT SCALE (ECAAS)

Creating Open Agricultural Maps and Ground Truth Data to Better Deliver Farm Extension Services

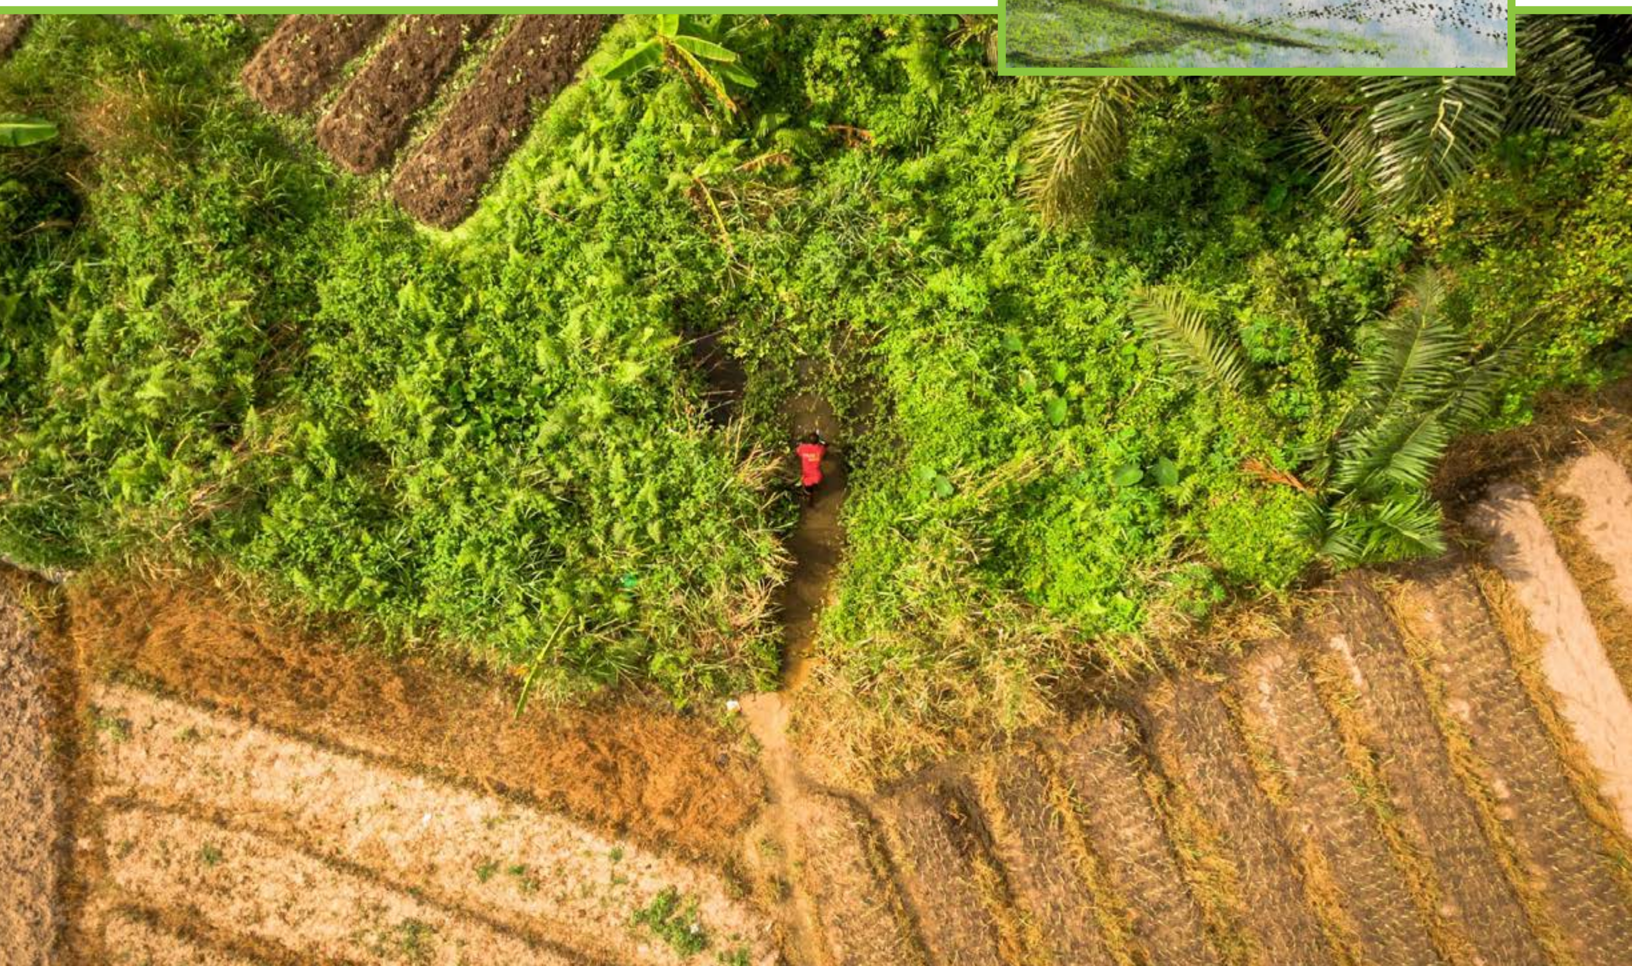

PREPARED BY

CLARK UNIVERSITY

IN COLLABORATION WITH

Farmerline

APRIL 2021

Contents

1

Overview

Expanding and improving the quality of agricultural extension services available to smallholder farmers is an important part of the effort to improve their productivity and resilience to weather extremes and other shocks. Delivering such services effectively requires accurate data on where farmers' fields are located, what crops they are growing, and what their yields are. These data are generally lacking in many smallholder-dominated agricultural systems and therefore depend on satellite remote sensing to generate. However, remote sensing of smallholders' fields is a major challenge, particularly with respect to mapping crop types. Without accurate crop type maps, it is hard to map yields-which requires linking crop-specific models to fields-and thus to see how they vary between fields and in response to different management practices (e.g. Jain et al, 2019; Jin et al, 2019). Crop type maps are also important for other applications, such as for improving national planted area estimates.

There have been several recent advances in the ability to map crop types. One important development is the increasing number of satellites that can capture imagery that is both frequent and detailed enough to track seasonal changes in crop growth within the boundaries of small crop fields (<1 ha). Chief among these are the two Sentinel missions (European Space Agency), which provide free weekly to bi-weekly optical and radar imagery at <20 m resolution. Planet, a commercial satellite company, offers daily 3.7 m resolution imagery and is implementing a fusion product (Houborg and McCabe, 2018) that produces daily, Landsat-quality reflectance data at <5 m resolution. At the same time, new machine learning approaches are rapidly improving the ability to accurately distinguish between different crop types in these different types of imagery. Models such as Random Forests, which remains a go-to model for classifying crop types (e.g. Azzari et al, 2021), are being out-performed by deep learning models (e.g. Rustowicz et al, 2020). There are a variety of architectures being used, which vary in their ability to learn from both the temporal and spatial information provided by satellite image time series (e.g. Rustowicz et al, 2020; Rußüwurm and Kß∂rner; 2017, 2018), and in how they are transferred between mapping tasks and domains (e.g. Tseng et al, 2021).

Despite these advances, the major bottleneck to their implementation is the inescapable need for georeferenced observations of crop types. Unlike field boundaries (e.g. Lesiv et al, 2019; Estes et al, 2021), most crop types cannot be recognized visually in high resolution satellite data, and thus must be collected on the ground. Collecting crop types on the ground is expensive and hard to sustain, particularly across large areas and over recurring seasons, which is the scale of coverage needed to develop reliable and repeatable crop type maps.

For this project, we developed an approach designed to address the challenge of sustainably collecting crop-type data over large areas. This approach embeds crop type collection within an existing farm extension service that is connected to a large number of smallholder farmers. The observations, collected in accordance with recommended best practices, are used to create crop type maps using machine learning model applied to Sentinel and PlanetScope imagery, with predictions filtered a high-resolution cropland map. The aim is to use the resulting maps to develop additional extension services, thereby boosting revenues and providing a means for expanding and sustaining the collection groundtruth data.

This report provides an overview of the methods and results from the initial crop type maps for maize and rice developed from groundtruth data collected during the first season of this project in the Sekyere West and Ejura Sekyedumase districts (red in Figure 1.1A).

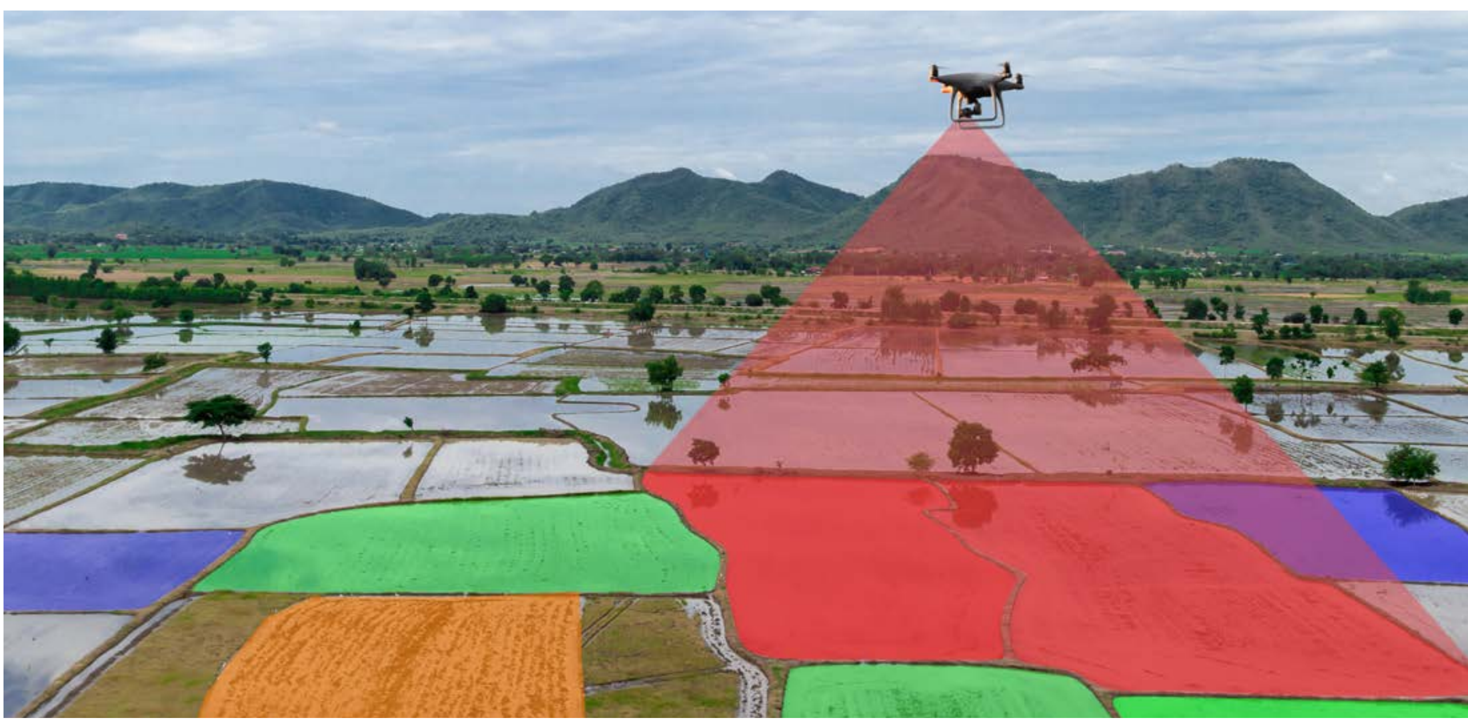

2

Mapping Approach

The mapping process had four key components (Figure 2.1), which included the processing of the sample data used to define the location and type of crop types, the pre-processing of the satellite imagery needed to predict those crop types, the training and validation of the model developed to map crop types, and the application of the trained model to map crop types onto the processed satellite imagery.

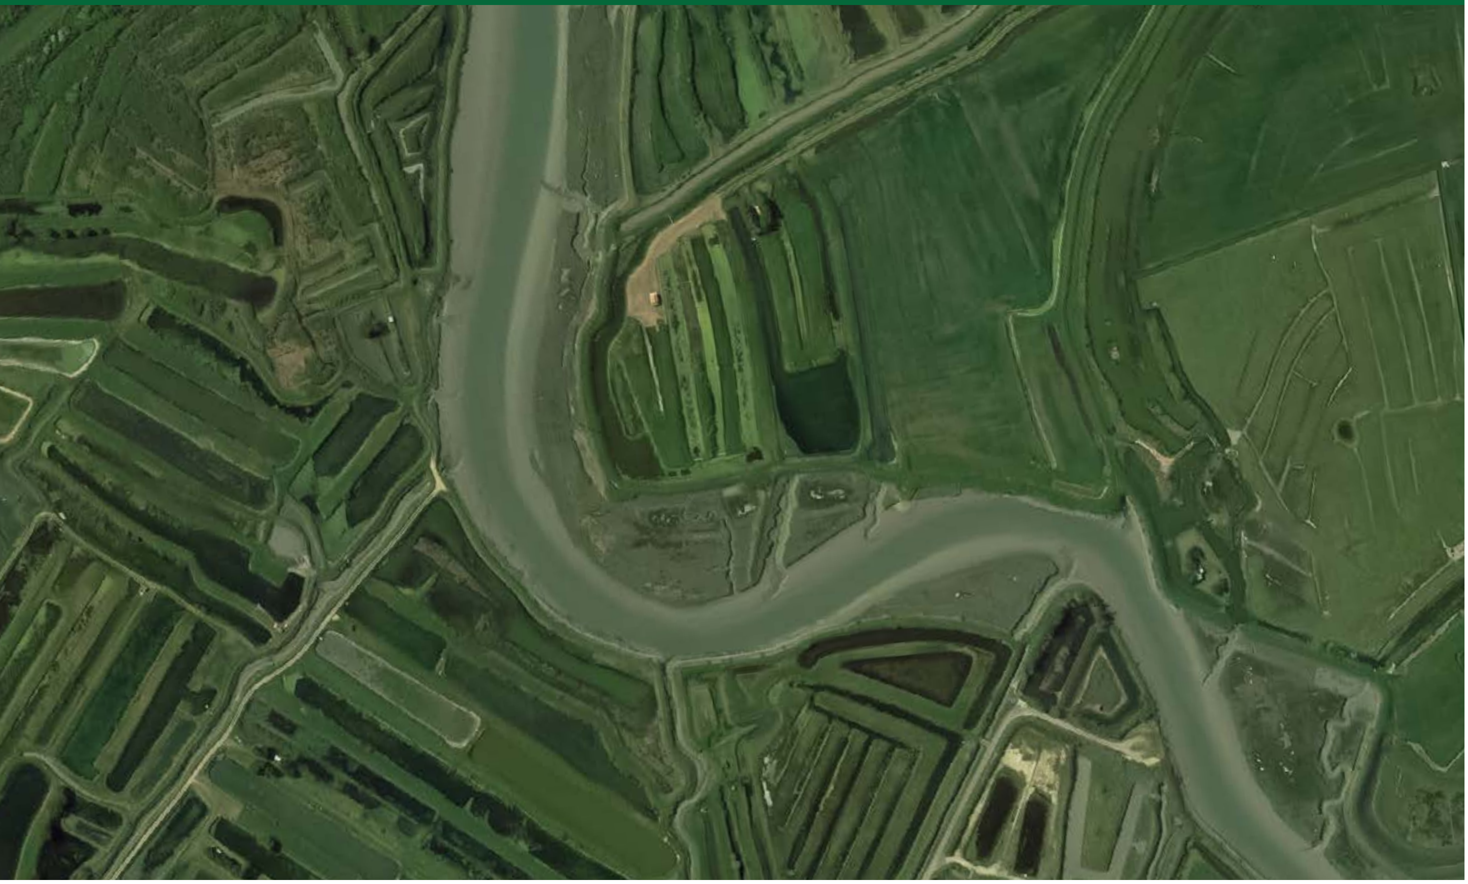

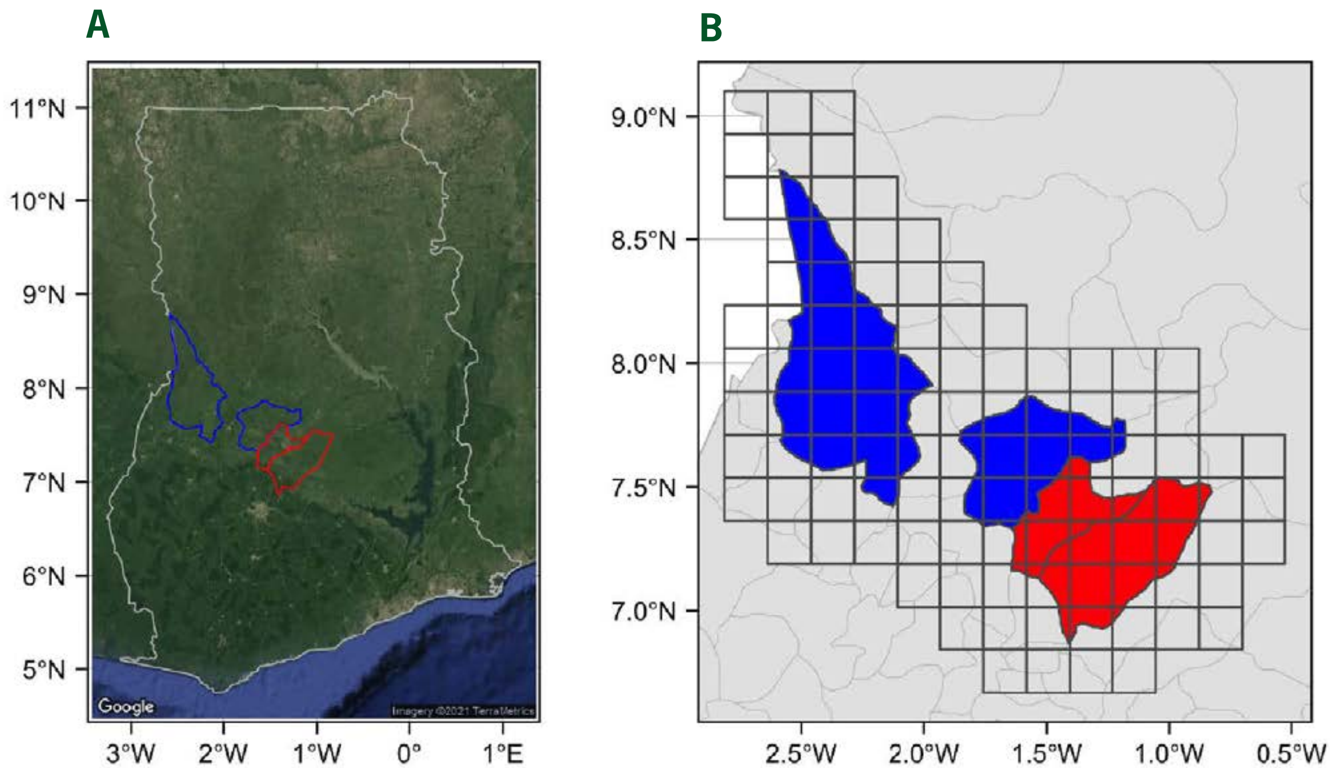

Figure 1.1:
Location of the four study districts (red = districts collected in year 1; blue = districted to be collected in year 2) in relation to the rest of Ghana (A), and an overlay of the image tiling grid used to process Sentinel-1 and 2 imagery (B).

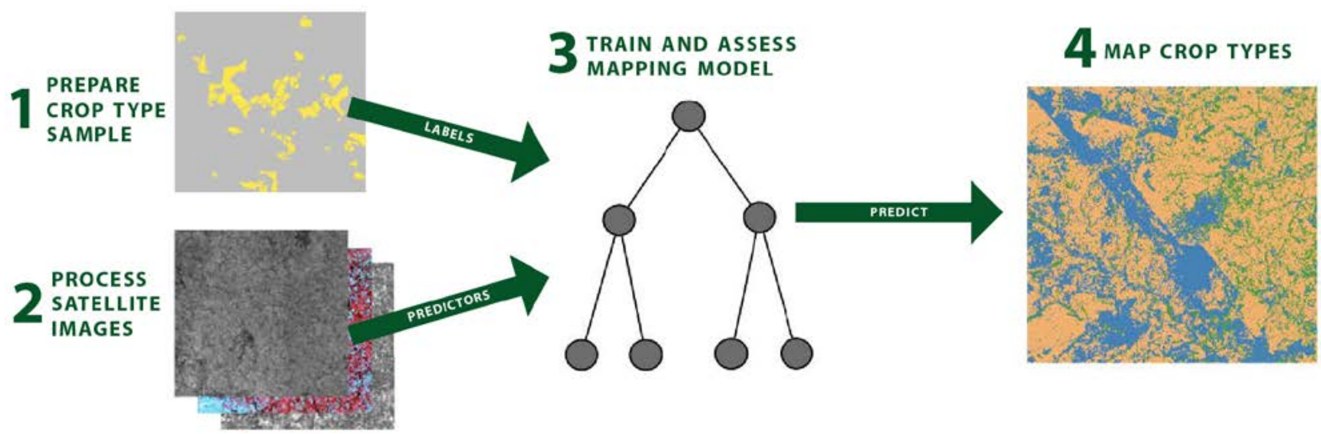

Figure 2.1:
The key components of the mapping approach.

2.1 Preparing the crop type labels

To prepare a set of data for training and validating a crop type model (labels), we followed three steps (Figure 2.2). We first cleaned and verified the crop-type polygons collected in the field (the groundtruth), which included maize, rice, and a variety of other crops grouped into a broad "other crops" class. We then created a set of sample points that indicated areas that are not cropland (non-cropland), and finally evaluated the characteristics of the sample, in terms of its spatial and temporal distribution, and the relative frequency and abundance of each class. These characteristics determine how representative the sample is of the broader region, and how effectively it can be used to train and validate mapping models.

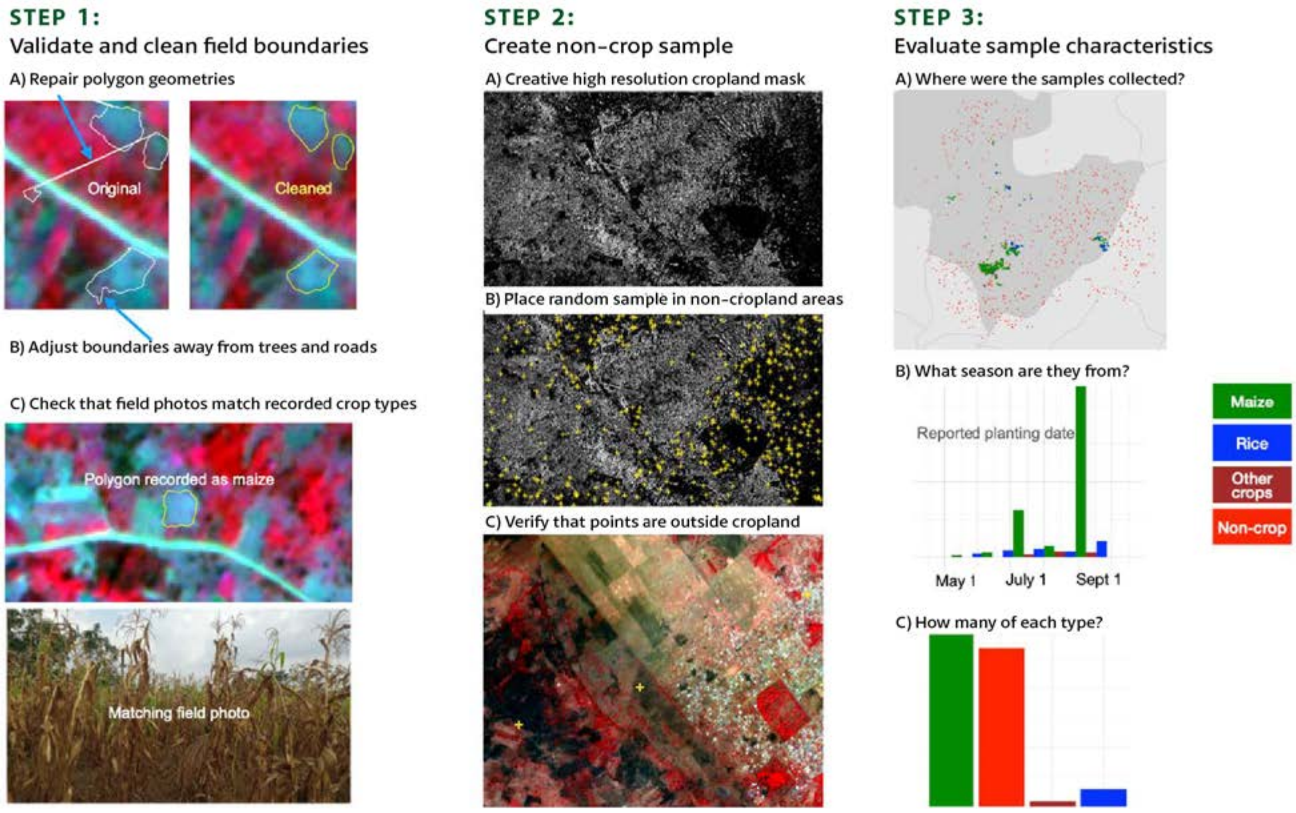

Figure 2.2:
An overview of the steps taken to prepare the crop-type sample.

The first step was performed to ensure that the field-collected boundaries did not have any invalid geometries (stray points, self-intersections, and overlaps) or misclassifications that would introduce errors into the machine learning process, or confound the ability to assess map accuracy. We repaired obviously damaged polygons where possible, or else removed them, and compared the boundaries of the polygons to field boundaries visible in PlanetScope base map imagery collected during the same season (see section 2.2). Misaligned boundaries were adjusted to avoid overlaps with adjacent field polygons, and to avoid adjacent stands of trees, roads, and other non-crop features, in order to minimize the contamination of the signal related to that crop type within the imagery. We also checked the ground-collected photos captured for each field, to ensure that the crop recorded for each polygon was correct. Table 2.1 describes several common issues encountered during this process.

Table 2.1:

A list of commonly occurring errors in the groundtruth polygons and their likely causes.

Category	Error	Cause
Geometric	› Spikes and indents › Partial polygons › Overlapping polygons › Polygon contains non-crop › cover (e.g. road) › Field photos pointed at ground or not at crop	› GPS mis-calibration, Partial polygons weak signal › GPS positional error; agent didn't follow the field edge › Photo capture didn't follow protocol
Data Entries	› Crop in photos doesn't match recorded type › Same photos for multiple fields	› Mis-coded entry; record capture or merge error › Data entry error; merge error

In the second step, we created a set of non-crop samples to help distinguish cropped from non-cropped areas. We used field boundary maps, developed for the year 2018 using our high-resolution mapping platform, to identify non-cropped areas, and placed a random sample of points within these areas. Two observers then examined each of these points (converted to ~0.1 ha polygons) within the current season's PlanetScope imagery to verify that they did not fall in cropland. The verified points were combined with the cleaned groundtruth polygons to create a full sample, which included 589 maize fields, 58 rice fields, 18 fields containing other crops (e.g carrots or cabbage), and 543 non-cropland samples located within ~15 km of the groundtruth data. The groundtruth polygons were grouped in three primary concentrations in the two districts (Figure 1.1). The reported planting dates for most fields were August-September, 2020. The appendix contains further details on the crop-type samples and their processing.

2.2 Pre-processing of satellite imagery

We used three different sources of satellite imagery known to be effective for agricultural mapping: PlanetScope, Sentinel-1, and Sentinel-2. The key characteristics and advantages and limitations of each kind of imagery are listed in Table 2.2 on the below.

Table 2.2:
The characteristics and advantages and limitations of the different satellite image sources used.

Satellite	Sensor Type	Characteristics	Advantages	Limitations
PlanetScope	Optical	<4 m resolution; Daily-coverage; 4 bands (visual and near-infrared)	Can distinguish boundaries of most small fields; High frequency good for tracking crop growth; Can develop composites even in cloudy regions; Free use for non-commercial, sustainability-oriented purposes	Spectral precision is variable; Daily imagery requires purchase; Free use limited to monthly/6-monthly coverage over tropics.
Sentinel 1	Radar	5-20 m resolution; 12-daily coverage; Dual-polarity in C-Band	Unaffected by cloud; provides uninterrupted time series of information related to vegetation physical structure	Back-scatter speckle requires filtering, limits effective resolution, and introduces error in models
Sentinel 2	Optical	10-20 m resolution; 5-daily coverage; 10 bands (visual, near and shortwave-infrared)	Can distinguish boundaries of many fields; High frequency captures vegetation phenology; High spectral depth and precision provides important information on crop type and health	Can miss smallest fields; hard to delineate field boundaries for training; 5-day coverage may be too limited for cloudiest regions

For this study, we used PlanetScope monthly to six-monthly basemaps (available through Norway's International Climate and Forests Initiative [NICFI]) covering the period June, 2020 through January, 2021. We collected Sentinel-1 and Sentinel-2 radar images for all available time periods between January, 2020 and February, 2021. For Sentinel-1, we used the ESA's algorithms for reducing noise and removing terrain effects. For Sentinel-2, we applied algorithms to atmospherically correct and detect clouds in the image time series, and then combined images into two separate temporal composites (February-October, 2020 and November, 2020 through January, 2020) to reduce contamination due to heavy cloud cover.

After pre-processing, we fit a harmonic regression to the Sentinel-1 time series, reducing the time series to 6 coefficients that provide information on seasonality, and calculated 10 vegetation indices from the Sentinel-2 seasonal composites. The Appendix contains further details on image processing.

2.3 Creating and assessing the crop type maps

The process for creating the crop type model and the resulting maps are shown in Figure 2.3 (which combines components 3 and 4 in Figure 2.1).

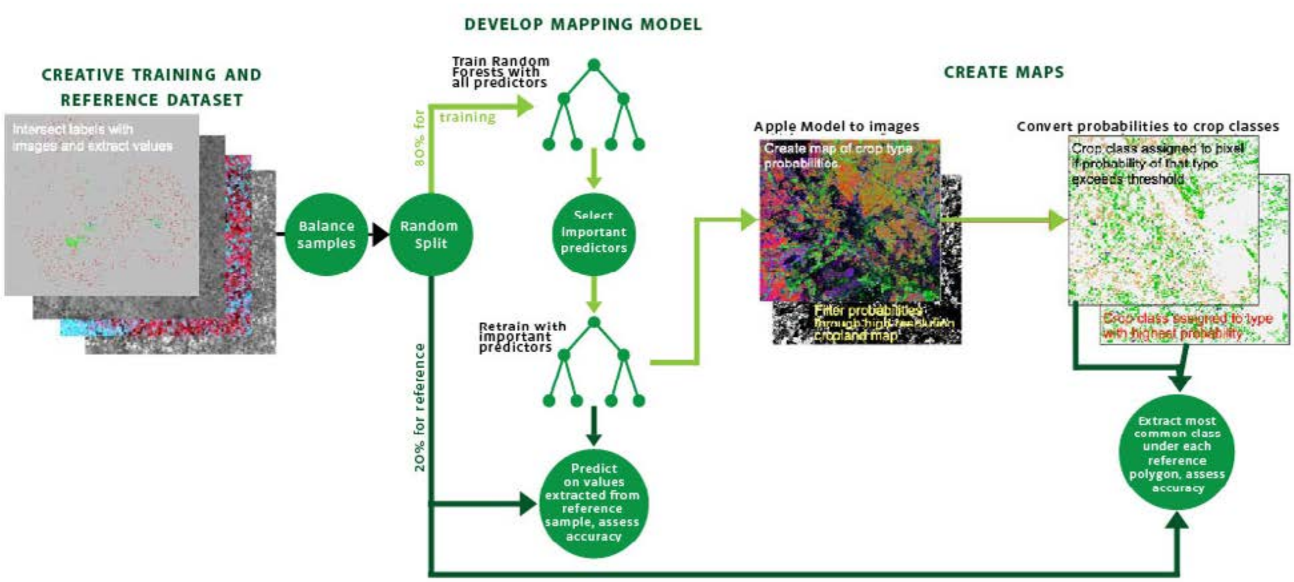

Figure 2.3:
The processing for developing and assessing the accuracy of the crop type maps.

In the first step, the crop-type samples were intersected with the processed satellite image features (80 in total), and the pixels (10 m) intersecting each polygon were extracted. We averaged the pixel values for each image feature intersecting each maize field and non-crop point but kept pixel values separate for rice and other crops because of their low numbers. We then boosted the sizes of rice and other crops by resampling and averaging subsets of their pixels within each field until their sizes matched those of the maize and non-crop samples (~550 each). We then randomly split the resulting data table, setting aside 20% of the records as a map reference data set. We used the other 80% to train a multi-class Random Forests model. We used this initial model to select the most influential image predictors of crop type and retrained the model with 53 selected features. We then applied the trained model to the full stack of selected image features to generate probability maps for each class over the study area, and filtered the probability maps through an improved, deep learning generated version of the cropland layer, to confine predictions to areas more likely to be croplands. We created two versions of the crop-type maps (Figure 2.4). The first used a thresholding approach to classify pixels as maize or rice where their respective probability maps were higher than 0.6, with all other cropland pixels defaulting to the other crops class. The second version was based on maximum probability, in which each pixel was classified to the type that had the highest predicted probability for that pixel. The first map is more conservative, as it requires higher confidence to classify pixels as maize or rice. The second map is more inclusive but its classifications have lower confidence because the maximum probability for a given pixel can be less than 0.5.

The two maps show substantially different distributions for maize, rice, and other crops. Maize and rice distributions are much sparser in the threshold map, and the other crop class is much more abundant. In comparison, the maximum probability (max-class) approach resulted in most cropland pixels being classified as maize, a larger area of rice, but relatively few cropland pixels classified as other crops.

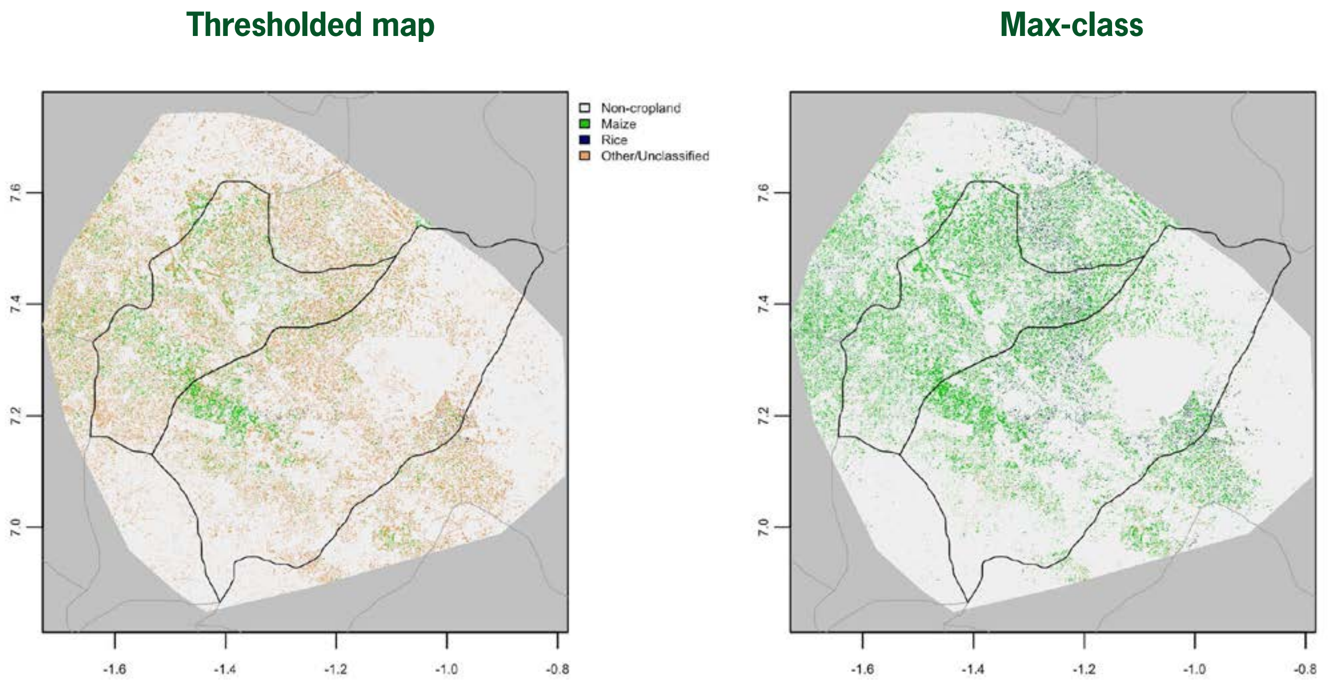

Figure 2.4:
Predicted crop type distributions based on the two mapping approaches: 1) thresholding the Random Forests predicted probabilities (left); 2) assigning the type having the maximum predicted probability per pixel (right). The predicted probabilities for both maps were first filtered using a high-resolution cropland mask. The boundaries for Ejura Sekyedumase and Sekyere West districts are shown in black.

2.3.1 Map accuracy

We calculated three estimates of model/map accuracy. First, we used the trained Random Forests model to predict crop types on the extracted and per-field averages of the image features within the map reference sample, and then cross-tabulated the predicted and observed classes. The accuracy calculated from the resulting error matrix was 96.7%, which is overly optimistic. We also calculated the accuracy of each map in distinguishing maize from non-maize areas.

To do that, we re-classed the maps to maize or non-maize, used the reference polygons for maize and noncropland (excluding rice and other crop types because resampling meant they not independent from the training sample) to extract the classified pixels from each map, and selected the dominant class in each polygon as the map class. We followed Olofsson et al., (2014), to calculate area-adjusted error matrices and accuracies (Table 2.3).

Table 2.3:

The Error matrix and accuracy measures for the maize and non-maize classes of the thresholded and maximum probability (max-class) crop type maps.

		Non-Maize	Maize	n	U
Threshholded Map	Non-maize	55.17	40.49	189	57.67
	Maize	0	4.34	38	100.00
	n	109	118		
	P	100.00	9.68		
Max-class map	Non-maize	80.57	6.02	115	93.04
	Maize	0.24	13.17	112	98.21
	n	109	118		
	P	99.70	68.61		

The overall (O), User's (U), and Producer's accuracies are provided, as are the sample size (n). The error matrix lists reference values in columns, map values in rows, and the matrix cells represent the percent of areas.

As expected, the thresholded map had the lowest overall accuracy (60%), because the high threshold resulted in most maize reference fields being missed. These omissions are reflected in the low Producer's accuracy (the complement of omission error) of 10%. Conversely, the User's accuracy for maize was 100%, meaning that no non-maize areas were misclassified as maize. We thus have higher confidence that those fields classified as maize did in fact have maize growing in it during that season. The maximum probability map's accuracy was much higher at 94%, as fewer maize fields were omitted (maize Producer's accuracy = 69%). As with the first accuracy assessment, the second map's accuracy measures almost certainly over-state its map-wide accuracy, since the groundtruth data did not conform to the design requirements for a probability sample (Stehman and Foody, 2019). The crop type polygons were geographically clustered, and are therefore unlikely to be representative of the entire mapped region. In other words, we should be less confident in a maize classification that is far away from a cluster of ground truth polygons than one that is closer to it. The most likely accuracy values for a maize/non-maize map is therefore likely to fall between the two estimates.

Both maps may be useful for different purposes. The thresholded map can be used for targeting future groundtruth collection, as it may increase the chances that sampled areas contain the crop type of interest. The maximum probability map may better represent the relative abundance of maize and rice in the mapped regions. Both maps should be considered version 1 maps, which can be greatly improved as the size and representativeness of the sample is increased (see next section for recommendations).

3

Key Findings and Recommendations

The results from this first crop type mapping exercise have highlighted several key lessons related to each component of the work, with corresponding recommendations for improvements.

3.1 Field collection of crop type data

3.1.1 The sample was unbalanced and unrepresentative

The main finding was that the small number of samples for rice and other crops relative to maize, and the clumpy geographic distribution of the sample, made it hard to develop a robust mapping model, and complicated the ability to objectively assess map accuracy. The class imbalance and lack of geographic representativeness were due in large part to logistical reasons described in the two subsequent findings.

Recommendation: Recent country-scale work found that 2,000-3,000 samples per crop class are sufficient for achieving close to the maximum possible model performance (Azzari et al, 2021). Our mapped region was <10% of the area in that study, thus the size of the maize sample in our analysis may be sufficient, but future efforts should focus on increasing samples from the other crops so that they are closer in size to that of maize. Increasing the spread of the sample across the entire mapping region is also critical. We, therefore, propose an updated sampling design that combines drone-based sampling with ongoing field efforts (see scaling-up plan in D2.5), which can increase both class balance and geographic coverage.

3.1.2 Farmer sensitization may undermine sample balance and representativeness

One of the factors contributing to the class imbalance and geographic clustering was the need to sensitize and identify farmers whose fields could be included in the sample. This preparatory work is vital and necessary for obtaining farmers' consent, but it can constrain the geographic scope of sampling and cause particular crops (maize) to be over-represented (e.g. when sensitization is done through a grower's coop).

Recommendation: This problem can be minimized by devoting more time to sensitization efforts, which can allow more farmer groups to be contacted before field efforts commence. The time required for such efforts should decline during subsequent field campaigns when previously sampled groups are revisited.

3.1.3 Transportation limited the geographic reach of sampling

Transportation hurdles, including poor roads and flooded river crossings, posed a significant impediment for field teams and limited the range of sampling efforts beyond the towns in which those teams were based. Field agents primarily use motorcycles for transport, thus traveling more than one hour away from town is impractical, or requires arranging overnight accommodation, which adds to costs and introduces other constraints (see next point).

Recommendation: Increasing the number of active field teams, selected from towns spread across the mapping region, may address this challenge (and has already been implemented for the Nkoranza and Tain districts).

3.1.4 Cell and electricity networks delay data collection and delivery

Poor cellphone coverage in certain areas (e.g. around Berem, in Sekyere West) prevented field teams from uploading their data in a timely fashion. Lack of electricity in areas where teams were sampling limited the number of fields they could collect to the length of time their batteries could last on a single charge. This was a particular problem for teams that had to stay overnight in remote areas away from their home bases.

Recommendation: The first problem cannot be readily addressed, but the latter constraint can be (and was) partially addressed by providing additional battery packs to field teams.

3.1.5 GPS devices were often miscalibrated

Geometric errors in the field polygons (Table 2.1) indicated that the GPS devices used with the Mergdata platform were prone to calibration errors, resulting in data loss and the need for substantial post-collection cleaning, which prevents the automation of mapping pipelines.

Recommendation: Calibration errors may be improved by adopting GPS software that prevents track points from self-intersecting or overlapping, or by processing geometries stored in the database with cleaning algorithms (an example is ST_MakeValid in PostGIS). Field collection protocols should also ensure that each agent verifies that collected polygons are closed and saved as soon as a field perimeter has been walked. An additional protocol suggested by NASA HARVEST (Kerner, pers. comm.) is to ensure that collection occurs only when GPS error is below 3–5 m.

3.1.6 Reference photos were often improperly registered or collected

A number of collected fields had one or more shared photos, making it impossible to verify which field they belonged to, and thereby undermining confidence in the identity of their recorded crop types. A number of photos were captured at inconsistent angles or pointed towards the ground, which limited understanding of the field context.

Recommendation: Photo capture should be more tightly integrated with the Mergdata platform. Field-collected images should retain the ExIf data containing GPS coordinates, which can help in verifying that images are correctly assigned to their recorded field polygons. Field protocols should ensure that reference images are captured consistently in a landscape mode and with a view that contains both the crop and the opposite field boundary, from each of the four major sides of the field.

3.2 Satellite image processing

3.2.1 Cloud cover is an importantlimit on the number of predictors

The high cloud cover over this area of Ghana, particularly during the May-October period, makes crop type mapping more dependent on radar-derived (Sentinel-1) predictors than in less cloudy regions. Since time series of optical imagery (particularly Sentinel-2) provide the most effective predictors (Azzari et al, 2021), the reliance on radar data will reduce the skill of mapping models.

Recommendation: Evaluate the ability of radar-optical fusion approaches to mitigate this limitation. Another approach is the CESTEM algorithm (Houborg et al., 2018) that is being operationalized by Planet, which fuses PlanetScope with other optical sensors (e.g. Landsat and MODIS) to create daily time series that may increase the frequency of cloud-free optical observations during the growing season.

3.3 Model development

3.3.1 Random Forests are less transferable

The Random Forests model we used here is less transferable than newer deep learning approaches, and may therefore require larger and more frequently updated groundtruth data to reliably map crop types.

Recommendation: Examine whether deep learning models trained with this sample and open datasets from other regions (available of Radiant MLHub), and combined with transfer or meta-learning techniques (e.g. Tseng et al, 2021), improve performance and map reliability.

4

Appendix

This appendix contains further details on methods and results.

4.1 Groundtruth data

The data processed for this study were confined to the samples of fields collected in the Ejura Sekyedumase and Sekyere West districts, representing crops planted mostly in the second half of 2020. Collections planned to capture the tail-end of the same season in the Nkoranza and Tain districts began too late, and thus coincided with the start of the subsequent growing season.

4.1.1 Field boundary validation

To evaluate field boundaries, we compared their shapes against PlanetScope analytic basemaps (4.8 m resolution) for the months of November, 2020, and January, 2021 (see Section 2.2.1), to detect misalignments between boundaries visible in the satellite imagery and GPS-collected boundaries. We removed overlapping fields, deleted "spike" vertices, and repaired self-intersections. Partial polygons caused by GPS miscalibration were edited where possible, by adjusting geometries to align with boundaries visible within the PlanetScope imagery, but we removed records where the correct location of the field could not be ascertained. In some cases, we shrank boundaries to avoid portions of fields that substantially overlapped with trees or bushed areas, to minimize contamination of the cropland signal by reflectance from non-crop vegetation. We also examined the ground-collected reference images for duplicates, and for concordance with landscape conditions visible in the satellite imagery.

4.1.2 Sample characteristics

The crop type data were collected in three primary concentrations. The largest concentration was located in the southernmost portion of Sekyere West district (near the town of Mampong), a smaller cluster in the east of this district near the town of Berem, and a more diffuse area in Ejura Sekyedumase district (Figure 4.1). We selected the portion of the non-cropland that fell within a ~15 km distance from the nearest field boundary for inclusion in the training dataset. The outermost null samples define the modeling domain used for this deliverable (Figure 4.1B).

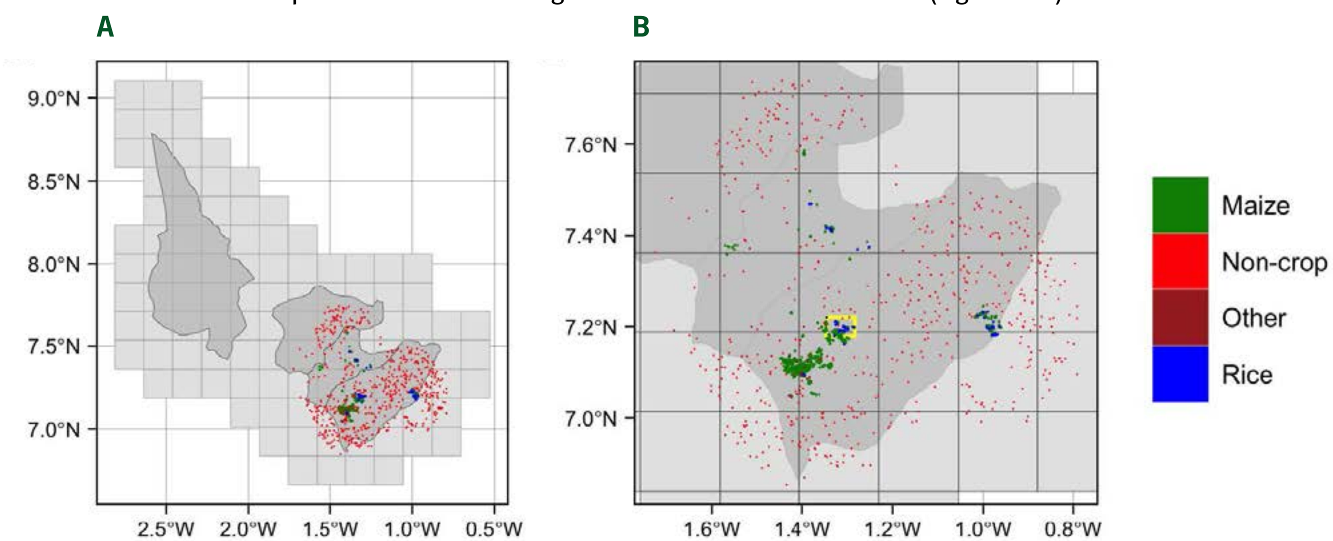

Figure 4.1:
Distribution of collected samples (Farmerline-collected field samples and image-interpreted non-cropland samples) in relation to the four districts (A), and zoomed in to the area defining the mapping bounds (B). The yellow box in (B) indicates the area mapped in Figure 2.2A.

A

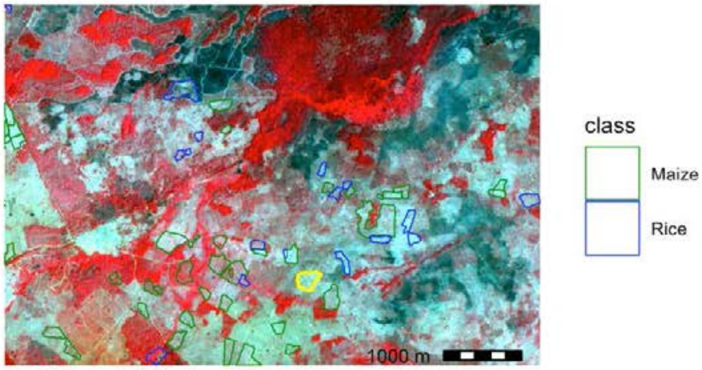

B

Figure 4.2:
Close-up of Farmerline-collected field samples and image-interpreted non-cropland samples (A; see yellow box in Figure 2B for location), and B) one of four collected reference images for a maize field indicated with the yellow outline in (A).

Table 4.1 shows the frequency and average field size of the cleaned field sample. The mean-field sizes for the polygons in each class showed that maize fields were the largest on average, although it should be noted that this is an underestimate given the boundary reductions that occurred during the cleaning process.

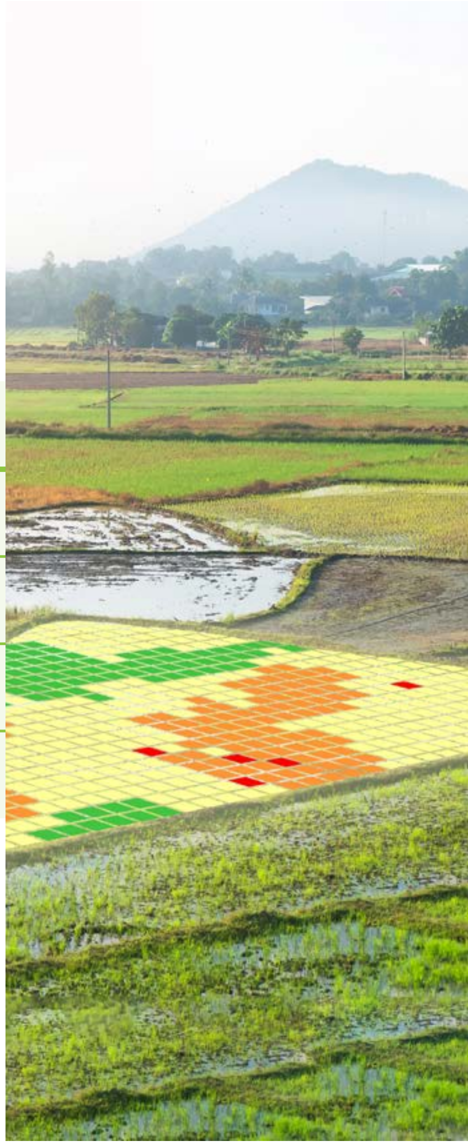

Table 4.1:
The count of the four available classes in the sample. The mean (in ha) and standard deviation (in parentheses) for the areas of polygons in each class are also provided. Class N Area Maize 589 1.84 (1.95) Rice 58 1.42 (1.49) Other 18 0.37 (0.26) Non-crop 5430.01 (0)

Class	N	Area
Maize	589	1.84 (1.95)
Rice	58	1.42 (1.49)
Other	18	0.37 (0.26)
Non-crop	543	0.01 (0)

The majority of collected fields were reported to have been planted in August-September, 2020, with a smaller number planted in June/July 4.3). A handful of records suggested planting between December 2019 and February/March, 2020, but these dates possibly reflect misreports or transcription errors.

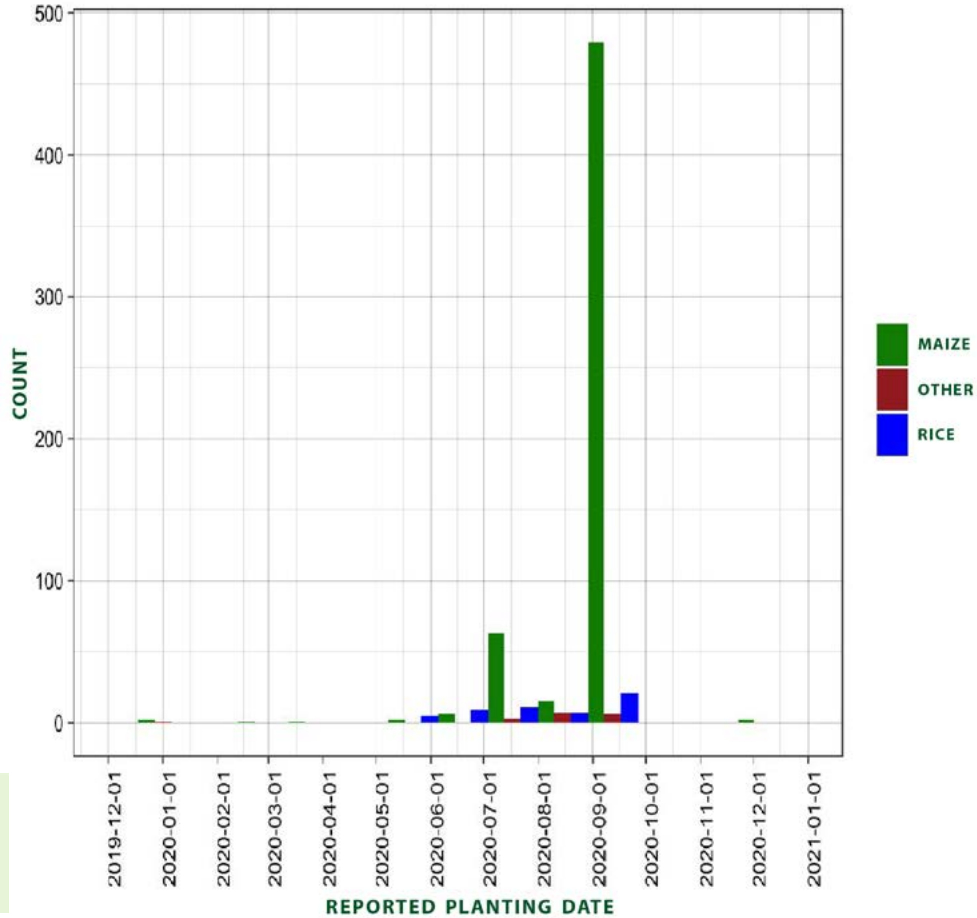

Figure 4.3:
The distribution of reported planting dates by crop type.

During the cleaning process, we also noted several commonly occurring issues within the data. The first was a relatively high frequency of topological errors, manifesting as self-intersected, overlapping, and incomplete polygons, in addition to lines connecting polygons over large distances (reported in the previous deliverable) or "spikes" protruding away from or into field boundaries. Such errors resulted from GPS miscalibration. We were able to repair a number of these errors through manual digitization or GIS post-processing, but many errors could not be fixed and the corresponding records were thus removed.

The most difficult cases were locations that had partial overlapping polygons that appeared (within satellite imagery) to fall within a single field, yet contained different data. In these cases, we used our best judgment to assign the closest matching polygons to the field in question, and removed the other records. Another issue affecting many records were cases where the same images appear in the records for two or more different fields. As the field images lack ExIf data containing GPS coordinates, the fields to which such duplicated images belonged could not be identified. This resulted in the removal of some records where it was not possible to distinguish the originating field. In other cases, images appeared to show a different landscape than what was clearly visible in the satellite imagery over the field. For example, the field images might show an open field, whereas the polygon was situated in a clearly wooded landscape. In such cases, the record was removed from the database. A number of images were also taken pointed towards the ground, which made them less useful for understanding the field's context, although they were helpful for identifying crop types.

A number of polygons did not have data entries associated with them, beyond the crop type contained in the field. We retained some of these after inspection, in order to maintain as many records of rice and other crops as possible. These records are of lower confidence, given the lack of corroborating data.

Further reconciliation for a number of records will still be required.

4.2 Imagery

4.2.1 PlanetScope

We collected PlanetScope basemaps available through Norway's International Climate and Forests Initiative (NICFI). The basemaps represented the period June-August, 2020 (with the longer period necessary for cloud cover), and monthly basemaps for October, November, and December, 2020, and January, 2021. We did not derived any indices from these images, but used their boundaries to define the tiling unit for pre-processing Sentinel-1 and 2 imagery.

4.2.2 Sentinel-1

We collected level 1 granules in dual-polarization (VV+VH) over the mapping region for the time period January, 2020-February, 2021, and undertook several pre-processing steps. We first applied the orbit file to update the satellite orbit information, then normalized the backscatter signals within the entire Sentinel-1 scene to reduce the thermal noise effects. We removed noisy and invalid data near scene edges, and converted DN values to SAR backscatter values, applying a modified Lee filter to reduce speckling, and a terrain correction algorithm to reduce topographic distortions in the images. We converted the unitless backscatter coefficient back to dB, and applied a Guided filter to further reduce speckling.

4.2.3 Sentinel-2

We collected and processed Level-1C Ttop-of-atmosphere Sentinel-2 imagery for the same time period. We applied atmospheric correction and cloud detection using the MAJA software (MACCS-ATCOR Joint Algorithm), producing a Level-2A product for each available image in the time series. We created cloud-free temporal composites using the WASP (Weighted Average Synthesis Processor), which calculates the average pixel value from an image times series using weights that include distance to clouds/shadows, aerosol optical thickness, and distance to synthesis date. The time interval used in ESA's Sen2agri platform is 45 days, but the study region in Ghana is very cloudy, therefore we used a longer multi-month interval to create the composites. The first composite period was 21 February, 2020 to October 28, 2020, and the second was 31 October, 2020 to 29 January, 2021.

4.2.4 Derived features

From the processed imagery, we derived a number of additional features to be used as predictors. We fit a "Least Absolute Shrinkage and Selection Operator" (LASSO) regression to the Sentinel-1 time series, resulting in 6 coefficients extracted from the full annual time series, which are informative about vegetation phenology. We also derived the following vegetation indices from Sentinel-2 data (following Jin et al; 2019): indices from Sentinel-2 data (following Jin et al., 2019):

Table 4.2:

Vegetation indices derived from the bands of Sentinel-2. All indices were derived for each seasonal composite.

Index	Formula
NDVI	(NIR – Red) / NIR + Red)
GCVI	(NIR/Green) – 1
RG1_GCVI	(NIR/RedEdge1) – 1
RG2_GCVI	(NIR/RedEdge2) – 1
MTCI	(NIR – RedEdge1) / (RedEdge1 – Red)
MTCI2	(RedEdge2 – RedEdge1) / (RedEdge1 – Red)
REIP	700 + 40 * ((Red + RedEdge3) / 2 – RedEdge1) / (RedEdge3 – RedEdge1)
NBR1	(NIR – SWIR1) / (NIR + SWIR1)
NBR2	(NIR – SWIR2) / (NIR + SWIR2)
NDTI	(SWIR1 – SWIR2) / (SWIR1 + SWIR2)
CRC	(SWIR1 – Green) / (SWIR1 + Green)
STI	SWIR1 / SWIR2

4.3 Development of the Random Forests model

To prepare the sample for training and testing the model, and to make up for the substantial class imbalance between the crop types, we extracted the features for each 10 m image pixel falling primarily within a given field boundary, and then averaged the feature values within each field for maize pixels, but maintained as separate the extracted feature sets underlying each pixel falling within rice or other crop types. We then increased the number of samples for these two crop types by randomly selecting pixels from each field for each type, randomly sampling (with replacement) subsets of pixels within each field and averaging them, and repeating this until the total number of samples for each crop was ~550 (close to the size of maize sample). We then randomly selected 80% of samples from each class, including the non-crop sample, and reserved the remaining 20% for testing. The resampling effort for rice and other crops meant that the 20% of these assigned to the reference sample were not independent.

The final Random Forests model had 100 trees and a tree depth of 1000. Of the 80 predictor variables, 53 were retained based on causing a mean decrease in accuracy >0.01 when removed from the model. These variables and their relative importance are shown in Figure 4.4.

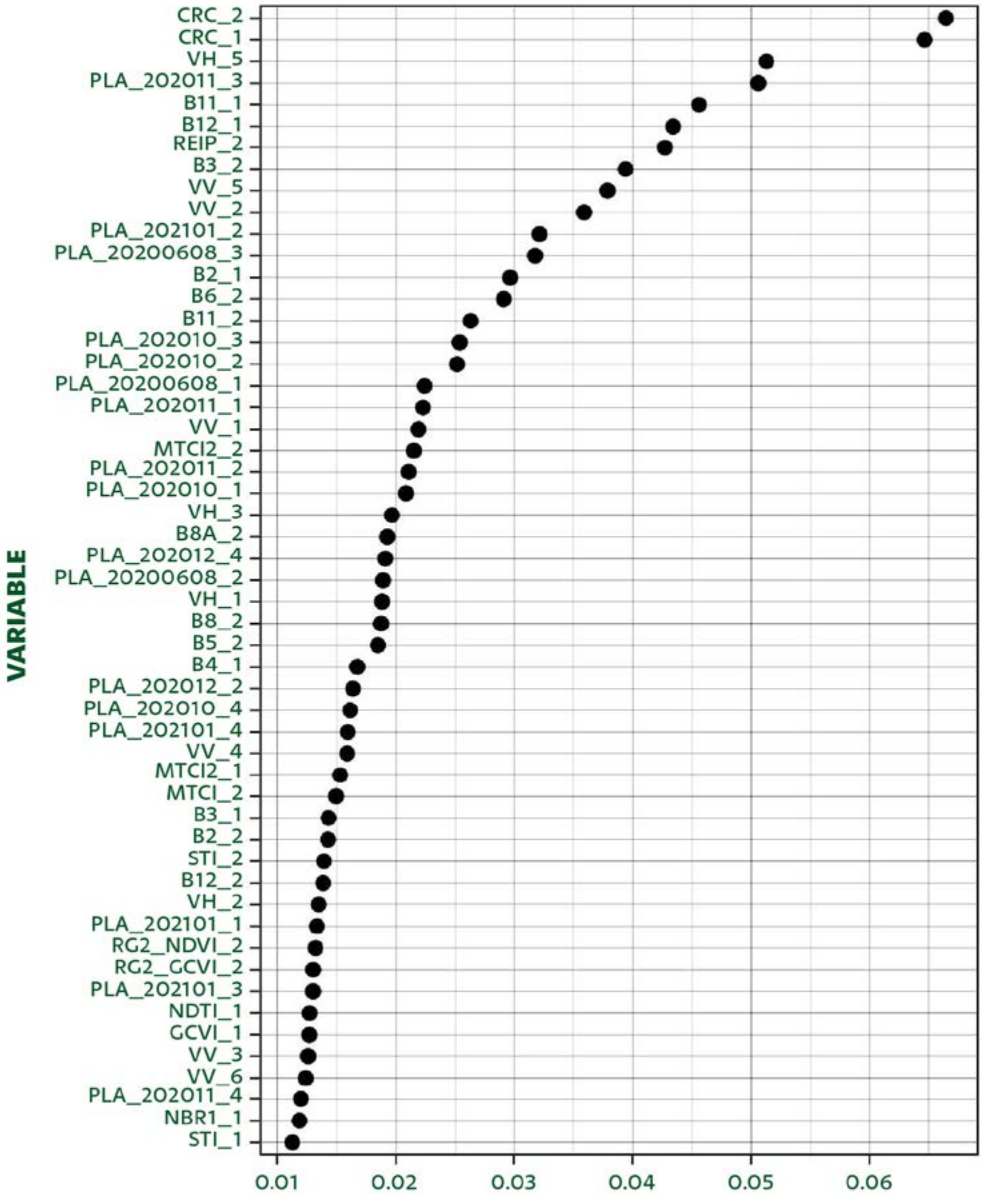

Figure 4.4:

The relative importance of predictor variables retained in the Random Forests model. Variable importance is assessed in terms of the mean decrease in model accuracy when the variable is removed. Variables containing ‚ÄòPLA' refer to PlanetScope basemaps, followed by image time period, and the band number (indicated after the underscore). Sentinel-2 bands are those beginning with ‚ÄòB', followed by a channel number and interval number (1 = February – October, 2020; 2: October, 2020 – January, 2021). Vegetation indices are described in Table 2.2. Sentinel-1 harmonic coefficients for the VV or VH polarizations are indicated by their coefficient number.

We used the trained model to generate probability maps for each crop class over the study area. We then filtered these probabilities using a second-generation cropland layer that was generated using deep-learning (U-Net) to classify cropland at 3.7 m resolution. This newer map replaces the initial version created using the original active learning approach (Estes et al, 2021). The predicted crop type probabilities were masked using this cropland layer, in order to confine predictions to more likely field boundaries. Because the map was generated from 2018 and has more omission than commission error, and because field boundaries may shift substantially between years, we merged the crop-type polygons with the cropland mask so that groundtruth samples wouldn't be masked out.

5

References

Azzari, G., Jain, S., Jeffries, G., Kilic, T. & Murray, S. (2021) Understanding the requirements for surveys to support satellite-based crop type mapping: Evidence from Sub-Saharan Africa. World Bank Policy Research Working Paper, 9609, 1-43.

Estes, L.D., Ye, S., Song, L., Luo, B., Eastman, J.R., Meng, Z., Zhang, Q., McRitchie, D., Debats, S.R., Muhando, J., Amukoa, A.H., Kaloo, B.W., Makuru, J., Mbatia, B.K., Muasa, I.M., Mucha, J., Mugami, A.M., Mugami, J.M., Muinde, F.W., Mwawaza, F.M., Ochieng, J., Oduol, C.J., Oduor, P., Wanjiku, T., Wanyoike, J.G., Avery, R. & Caylor, K. (2021) High resolution, annual maps of the characteristics of smallholder-dominated croplands at national scales.

Houborg, R. & McCabe, M. (2018) Daily Retrieval of NDVI and LAI at 3 m Resolution via the Fusion of CubeSat, Landsat, and MODIS Data. Remote Sensing, 10, 890. Jain, M. (2019) The impact of agricultural interventions can be doubled by using satellite data. 2, 5.

Jin, Z., Azzari, G., You, C., Di Tommaso, S., Aston, S., Burke, M. & Lobell, D.B. (2019) Smallholder maize area and yield mapping at national scales with Google Earth Engine. Remote Sensing of Environment, 228, 115-128.

Lesiv, M., Laso Bayas, J.C., See, L., Duerauer, M., Dahlia, D., Durando, N., Hazarika, R., Kumar Sahariah, P., Vakolyuk, M., Blyshchyk, V., Bilous, A., Perez-Hoyos, A., Gengler, S., Prestele, R., Bilous, S., Akhtar, I. ul H., Singha, K., Choudhury, S.B., Chetri, T., Malek, Ž.., Bungnamei, K., Saikia, A., Sahariah, D., Narzary, W., Danylo, O., Sturn, T., Karner, M., McCallum, I., Schepaschenko, D., Moltchanova, E., Fraisl, D., Moorthy, I. & Fritz, S. (2019) Estimating the global distribution of field size using crowdsourcing. Global Change Biology, 25, 174-186.

Olofsson, P., Foody, G.M., Herold, M., Stehman, S.V., Woodcock, C.E. & Wulder, M.A. (2014) Good practices for estimating area and assessing the accuracy of land change. Remote Sensing of Environment, 148, 42-57.

Rußüwurm, M. & Körner, M. (2017) Multi-temporal land cover classification with long short-term memory neural networks. ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLII-1/W1, 551-558.

Rußüwurm, M. & Körner, M. (2018) Multi-temporal land cover classification with sequential recurrent encoders. ISPRS International Journal of Geo-Information, 7, 129.¬†

Rustowicz, R.M., Cheong, R., Wang, L., Ermon, S., Burke, M. & Lobell, D. (2019) Semantic Segmentation of Crop Type in Africa: A Novel Dataset and Analysis of Deep Learning Methods. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, 75-82.¬†

Stehman, S.V. & Foody, G.M. (2019) Key issues in rigorous accuracy assessment of land cover products. Remote Sensing of Environment, 231, 111199. Tseng, G., Kerner, H., Nakalembe, C. & Becker-Reshef, I. (2021) Learning to predict crop type from heterogeneous sparse labels using meta-learning. To appear in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops.

Enabling Crop Analytics At Scale

Unlocking the Potential of Satellite-based Data and Analytics for Smallholder Farmers

The Enabling Satellite-based Crop Analytics at Scale (ECAAS) Initiative is a multi-phase project that aims to catalyze the development, availability, and uptake of agricultural remote-sensing data and subsequent applications in smallholder farming systems. The initiative is funded by The Bill & Melinda Gates Foundation and implemented by Tetra Tech.

info.ecaas@tetratech.com
cropanalytics.net