A MULTISCALE SPATIOTEMPORAL APPROACH FOR TRANSFERABLE IRRIGATION DETECTION

Terence Conlon*
Department of Mechanical Engineering
Columbia University
New York, NY, USA
terence.conlon@columbia.edu

Christopher Small
Lamont Doherty Earth Observatory
Columbia University
Palisades, NY, USA
csmall@columbia.edu

Vijay Modi
Department of Mechanical Engineering
Columbia University
New York, NY, USA
modi@columbia.edu

February 5, 2022

ABSTRACT

In presenting an irrigation detection methodology that leverages multiscale timeseries of satellite imagery, this paper introduces a scalable process to supplement limited ground-collected labels and ensure classifier transferability. Spatiotemporal mapping at 250m resolution is first performed to characterize dominant vegetation phenologies in an area of interest and to guide supplementary label collection. Validated dry season greening and senescence cycles observed in 10m Sentinel-2 imagery are then used to train a classifier for automated mapping of potential dry season irrigation. Strategies to improve model robustness are also demonstrated, including a method of data augmentation that randomly shifts training samples; and an assessment of classifier types that produce the best performance in withheld target regions. In applying this irrigation detection methodology to the Ethiopian highlands, results show that a transformer-based neural network architecture allows for the most robust prediction performance in withheld regions, followed closely by a CatBoost random forest model. Over withheld ground-collection survey labels, the transformer-based model achieves 96.7% accuracy over non-irrigated samples and 95.9% accuracy over irrigated samples; over a larger set of samples independently collected via the introduced method of label supplementation, non-irrigated and irrigated labels are predicted with 98.3% and 95.5% accuracy, respectively. The irrigation detection model is then deployed over Tigray and Amhara regions, revealing crop rotation patterns and areas containing changes in vegetation phenology. Predictions suggest that irrigated area in these two administrative regions has decreased by approximately 40% from 2020 to 2021.

Keywords irrigation detection · spatiotemporal modeling · multiscale imagery · transferability · Ethiopia

1 Introduction

Between 1970 and 2008, global irrigated area increased from 170 million to 304 million hectares [Vogels et al., 2019a]. In sub-Saharan Africa however, as little as 4-6% of cultivated area is irrigated, given the lack of electric grid infrastructure and the high cost of diesel [Wiggins et al., 2021]. Locating isolated irrigation identifies areas that can support higher quality energy provision services – e.g. a grid connection or minigrid installation – as these regions can sustain higher energy demands and the attendant electricity costs [Conlon et al., 2020]. Facilitated through informed

*Corresponding author

planning, irrigation expansion has a direct impact on poverty reduction: In Ethiopia, one study found that the average income of irrigating households was double that of non-irrigating households [Gebregziabher et al., 2009].

In data poor regions, satellite imagery provides a source of detailed synoptic observations of irrigated agriculture [Pervez et al., 2014]. A previous irrigation mapping effort in Ethiopia used three 1.5m SPOT6 images to distinguish between large-scale and smallholder irrigation in the Ethiopian rift [Vogels et al., 2019a]; this approach was then adapted to intake a timeseries of 10m Sentinel-2 imagery to predict irrigation presence across the horn of Africa [Vogels et al., 2019b]. While both studies demonstrated high accuracies over collected observations, limited labels precluded a more rigorous performance assessment over the entire study area of interest. Other studies have used multiscale imagery to detect irrigation, including one that fuses MODIS and Landsat imagery to identify irrigated extent, frequency, and timing in northwestern China [Chen et al., 2018]. Here, unique advantages of satellite imagery products at different resolutions are exploited: 250m MODIS imagery is valuable for characterizing vegetation over large areas [Huete et al., 1999], while decamemeter resolution imagery from Landsat or Sentinel-2 missions can better discern plot extent [Phiri et al., 2019].

Deep learning techniques have become widely used for land process classification, as they uncover intricate structures in large, complex datasets; and provide a robust method of handling phenological variability [Lecun et al., 2015, Zhong et al., 2019]. However, despite increasing availability of remotely sensed imagery, computing resources, and advanced algorithms for information extraction, high-quality labels remain scarce and expensive to acquire. Methods of overcoming label scarcity generally fall into one of four categories: 1) using pretrainined networks; 2) unsupervised and self-supervised learning; 3) data augmentation; or 4) additional label collection [Li et al., 2018]. Even as pretrained networks like ImageNet [Jia Deng et al., 2009] are highly effective for true-color image classification, these networks' weights do not translate to tasks that intake multispectral or hyperspectral imagery [Tao et al., 2020]. Unsupervised learning techniques, including those that ensemble different clustering methods – e.g. Banerjee et al. [2015] – have been shown to effectively organize unlabeled imagery. Existing work has also demonstrated that training a Generative Adversarial Network (GAN) – itself a type of unsupervised learning – has allowed for improved change detection performance on multispectral imagery, e.g. Saha et al. [2019]. For data augmentation, three techniques are often leveraged: image translation, rotation, and flipping [Yu et al., 2017, Stivaktakis et al., 2019]; however, these techniques do not have obvious analogues for pixel-based classification. Lastly, methods of collecting labels have included thresholding pixels based on normalized difference in vegetation index (NDVI) [Bazzi et al., 2021] and inspecting a single layer of high-resolution imagery [Vogels et al., 2019b].

Another lingering issue in land process mapping is that methods are often site-specific and not easily transferable to other regions or climes [Ozdogan et al., 2010, Bazzi et al., 2020]. Since labeled samples that do exist often belong to one of a few highly-studied datasets covering the US or Europe, non-transferability of methods and models means that regions like sub-Saharan Africa receive limited attention in the deep learning literature. To increase model transferability, researchers are exploring transfer learning, defined as the process of adapting models trained in an existing scene to a learning problem in a new scene [Pan and Yang, 2010]. A review of transfer learning approaches in remote sensing finding that 1) transfer learning is a recommended best-practice for novel classification tasks, and 2) models trained on larger, more generic datasets outperform those trained directly on smaller datasets [de Lima and Marfurt, 2020].

Accordingly, there is a need for deep learning methods that use limited labels to train transferable models. As current methods primarily focus on already well-understood areas of interest with existing datasets, new techniques and products developed for parts of the world lacking labeled data can provide new insights of significant value. In the realm of irrigation detection, new methodologies and mapping products can help identify regions for further energy system planning and investment, as these areas contain latent energy demands that can make higher quality energy services cost-effective and increase incomes. To this end, the following manuscript presents a multiscale methodology for robust irrigation mapping that is applied to the Ethiopian highlands, introducing a novel method of label collection, an evaluation of training strategies that ensure model transferability to unseen areas, and an assessment of irrigated area in the Tigray and Amhara regions of Ethiopia for 2020 and 2021.

2 Background

Identification of dry season greening as potentially irrigated agriculture must take into account spatiotemporal variations in native vegetation phenological cycles. The complex topography of the Ethiopian highlands and East African rift system, combined with the latitudinal movement of the InterTropical Convergence Zone (ITCZ) and seasonal upwelling of the Somali current in the Arabian Sea produces a diversity of rainfall patterns that control annual vegetation phenological cycles in the study area. In order to provide a phenological context with which to identify anomalous dry season greening, a regional vegetation phenology map is derived from spatiotemporal analysis of timeseries of vegetation abundance maps. Using the spatiotemporal characterization and temporal mixture modeling approach given

by [Small, 2012] applied to timeseries of MODIS enhanced vegetation index (EVI) maps, four temporal endmember (tEM) phenologies are identified that bound the temporal feature space of all vegetation phenology cycles observed on the East African Sahel. These four tEM phenologies form the basis of a linear temporal mixture model that can be inverted to provide tEM fraction estimates for each pixel's vegetation phenology in an imagery cube. Figure 1 presents a spatiotemporal phenological characterization for the country, created by applying a temporal mixture model to an image cube containing 16-day 250m MODIS EVI imagery between June 1st, 2011 and June 1st, 2021.

The four tEMs extracted for Ethiopia are as follows: a *single cycle* tEM, representing a single annual vegetation cycle per year that peaks in September/October; an *evergreen* tEM, representing perennial vegetation; a *double monsoon* tEM, representing semiannual vegetation cycles cycles observed on the Somali peninsula; and a *non-vegetated* tEM, representing barren or non-existent vegetation. The ensuing phenology map in Figure 1 contains unmixing root mean square (RMS) error less than 10% for 90% of the pixels; additional unmixing error statistics and the locations of the extracted tEMs in PC feature space are shown in Supplementary Figures S1-S2.

Figure 1 roughly divides into 4 quadrants. In the northeast quadrant, the Afar region appears as dark green, indicating that none of the 4 tEMs contribute significantly to phenologies in this part of the country: The vegetation that does exist in this mostly barren region is represented by low evergreen tEM abundances. In the southeast quadrant, dominated by the Somali region and a portion of Oromia, vegetation patterns cycle with the double monsoon. However, it is worth noting that the *double monsoon* tEM contains peak vegetation abundances lower than those of the *single cycle* and *evergreen* tEMs, denoting the existence of sparser vegetation; it follows that these regions are more pastoral and not as agriculturally cultivated as the other parts of the country.

The southwest quadrant – covering Southern Nations, Nationalities, and Peoples (SNNP) Region, Sidama, and the western portion of Oromia – contains significant amounts of evergreen vegetation, as is demonstrated by its bright green hue. In contrast, the northwest quadrant of the phenology map contains red-dominant color gradients, indicating phenologies similar to the *single cycle* tEM. This portion of the country, known as the Western Ethiopian Highlands and comprising of the regions Amhara, and Tigray, is highly agricultural; the main cropping season lasts from June to October and coincides with the primary *kiremt* rains, with some secondary cropping following the lighter *belg* rains from March to May.

In presenting a map of dominant vegetation phenologies in Ethiopia, Figure 1 provides a guide for land cover classification transferability within the country. For instance, a dry season crop detector trained in Amhara will perform poorly in SNNP, as phenological patterns differ significantly across these regions, and dry season crop cycles exhibit different vegetation signatures. In contrast, a dry season crop detector developed across Amhara can be transferred to Tigray or Benishangul-Gamuz, due to the phenological similarities between these regions.

The named, italicized polygons in Figure 1 represent the 8 regions containing labeled samples used in this manuscript: The yellow polygon indicates an area where labeled data were collected via a ground survey, and the purple polygons indicate areas where labels were collected by means of visual interpretation and timeseries inspection. Full information on the labeled data collection process is presented in Section 3.

3 Materials and Methods

The data collection portion of this manuscript's methodology consists of pairing Sentinel-2 imagery with labeled polygons to train an irrigation detector. Irrigation is defined as such: A pixel is irrigated if its phenology includes at least one non-perennial vegetation cycle during the dry season, December 1st to April 1st for the West Ethiopian Highlands. Conversely, a pixel is non-irrigated if its phenology demonstrates only vegetation growth that can be attributed to the region's known rainy seasons. Irrigated areas are only of interest if they contain multiple annual vegetation cycles; this strict definition of irrigation excludes supplemental irrigation practices and perennial crops that may be consistently irrigated throughout the year.

3.1 Sentinel-2 imagery collection

The following analysis uses bottom-of-atmosphere corrected (processing level L2A) Sentinel-2 temporal stacks – four dimensional arrays created by stacking a set spatial extent of imagery bands over multiple timesteps – using the Descartes Labs (DL) platform, a commercial environment for planet-scale geospatial analysis. Images are collected at a 10-day time resolution. To focus on the 2020 and 2021 dry seasons, the time period of interest is defined as between June 1st, 2019, and June 1st, 2021. Given the 10-day timestep, 72 image mosaics are collected – 36 per year. Additional information on the imagery download process is available in the Supplementary Materials.

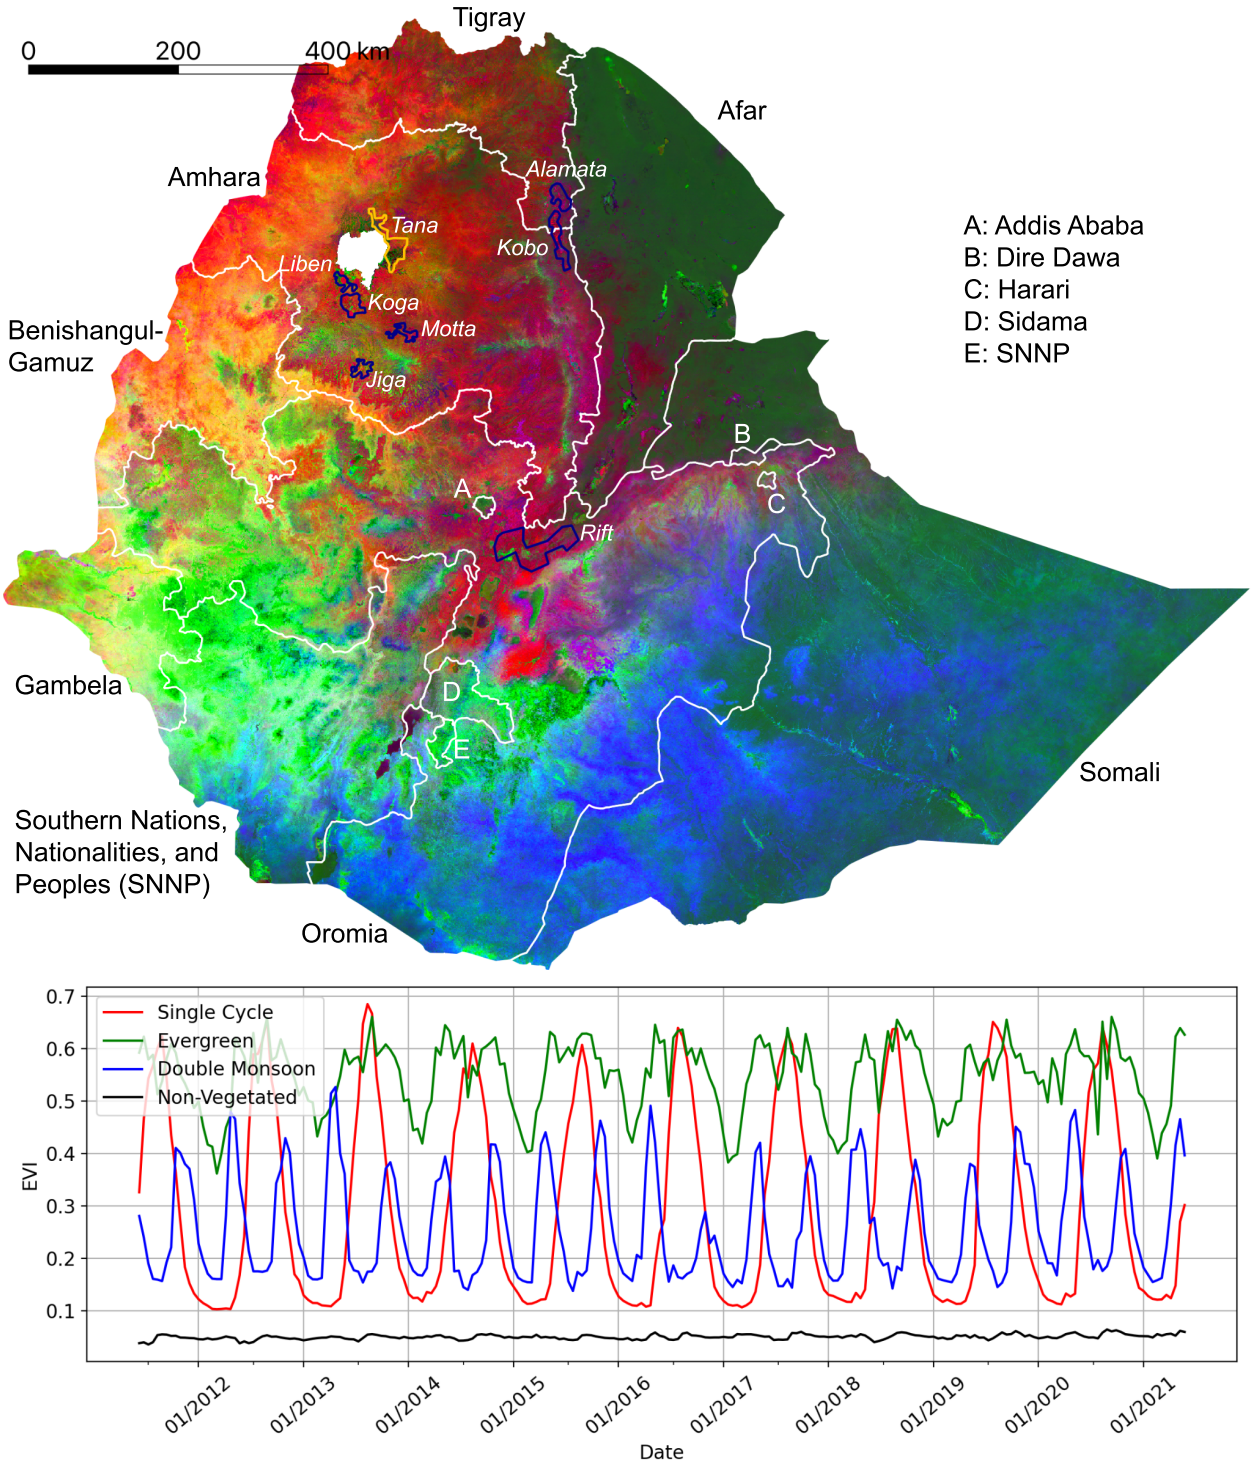

Figure 1: Continuous temporal endmember fraction map derived from a temporal mixture model of 250m MODIS enhanced vegetation indices (EVI). Continuous gradients and abrupt transitions in phenology are primarily related to topography and variations in precipitation. Region names showing locations of labeled polygons are italicized: The region containing ground collection (GC) labels is delineated in gold; the regions containing visual collection (VC) labels are delineated in blue.

3.2 Label collection

Two types of labeled data are leveraged for irrigation mapping: *ground collection* (GC) labels, acquired via an in-person survey; and *visual collection* (VC) labels, acquired via visual identification of dry season vegetation from Sentinel-2 imagery using the DL platform and subsequent cleaning via timeseries clustering. The locations of these GC and VC regions are shown in italics in Figure 1, with all labels collected for the 2021 dry season. A description of the ground collection survey is presented in the Supplementary Materials. As the GC labels constitute our highest quality irrigation observations, verified by in-situ visits to individual plots, we do not use them for training during the model sensitivity analysis, instead reserving them for validation of classifier performance.

3.2.1 Visual label collection

To supplement the GC labels located in Tana, visually collected labels are acquired for seven separate regions via a three-step process of 1) visual inspection, 2) EVI timeseries confirmation, and 3) cluster cleaning; each of these steps are described in the eponymous subsections below.

Visual inspection The first step in the VC labeling process involves drawing polygons around locations that either: a) present as cropland with visible vegetation growth (for the collection of irrigated samples), or b) present as cropland with no visible vegetation growth (for the collection of non-irrigated samples), based on dry-season, false-color Sentinel-2 imagery presented on the DL platform. Sub-meter resolution commercial satellite imagery from Google Earth Pro is also used to confirm the existence of cropland in the viewing window. For the collection of non-irrigated samples, polygons are restricted to areas that contain non-perennial cropland; however, because only phenologies that contain dry season vegetation cycles are considered irrigated, non-irrigated polygons occasionally overlap other types of land cover – e.g., perennial crops, fallow cropland, or areas with human settlement – with any overlap likely to improve training robustness.

EVI timeseries confirmation After drawing a polygon around a suspected irrigated or non-irrigated area, the second step in the VC label acquisition process entails inspection of the median Sentinel-2 EVI timeseries of all pixels contained within the polygon; this step is shown in the plot windows of Figure 2. Here, all available Sentinel-2 imagery with less than 20% cloud cover between June 1, 2020, and June 1, 2021 is retrieved; a cubic spline is then fit to all available samples to generate continuous EVI timeseries. For potential irrigated polygons, if the EVI timeseries shows a clear peak above 0.2 during the dry season, it is confirmed as irrigated. Similarly, for potential non-irrigated polygons, an EVI timeseries that demonstrates a single vegetation cycle attributable to Ethiopia's June to September rains is taken as confirmation of a non-irrigated VC polygon. However, if the EVI timeseries does not confirm the expected irrigated/non-irrigated class, or if the plotted EVI error bars (representing $\pm$one standard deviation of the EVI values at that timestep) indicate a level of signal noise within the polygon that prevents the identification of a clear vegetation phenology, the polygon is discarded.

Figure 2(a) demonstrates an example of irrigated VC label collection in the Koga region – here, the double vegetation peak present in the EVI timeseries confirms the purple polygon in the center of the window as irrigated (blue polygons indicate areas already saved as irrigated VC labels). Figure 2(b) demonstrates the same process for non-irrigated VC labels, also in Koga: The single EVI peak in October 2020 confirms the pink polygon in the top left of the window as non-irrigated (red polygons indicate areas already saved as non-irrigated VC labels).

Cluster cleaning The third step in the VC label acquisition process involves bulk verification of the collected timeseries by means of cluster cleaning. For each VC region, all pixels that reside within labeled polygons are collected and split based on the irrigated/non-irrigated class labels of the polygons. Fifteen-component Gaussian mixture models are fit to each class's data to extract the dominant phenologies contained within the region's samples; the EVI timeseries representing the cluster centroids are then plotted, with the plot legend displaying the number of samples per cluster. Figure 3(a) presents the results of this initial clustering for the Koga region.

From the initial cluster timeseries, an iterative process begins to ensure that all cluster timeseries align with the specified class label. For an irrigated cluster timeseries to be kept, it must contain multiple successive EVI values above and below 0.2, and it must contain a clear EVI peak above 0.2 during the dry season. Analogously, non-irrigated cluster timeseries are discarded if they display a clear dry-season EVI peak above 0.2 . If these conditions are not met – as is the case for Clusters 12, 13, and 14 of the Koga irrigated samples, which do not contain a clear EVI peak above 0.2 between December 1, 2020 and April 1, 2021 and/or do not senesce below an EVI threshold of 0.2 for successive timesteps – all pixels associated with that cluster are discarded from the labeled data. This process is repeated until all 15 clusters for both classes demonstrate EVI signals that meet the non-irrigated/irrigated class definitions. The final, cleaned cluster timeseries for the Koga region are shown in Figure 3(b).

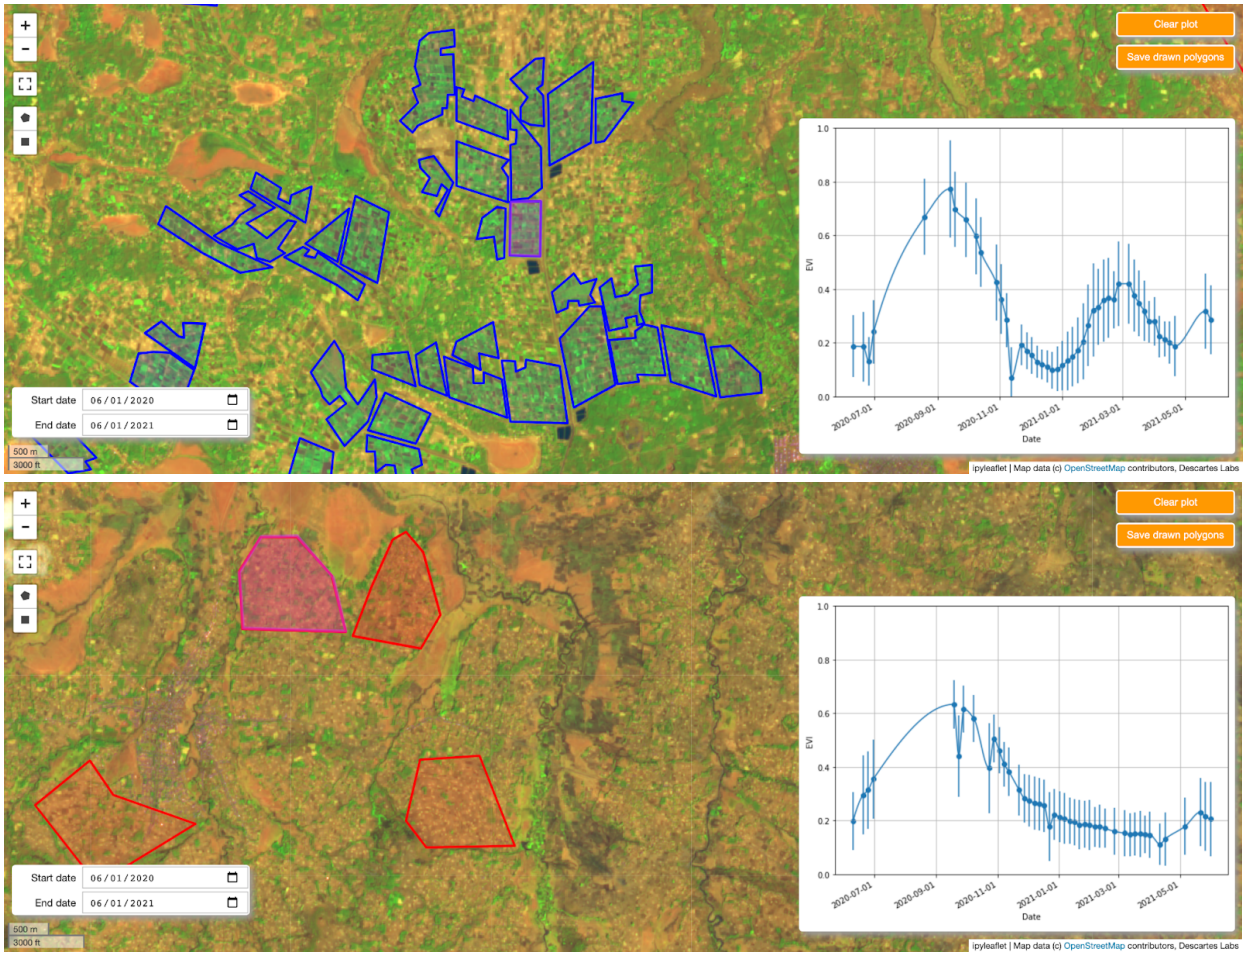

Figure 2: Example of the visual collection (VC) labeling process in Koga using the Descartes Labs platform. Blue polygons denote areas determined to be irrigated; red polygons are determined to be non-irrigated. Background imagery is a false-color Sentinel-2 image taken in March 2021: red, near-infrared, and blue bands are presented in the RGB channels, respectively. In (a), the Sentinel-2 enhanced vegetation index (EVI) timeseries is shown for the drawn purple rectangle in the middle of the window; in (b), the Sentinel-2 EVI timeseries is shown for the drawn pink, semi-octagonal polygon in the top left of the window. Both timeseries present the median EVI values for all pixels contained within the drawn polygon; the error bars show one standard deviation of these values above and below the median. In both figures, the drawn polygons are confirmed as VC labels, since they meet the definitions of irrigation/non-irrigation, respectively.

Cluster-cleaning is performed for all regions' labeled data, including labeled data collected from the GC region, Tana. For increased visibility into the labeled data collected and used for training, these regions' clusters before and after cleaning are included in Appendix A of the Supplementary Materials.

A summary of the number of collected polygons and cleaned pixel timeseries samples in each region is shown in Supplementary Tables S1-S2: In total, 1,207,233 non-irrigated samples and 907,887 irrigated samples are used, taken from 1702 and 750 labeled polygons, respectively. For model training and evaluation, data are divided among training, validation, and test splits[2]. Here, polygons in each region are split according to a 70/15/15 training/validation/test ratio; this method ensures that highly similar pixels from within the same polygon do not exist across training configurations, a division of data that would artificially inflate model performance for the task of predicting irrigation over pixel timeseries unseen by the model. All training, validation, and testing is performed pixelwise (i.e., having removed the spatial relationships of samples).

[2]In splitting the labeled data, the training/validation/testing terminology standard in machine and deep learning literature is adopted.

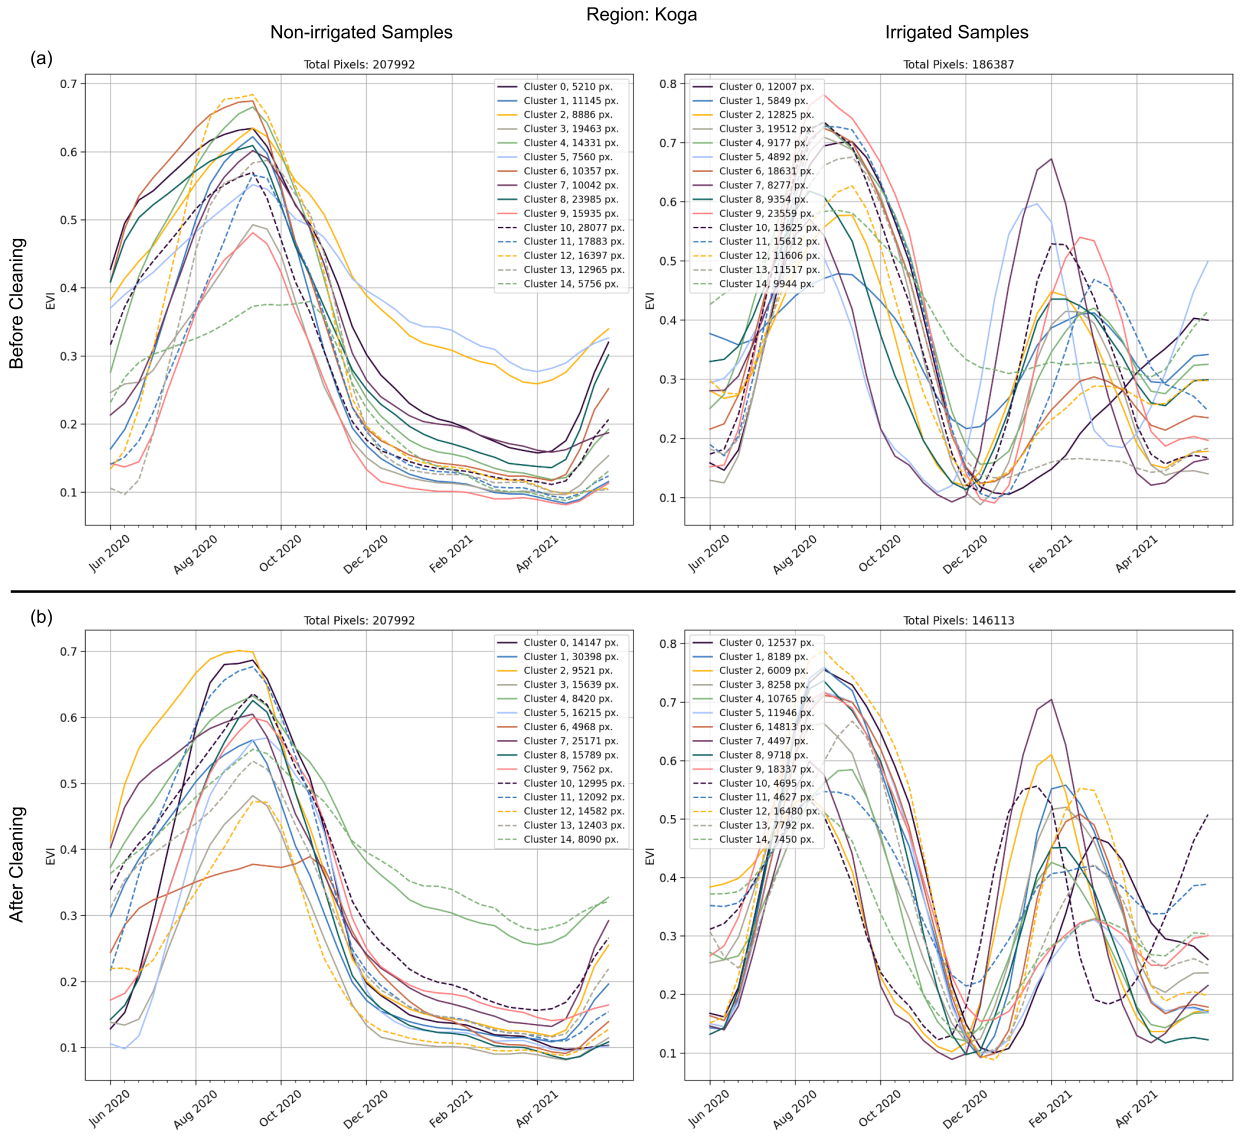

Figure 3: Clustered enhanced vegetation index (EVI) timeseries before and after cluster cleaning for the Koga visual collection (VC) region. Before and after cleaning, pixels are grouped into one of 15 randomly-indexed clusters. In (a), Clusters 12, 13, and 14 of the irrigated samples do not achieve an EVI peak of 0.2 during the dry season (December 1st to April 1st), and are discarded. All non-irrigated clusters display a single vegetation peak aligned with the main rainy season, and the irrigated clusters after cleaning (b) all display a vegetation cycle during the dry season.

The Supplementary Materials contain additional information about the labeled data distributions, including a statistical evaluation of the similarity of labeled samples across region and class (see Supplementary Tables S3-S4)

3.3 Prediction admissibility criteria

Given that irrigated phenologies exist over a small fraction of the total land area of the Western Ethiopian Highlands, and that there are many types of land cover that do not fall within this manuscript's non-irrigated/irrigated cropland dichotomy, a set of criteria are imposed to exclude pixel phenologies that are not cropland or are highly unlikely to be irrigated. Table 1 presents five criteria that must all be met for a pixel to be potentially irrigated and the motivation behind each.

Table 1: Prediction admissibility criteria. All criteria need to be satisfied for a prediction to be admitted as irrigated.

Criteria	Motivation
10^{th} percentile of EVI timeseries < 0.2	Remove evergreen samples
90^{th} percentile of EVI timeseries > 0.2	Remove barren/non-vegetated samples
Maximum of the EVI timeseries during the dry season (Dec 1 – Apr 1) > 0.2	Remove samples with no vegetation growth in the dry season
Ratio of the 90^{th}:10^{th} percentile of the EVI timeseries > 2	Remove evergreen samples
Shuttle Radar Topography Mission slope measurement $< 8\%$	Remove samples in highly sloped settings where cropping is impractical

Theses vegetation-specific criteria are informed by the EVI distributions of labeled irrigated samples for all regions: Supplementary Figure S3 contains cumulative distribution functions (CDFs) for the 10^{th} and 90^{th} EVI timeseries percentiles, the 90^{th}:10^{th} EVI timeseries percentile ratio, and the maximum EVI value during the dry season. CDFs are presented for all regions' irrigated samples, including for a set of polygons collected over evergreen land cover areas.

The criteria in Table 1 are also used to create a reference irrigation classifier that does not rely on machine learning. For this reference classifier, if all 5 conditions are met, the sample is deemed irrigated; if any of the conditions is not satisfied, the sample is deemed non-irrigated.

3.4 Model training

3.4.1 Model architectures

Five separate classifier types are compared to determine the model architecture with the most robust irrigation detection performance across regions. The first two classifiers are decision tree-based: A random forest with 1000 trees [Breiman, 2001]; and a CatBoost model that uses gradient boosting on up to 1000 trees [Prokhorenkova et al., 2019]. The other three classifiers are neural networks (NN): A baseline network, an LSTM-based network, and a transformer-based network. For comparability, these three classifier architectures are designed to have similar structures, based on the strong baseline model structure proposed in [Wang et al., 2017]; as seen in Figure 4, they differ only in the type of encoding blocks used.

3.4.2 Model training strategy

The implemented model training strategy addresses two potential pitfalls among training processes: 1) imbalanced samples across region and class; and 2) high similarity among samples within a region that may not reflect the sample distributions across all regions. Consistent with best practices in dealing with imbalanced data, this first issue is addressed with a) class balancing weights specific to each region, based on the "balanced" heuristic inspired by King and Zeng [2001]; and b) a region-specific weight equal to the ratio of the maximum number of samples in any region to the number of samples for the region in question. Both class-balancing and region-balancing weights are used in all training configurations.

To address address potential redundancy and time-specificity among samples within a region, random shifts are applied to all input timeseries. The sizes of these random shifts vary between -3 and +3 timesteps (corresponding to between -30 and +30 days), with an equal probability of each shift occurring (including a shift by 0 timesteps). Random shifts are applied to all samples in the training and validation sets and differ for each sample every time it's seen by the model. No shifts are applied to the samples in the testing sets.

The primary metric for performance evaluation is the F_1 score on the test datasets of regions withheld from training. Accordingly, performance is assessed in a manner that prioritizes classifier robustness – i.e. performance in regions unseen during training – and not in a manner that could be inflated by close similarity of samples within a region. For reference, the F_1 score balances prediction precision and recall, and is calculated per Eq. 1.

$$F_1 = \frac{TP}{TP + \frac{1}{2}(FP + FN)} \tag{1}$$

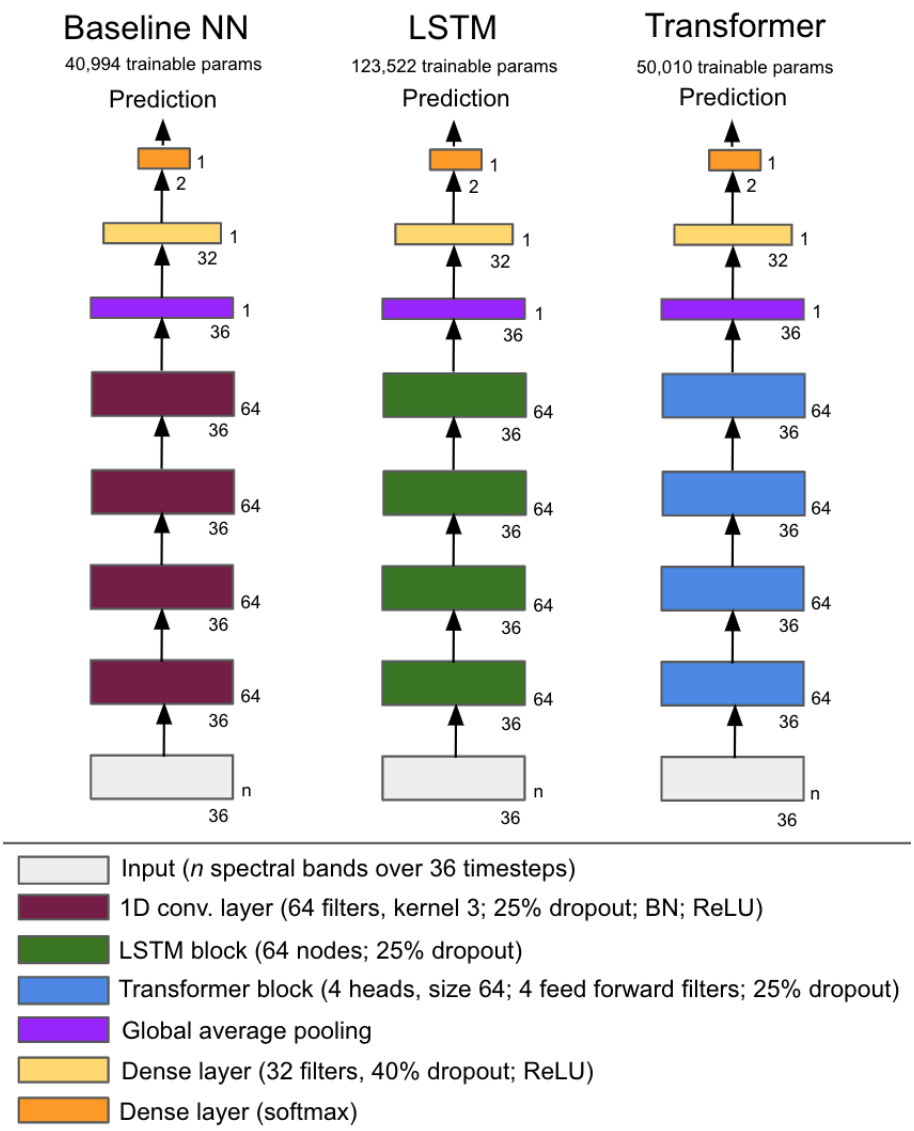

Figure 4: Neural network (NN) model architectures tested as irrigation detection classifiers. Model architectures are consistent by design; only encoding blocks differ across networks.

The training strategy differs for the tree-based classifiers and for the neural network-based classifiers. As training the tree-based classifiers occurs across a single batch with no iteration across epochs, there is no need for separate validation and testing datasets: The training and validation datasets of all included regions are therefore combined to create a single training dataset. After training on this combined dataset, performance is evaluated across the test datasets of all regions, both those included in training and those withheld.

In contrast, training neural network-based models takes place by batch across epochs, and a validation set is required to guide the training process. For a given training step, one batch from each region is concatenated, with the combined output shuffled before model intake. After the epoch is finished, performance is assessed on the validation set of each region included in training. If the minimum F_1 score among all regions' validation sets has increased from its previous maximum, the model weights are saved; however, if the minimum F_1 score has not increased from its previous high point, the model weights are discarded. A minimum F_1 score across all validation regions is selected as the weight update criteria to ensure model robustness: Consistent performance across all regions is desired, not high performance in one set of regions and poor performance in another. Training concludes once the minimum validation region F_1 score has not improved for 10 training epochs, or after 30 epochs have been completed. After training, model weights are

loaded from the epoch with the highest minimum validation region F_1 score; performance of this model on the test datasets of all regions is then reported. For all training runs, a binary cross-entropy loss, a learning rate of 1e-4, and an Adam optimizer [Kingma and Ba, 2015] are specified. Inputs are standardized to a mean of 0 and standard deviation of 1 using statistics from the entire set of labeled samples.

4 Results

4.1 Model sensitivity

Figure 5 presents withheld VC region test dataset F_1 scores for three different types of model input – one that includes all spectral bands for all timesteps; one that includes only the EVI layer for all timesteps; and one that includes only the EVI layer for all timesteps with the random sample shift applied. Here, the performance of models trained on all combinations of VC regions is evaluated; these results are organized along the *x*-axis by the number of VC regions included during training. Each *x*-axis tick label also includes in parentheses the number of withheld VC region test dataset evaluations for all models trained on *x* included VC regions. All results are presented for the transformer model architecture; however, these findings are agnostic to the classifier architecture selected.

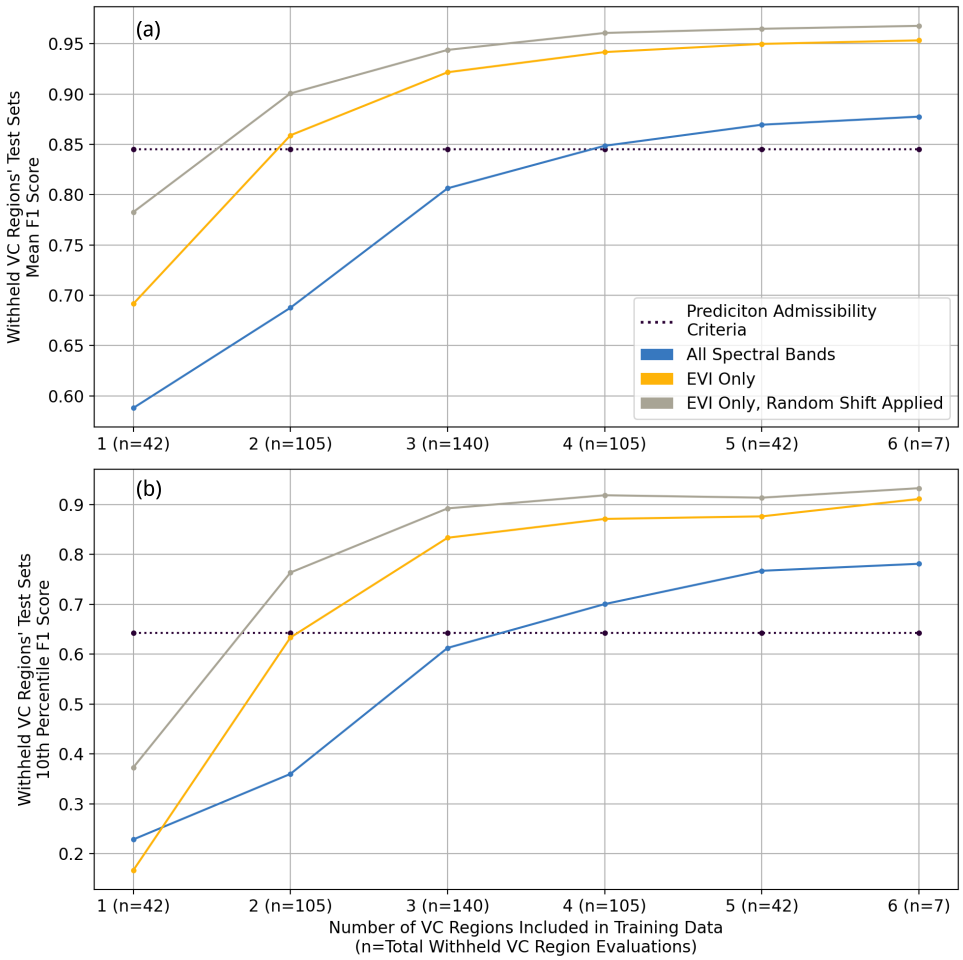

Figure 5: Withheld region test dataset performance for different types of model input, organized along the *x*-axis by the number of regions included during training. (a) presents mean F_1 score over the withheld regions; (b) presents the 10^{th} percentile F_1 score over the withheld regions. Results indicate that model inputs of randomly-shifted enhanced vegetation index (EVI) timeseries yield the best classifier performance. F_1 scores from classification based on the prediction admissibility criteria are presented for reference.

Figure 5 demonstrates that models trained on samples containing only EVI timeseries outperform those that include all spectral bands at all timesteps, both on average (a) and in low performing regions (b). The 10th percentile of withheld regions' F_1 scores is shown in order to understand the low-end of model performance without accounting for outliers. Accordingly, explicitly feeding classification models information about samples' vegetation content – i.e. feature engineering – allows for better performance compared to models that intake 10 Sentinel-2 spectral bands. Introducing a random temporal shift to the EVI timeseries further increases performance; by increasing the sample variance seen by the model, randomly shifting the input timeseries improves model transferability. For reference, classifier performance based on the prediction admissibility criteria is also included.

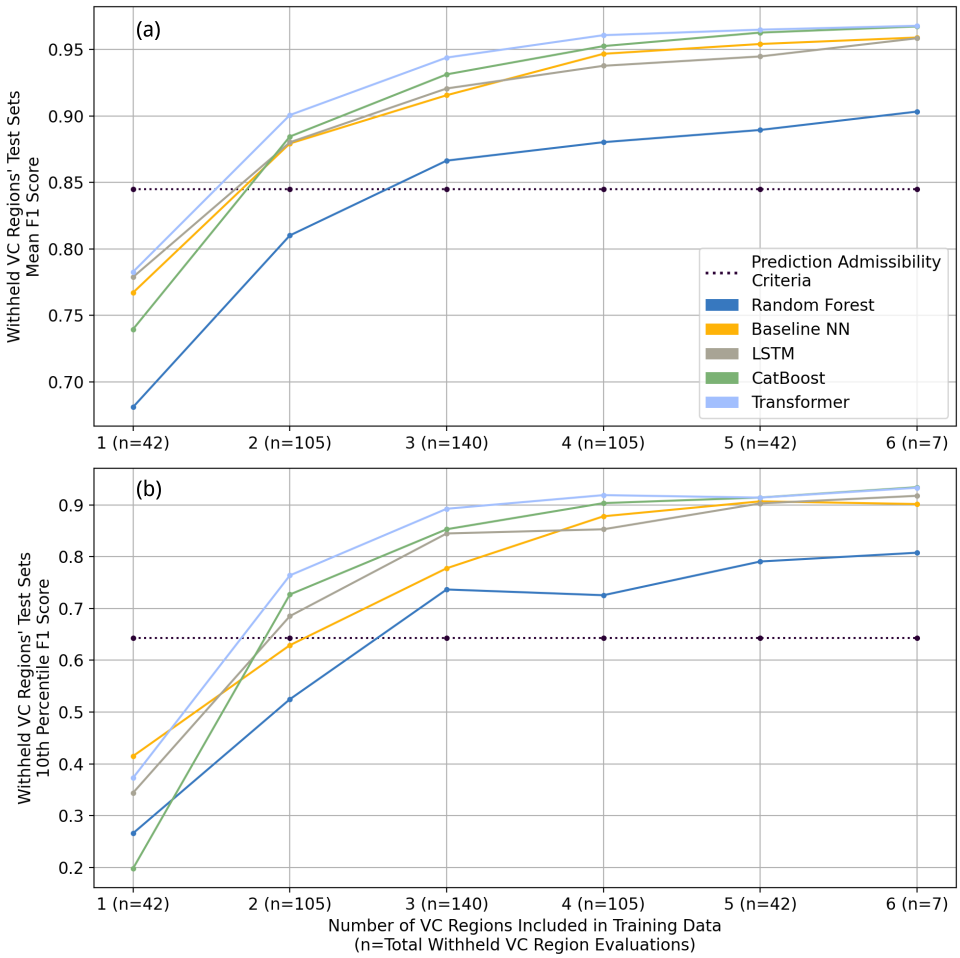

Figure 6: Withheld region test dataset performance for different classifier models, organized along the *x*-axis by the number of regions included during training. (a) presents mean F_1 score over the withheld regions; (b) presents the 10th percentile F_1 score over the withheld regions. Results indicate that the transformer based classifier yields the best performance, followed closely by the CatBoost model. F_1 scores from classification based on the prediction admissibility criteria are presented for reference.

Taken together, randomly shifted EVI timeseries increase withheld region F_1 scores by an average of 0.22 when only 2 VC regions are included in the training data, compared to models that use all spectral bands. As performance begins to plateau with 4 or more VC regions included in the training data, this gap shrinks to an improvement of 0.10. Similar results can be seen in Figure 5(b) for the low-end of performance: Extracting and randomly shifting EVI timeseries increase the 10th percentile of withheld region F_1 scores by 0.40 when 2 VC regions are included in the training data, a difference that shrinks to approximately 0.14 with 5 or more VC regions in the training data. Two additional findings are gleaned from the results for the model trained on randomly shifted EVI timeseries (i.e. the grey curve). First, a classifier trained on data from 2 VC regions or more outperforms the pixel filtering baseline. Second, increasing the

number of VC regions included in the training set improves withheld region prediction performance up around 4 VC regions before tapering off.

Figure 6 displays (a) mean and (b) 10^{th} percentile F_1 score for all combinations of VC regions included in training for the 5 classification models tested, along with the reference classifier based on the prediction admissibility criteria. Figure 6 demonstrates that the transformer architecture is most robust for all combinations of VC training regions, followed closely by the CatBoost architecture for all training configurations with 2 or more VC regions. Moreover, for models with 5 or 6 VC regions included in training, mean and low-end F_1 scores for these two architectures are practically indistinguishable at 0.97 and 0.92, respectively. The Supplementary Materials contain further comparisons between Transformer and CatBoost performance (see Supplementary Table S5), showing that when each model is trained on all 8 regions' training data, the two models demonstrate an average regional prediction alignment of 98.9%. Moreover, an ablation study on training dataset size finds that reducing the proportion of polygons in the training set from 70% to 15% has minimal impact on prediction performance (See Supplementary Figure S5). Lastly, Figure 6 shows that the LSTM architecture does not noticeably improve performance compared to the baseline neural network, and that the trained Random Forest models yield the worst performance in withheld regions.

Next, prediction performance over the unseen ground-collected labels in Tana is assessed. As the transformer model demonstrates the most robust performance over unseen regions' samples, it is selected for prediction, achieving 96.7% accuracy over irrigated samples (88,128/91,898) and 95.9% accuracy over non-irrigated samples (33,954/35,121) for an F_1 score of 0.932.

4.2 Model inference

For model inference, the transformer architecture is trained on the randomly shifted EVI timeseries of the labeled data from the 7 VC and one GC regions. The trained model is then deployed over Tigray and Amhara for the 2020 and 2021 dry seasons (using imagery collected between June 1, 2019 and June 1, 2020; and between June 1, 2020 and June 1, 2021, respectively). Two post-processing steps are then taken: 1) the prediction admissibility criteria are applied, and 2) contiguous groups of predicted irrigated pixels smaller than 0.1 Ha are removed in order to ignore isolated, outlier predictions.

During inference, another step is taken to verify the accuracy of irrigation predictions. Here, five additional enumerators collect 1601 labeled polygons for the 2020 and 2021 dry seasons – 1082 non-irrigated polygons covering 3807 Ha and 519 irrigated polygons covering 582 Ha – across the extent of Amhara via the same labeling methodology used to collect the training, validation, and testing data. The locations of these independently labeled polygons are shown in Supplementary Figure 6. After cluster cleaning and applying the prediction admissibility criteria, these polygons yield 361,451 non-irrigated samples and 48,465 irrigated samples. An F_1 score of 0.917 is achieved over these samples – 98.3% accuracy over non-irrigated samples and 95.5% accuracy over irrigated samples, performance that remains in line with the reported test dataset metrics from Figure 6 and accuracies over the withheld Tana ground-collected labels.

In Figures 7 and 8, bitemporal irrigation maps are presented at a resolution far coarser than their native 10m. The full resolution, georeferenced irrigation maps are available from the corresponding author upon request.

4.2.1 Tigray

Figure 7 presents predicted irrigated areas in Tigray for 2020 and 2021, with 2020 irrigation predictions in the red channel and 2021 irrigation predictions in the green and blue channels, yielding cyan. To better understand the nature of changing vegetation phenologies across this time period, the inset of Figure 11 contains example timeseries that produced an irrigation prediction in one of 2020 or 2021. These example timeseries show that a second crop cycle with vegetation growth peaking in January is associated with a positive irrigation prediction; in contrast, the non-existence of this cycle is associated with non-irrigated prediction. Table 2 displays the total predicted irrigated area for Tigray for 2020 and 2021, along with the total land area, organized by zone. Between 2020 and 2021, Table 2 quantifies a 39.8% decline in irrigated area in Tigray.

4.2.2 Amhara

Figure 8 presents a bitemporal irrigation map for Amhara. This map contains a predominance of irrigated predictions around Lake Tana in the zones of Central Gondar, South Gondar, and West Gojjam, an intuitive finding given the availability of water from Lake Tana and the rivers that extend off it. Irrigation is also detected in the portions of Amhara's easternmost zones that fall within the Main Ethiopian Rift (MER); as the valley formed by the MER extends north into Tigray, irrigation predictions in the North Wello, Oromia, and North Shewa zones align with irrigation predictions in the Southern zone of Tigray shown in Figure 7. Table 3 displays the total predicted irrigated area for

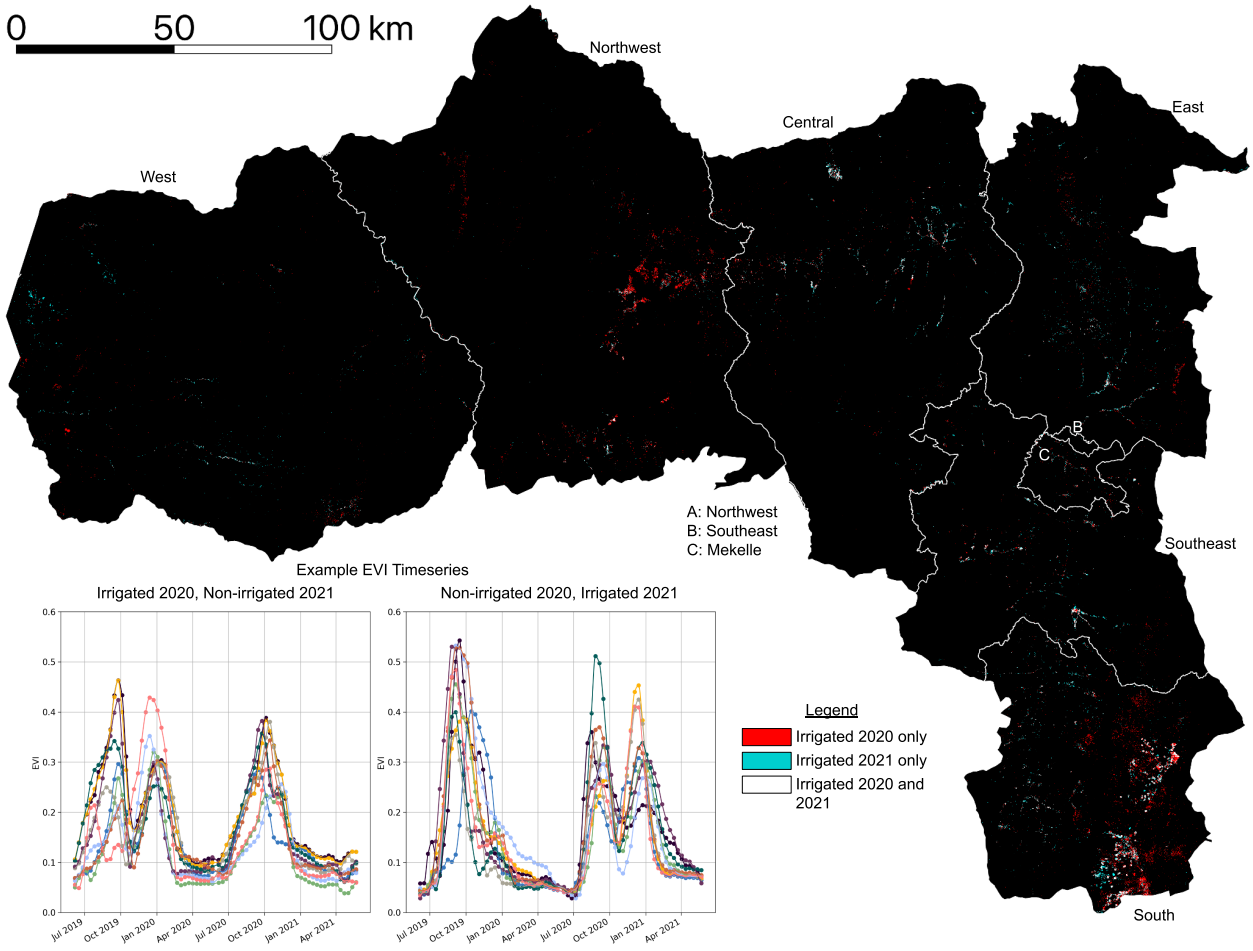

Figure 7: Bitemporal irrigation map for Tigray. Figure inset contains example EVI timeseries predicted as irrigated in either 2020 or 2021. A predominance of red indicates that many parts of Tigray contain irrigation detected in 2020 but not in 2021.

Amhara for 2020 and 2021, along with the total land area, organized by zone. Between 2020 and 2021, Table 3 quantifies a 41.6% decline in irrigated area in Amhara.

The inset of Figure 8 presents interannual irrigated cropping patterns for an area southwest of Choke Mountain. Interlocking red and cyan plots indicate the spatial rotation of irrigated crops from 2020 to 2021; no white plots are observed, which would signify dry season crop growth in both years.

5 Discussion

This manuscript makes a set of contributions to the literature for learning from limited labels. First, it demonstrates a process of collecting training data to supplement ground-collected labels, instead of using imagery from a single timestep or user-imposed vegetation content heuristics. This process improves on previous methods of sample collection as it verifies the existence or non-existence of full vegetation cycles during the dry season. Second, an evaluation of inputs, classifier architectures, and training strategies is presented for achieving irrigation classifier transferability. Results indicate that enhanced vegetation timeseries (EVI) timeseries outperform a full set of spectral bands as inputs; that randomly shifting input timeseries prevents classifier models from overfitting to region-specific input features; and that a transformer-based neural network produces the highest prediction accuracies in unseen target regions. Due to the close similarity of performance metrics and alignment of predictions, the faster training, more easily interpretable CatBoost architecture is also shown as a suitable alternative for irrigation mapping efforts.

Table 2: Predicted irrigated area statistics in Tigray for 2020 and 2021, organized by zone.

Zone	Irrigated Ha., 2020	Irrigated Ha. 2021	Total Ha.	Percent Change, 2020 to 2021	Percent Change as Fraction of Total Area, 2020 to 2021
Central	3710	3554	954,616	-4.2%	0.0%
Eastern	3068	2863	635,670	-6.7%	0.0%
Mekelle	556	397	52,313	-28.5%	-0.3%
North Western	7439	2062	1,246,715	-72.3%	-0.4%
South Eastern	2658	2301	533,334	-13.4%	-0.1%
Southern	16,474	8064	506,151	-51.1%	-1.7%
Western	2278	2557	1,331,652	12.3%	0.0%
Total	36,181	21,799	5,260,451	-39.8%	-0.3%

Table 3: Predicted irrigated area statistics in Amhara for 2020 and 2021, organized by zone.

Zone	Irrigated Ha., 2020	Irrigated Ha. 2021	Total Ha.	Percent Change, 2020 to 2021	Percent Change as Fraction of Total Area, 2020 to 2021
Awi	27,443	20,547	906,682	-25.1%	-0.8%
Central Gondar	73,450	50,954	2,095,018	-30.6%	-1.1%
East Gojam	44,975	33,888	1,405,689	-24.7%	-0.8%
North Gondar	7381	3367	684,247	-54.4%	-0.6%
North Shewa (AM)	62,933	21,362	1,622,197	-66.1%	-2.6%
North Wello	21,367	8250	1,110,856	-61.4%	-1.2%
Oromia	30,875	5285	380,773	-82.9%	-6.7%
South Gondar	72,682	43,046	1,406,698	-40.8%	-2.1%
South Wello	28,215	16,302	1,849,812	-42.2%	-0.6%
Wag Hamra	447	698	890,004	56.4%	0.0%
West Gojam	97,206	71,052	1,348,157	-26.9%	-1.9%
West Gondar	6180	1342	1,529,197	-78.3%	-0.3%
Total	473,155	276,093	15,229,329	-41.6%	-1.3%

Prediction results indicate strong classifier performance over sample timeseries from regions not seen during training. On data from withheld target regions, transformer-based classifiers achieve mean F_1 scores above 0.95 when four or more regions' data are included during training; using labels from all 7 visual collection (VC) regions, the transformer-based classifier achieves an F_1 score of 0.932 on the ground collection (GC) labels around Lake Tana. Over an independently collected set of more than 400,000 samples collected for performance assessment, the same classifier achieves 98.3% accuracy over non-irrigated samples and 95.5% accuracy over irrigated samples.

Deploying a transformer-based classifier trained on all 8 GC + VC regions over the Western Ethiopian Highlands yields insight into changing irrigation patterns. Results suggest that from 2020 to 2021, irrigation in Tigray and Amhara decreased by 40%. In Tigray, this decline was most precipitous in the Northwest and Southern Zones, which saw percent changes in irrigated area of -72.3% and -51.1%, respectively. The Western Zone of Tigray was the only zone to see an increase in irrigated area from 2020 to 2021; even so, this increase amounted to 279 Ha in a zone with a total area of 1,331,652 Ha. Amhara is predicted to have had similar decreases in irrigated area: Apart from the Wag Hamra zone, which was predicted to have less than 0.08% of its area irrigated in 2020 or 2021, all zones in Amhara experienced a change in irrigated area between -25% and -82.3%. The largest declines by area occurred in North Shewa (-41,572 Ha), South Gondar (-29,636 Ha), and West Gojjam (-26,154 Ha). Combined, results for Tigray and Amhara predict severe reductions in dry season crop growth from 2020 to 2021, findings that that align with recent reports of food insecurity following the eruption of civil conflict in Ethiopia in late 2020.

Despite presented performance metrics indicating high levels of prediction accuracy, there are a few limitations to the proposed methodology that are important to mention. First, the study area is limited to the Western Ethiopian Highlands, a highly agricultural, climatologically consistent region that is dominated by rainfed cropped phenologies. As the irrigation classifiers are only trained to separate dry season crop cycles from native vegetation cycles – associating

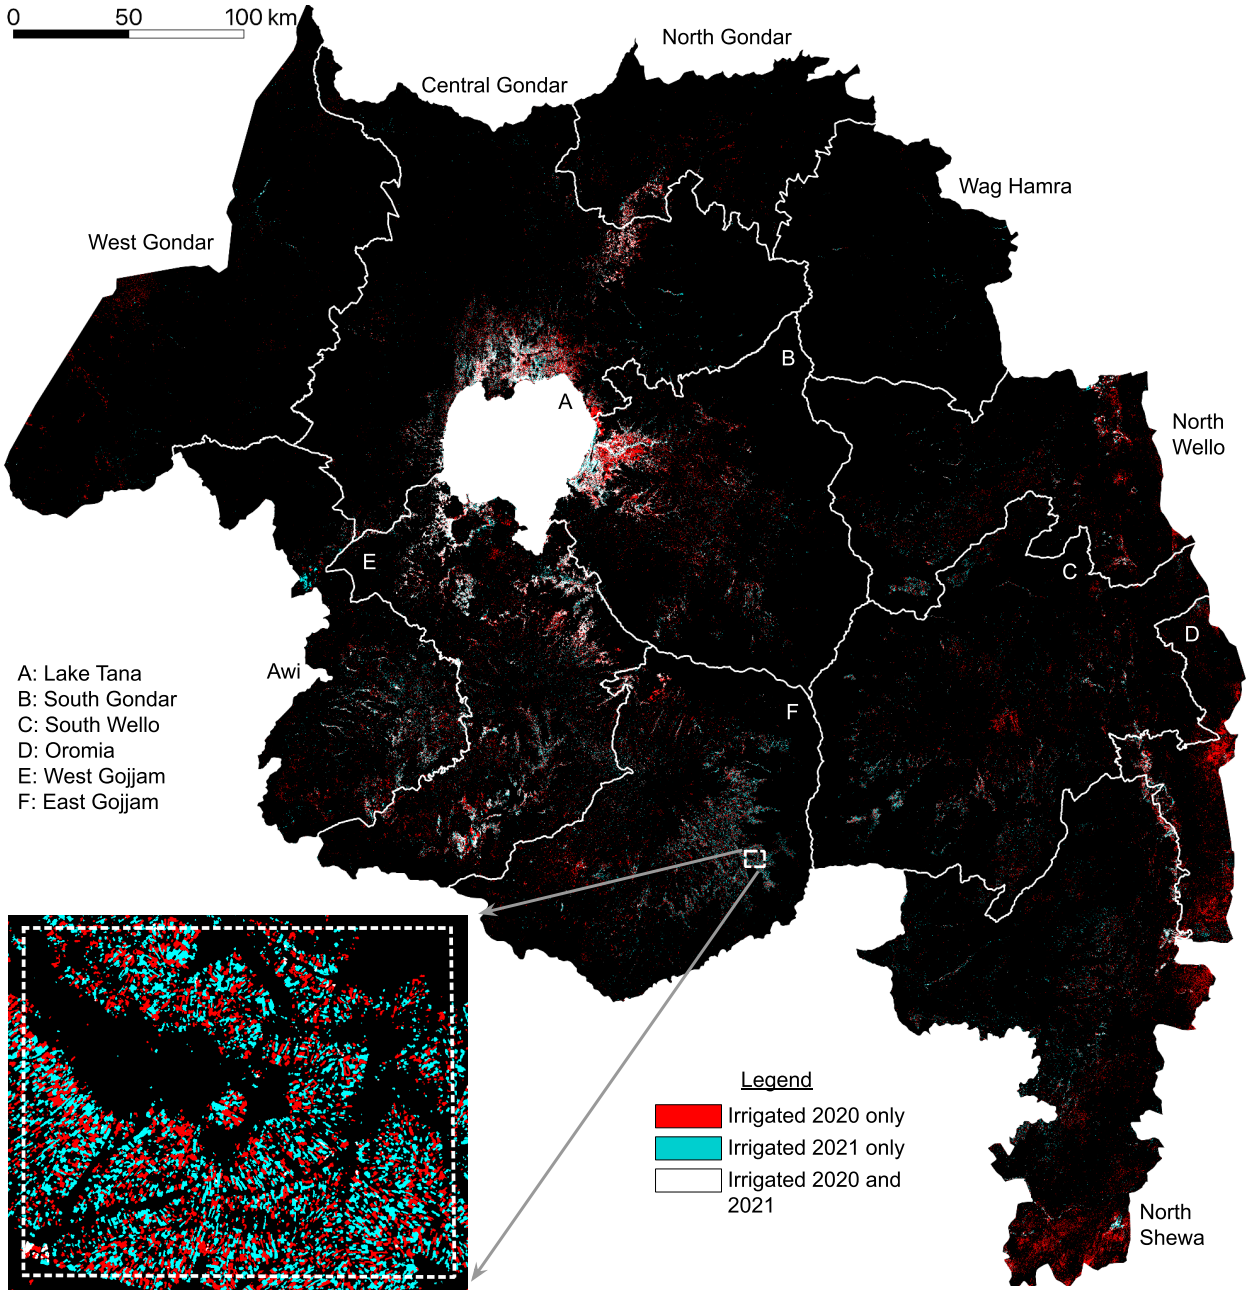

Figure 8: Bitemporal irrigation map for Amhara. Figure inset contains example predictions around Choke Mountain displaying interannual irrigation patterns. A predominance of red indicates that many parts of Amhara contain irrigation detected in 2020 but not in 2021.

identified dry-season cropping with irrigation presence – they will perform poorly in settings with different rainfall and phenological patterns. Relatedly, the trained irrigation classifiers do not identify irrigation used to supplement wet season precipitation, irrigation of tree crops, evergreen riparian vegetation, or irrigation that supports continuous cropping, as the phenological signatures of these types of vegetation are difficult to distinguish from evergreen, non-cropped signatures. This discrimination task is left for future work. Lastly, classifiers are trained only on cropped phenologies, which constitute a portion of the vegetation signatures that exist in the area of interest. To manage the other phenologies present at model inference, prediction admissibility criteria are implemented. Nevertheless, these

criteria are imperfect: There are surely irrigated pixels which have been mistakenly assigned a non-irrigated class label, along with non-cropped pixels which have evaded the cropland filters.

While the presented methodology is applied only for the task of irrigation identification in the Ethiopian highlands, the strategy of regional phenological characterization to provide context for geographically informed selection of training samples can be applied more broadly to a range of land cover mapping objectives. The applicability of this approach in the field of machine learning with limited labels is supported by results comparing classifier architectures and hyperparameter choice to assess the question of result uniqueness that overshadows all land cover classifications. As discussed by Small [2021], what is presented as *the* map is often just *a* map – one of many different products that can be obtained from the same set of inputs with different classifiers and hyperparameter settings. By assessing multiple classifier architectures and quantifying sensitivity of results to hyperparameter settings, this approach demonstrates consistency in results and indicates the uncertainty that can be expected of the resulting irrigation maps; as such, it provides a process for building robust, transferable classifiers in settings with scarce labeled data.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

The Author Contributions section is mandatory for all articles, including articles by sole authors. If an appropriate statement is not provided on submission, a standard one will be inserted during the production process. The Author Contributions statement must describe the contributions of individual authors referred to by their initials and, in doing so, all authors agree to be accountable for the content of the work. Please see here for full authorship criteria.

Funding

Partial support for this effort was provided by the National Science Foundation (INFEWS Award Number 1639214), Columbia World Projects, Rockefeller Foundation (eGuide Grant 2018POW004), OPML UK (DFID) and Technoserve (BMGF).

Acknowledgments

The authors are grateful to Jack Bott, Yinbo Hu, Hasan Siddiqui, and Yuezi Wu for their assistance in labelling. The authors would like to thank Gunther Bensch (RWI), Andrej Kveder (OPML) and Abiy Tamerat (EthioResource Group), Yifru Tadesse (ATA Ethiopia) and Esther Kim (Technoserve) for their assistance with field data collection efforts, Rose Rustowicz with use of Descartes Labs platform, and colleagues Jay Taneja (UMass), Markus Walsh (AfSIS) and Edwin Adkins (Columbia) with their continued stimulating discussions and guidance.

Data Availability Statement

Label data and 250m-resolution predictions are available from the corresponding author upon request.

References

M. F.A. Vogels, S. M. de Jong, G. Sterk, and E. A. Addink. Mapping irrigated agriculture in complex landscapes using SPOT6 imagery and object-based image analysis – A case study in the Central Rift Valley, Ethiopia –. *International Journal of Applied Earth Observation and Geoinformation*, 75(May 2018):118–129, 2019a. ISSN 1872826X. doi:10.1016/j.jag.2018.07.019. URL https://doi.org/10.1016/j.jag.2018.07.019.

Steve Wiggins, Dominic Glover, and Alex Dorgan. Agricultural innovation for smallholders in sub-Saharan Africa. Technical Report July, 2021. URL https://webarchive.nationalarchives.gov.uk/ukgwa/20211030121337/https://degrp.odi.org/publication/agricultural-innovation-for-smallholders-in-sub-saharan-africa/.

T. Conlon, Y. Wu, C. Small, H. Siddiqui, E. Adkins, and V. Modi. A Novel Method of Irrigation Detection and Estimation of the Effects of Productive Electricity Demands on Energy System Planning. In *AGU Fall Meeting Abstracts*, volume 2020, pages GC034–08, December 2020.

Gebrehaweria Gebregziabher, Regassa E. Namara, and Stein Holden. Poverty reduction with irrigation investment: An empirical case study from Tigray, Ethiopia. *Agricultural Water Management*, 96(12):1837–1843, 2009. ISSN 03783774. doi:10.1016/j.agwat.2009.08.004.

M. Shahriar Pervez, Michael Budde, and James Rowland. Mapping irrigated areas in Afghanistan over the past decade using MODIS NDVI. *Remote Sensing of Environment*, 149:155–165, jun 2014. ISSN 0034-4257. doi:10.1016/J.RSE.2014.04.008. URL `https://www.sciencedirect.com/science/article/pii/S0034425714001461`.

Marjolein F.A. Vogels, Steven M. de Jong, Geert Sterk, Harke Douma, and Elisabeth A. Addink. Spatio-temporal patterns of smallholder irrigated agriculture in the horn of Africa using GEOBIA and Sentinel-2 imagery. *Remote Sensing*, 11(2), 2019b. ISSN 20724292. doi:10.3390/rs11020143.

Yaoliang Chen, Dengsheng Lu, Lifeng Luo, Yadu Pokhrel, Kalyanmoy Deb, Jingfeng Huang, and Youhua Ran. Detecting irrigation extent, frequency, and timing in a heterogeneous arid agricultural region using MODIS time series, Landsat imagery, and ancillary data. *Remote Sensing of Environment*, 204(March 2017):197–211, 2018. ISSN 00344257. doi:10.1016/j.rse.2017.10.030. URL `https://doi.org/10.1016/j.rse.2017.10.030`.

Alfredo Huete, Chris Justice, and Wim Van Leeuwen. Modis vegetation index (mod13). *Algorithm theoretical basis document*, 3(213):295–309, 1999.

Darius Phiri, Matamayo Simwanda, Serajis Salekin, Vincent R. Ryirenda, Yuji Murayama, Manjula Ranagalage, Nadya Oktaviani, Hollanda A Kusuma, Tianxiang Zhang, Jinya Su, Cunjia Liu, Wen Hua Chen, Hui Liu, Guohai Liu, M. Cavur, H. S. Duzgun, S. Kemec, D. C. Demirkan, Radhia Chairet, Yassine Ben Salem, Mohamed Aoun, Zolo Kiala, Onisimo Mutanga, John Odindi, and Kabir Peerbhay. Remote Sensing Sentinel-2 Data for Land Cover / Use Mapping: A Review. *Remote Sensing*, 42(3):14, 2019. ISSN 16821750.

Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. ISSN 14764687. doi:10.1038/nature14539.

Liheng Zhong, Lina Hu, Hang Zhou, and Xin Tao. Deep learning based winter wheat mapping using statistical data as ground references in Kansas and northern Texas, US. *Remote Sensing of Environment*, 233:111411, nov 2019. ISSN 0034-4257. doi:10.1016/J.RSE.2019.111411. URL `https://www.sciencedirect.com/science/article/pii/S0034425719304304`.

Ying Li, Haokui Zhang, Xizhe Xue, Yenan Jiang, and Qiang Shen. Deep learning for remote sensing image classification: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(6):1–17, 2018. ISSN 19424795. doi:10.1002/widm.1264.

Jia Deng, Wei Dong, R. Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. pages 248–255, 2009. doi:10.1109/cvprw.2009.5206848.

Chao Tao, Ji Qi, Weipeng Lu, Hao Wang, and Haifeng Li. Self-Supervised Paradigm Under Limited Labeled Samples. 19:1–5, 2020.

Biplab Banerjee, Francesca Bovolo, Avik Bhattacharya, Lorenzo Bruzzone, Subhasis Chaudhuri, and B. Krishna Mohan. A new self-training-based unsupervised satellite image classification technique using cluster ensemble strategy. *IEEE Geoscience and Remote Sensing Letters*, 12(4):741–745, 2015. ISSN 1545598X. doi:10.1109/LGRS.2014.2360833.

Sudipan Saha, Yady Tatiana Solano-Correa, Francesca Bovolo, and Lorenzo Bruzzone. Unsupervised deep learning based change detection in Sentinel-2 images. *2019 10th International Workshop on the Analysis of Multitemporal Remote Sensing Images, MultiTemp 2019*, pages 0–3, 2019. doi:10.1109/Multi-Temp.2019.8866899.

Xingrui Yu, Xiaomin Wu, Chunbo Luo, and Peng Ren. Deep learning in remote sensing scene classification: a data augmentation enhanced convolutional neural network framework. *GIScience and Remote Sensing*, 54(5):741–758, 2017. ISSN 15481603. doi:10.1080/15481603.2017.1323377. URL `https://doi.org/10.1080/15481603.2017.1323377`.

Radamanthys Stivaktakis, Grigorios Tsagkatakis, and Panagiotis Tsakalides. Deep Learning for Multilabel Land Cover Scene Categorization Using Data Augmentation. *IEEE Geoscience and Remote Sensing Letters*, 16(7):1031–1035, 2019. ISSN 15580571. doi:10.1109/LGRS.2019.2893306.

Hassan Bazzi, Nicolas Baghdadi, Ghaith Amin, Ibrahim Fayad, Mehrez Zribi, Valérie Demarez, and Hatem Belhouchette. An operational framework for mapping irrigated areas at plot scale using sentinel-1 and sentinel-2 data. *Remote Sensing*, 13(13):1–28, 2021. ISSN 20724292. doi:10.3390/rs13132584.

Mutlu Ozdogan, Yang Yang, George Allez, and Chelsea Cervantes. Remote sensing of irrigated agriculture: Opportunities and challenges. *Remote Sensing*, 2(9):2274–2304, 2010. ISSN 20724292. doi:10.3390/rs2092274.

Hassan Bazzi, Nicolas Baghdadi, Ibrahim Fayad, Mehrez Zribi, Hatem Belhouchette, and Valérie Demarez. Near real-time irrigation detection at plot scale using sentinel-1 data. *Remote Sensing*, 12(9), 2020. ISSN 20724292. doi:10.3390/RS12091456.

Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010. ISSN 10414347. doi:10.1109/TKDE.2009.191.

Rafael Pires de Lima and Kurt Marfurt. Convolutional neural network for remote-sensing scene classification: Transfer learning analysis. *Remote Sensing*, 12(1), 2020. ISSN 20724292. doi:10.3390/rs12010086.

Christopher Small. Spatiotemporal dimensionality and Time-Space characterization of multitemporal imagery. *Remote Sensing of Environment*, 124:793–809, 2012. ISSN 00344257. doi:10.1016/j.rse.2012.05.031. URL `http://dx.doi.org/10.1016/j.rse.2012.05.031`.

Leo Breiman. Random Forests. *Machine Learning*, pages 1–28, 2001. doi:10.1201/9780429469275-8.

Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: unbiased boosting with categorical features. 2019.

Zhiguang Wang, Weizhong Yan, and Tim Oates. Time series classification from scratch with deep neural networks: A strong baseline. *Proceedings of the International Joint Conference on Neural Networks*, 2017-May:1578–1585, 2017. doi:10.1109/IJCNN.2017.7966039.

Gary King and Langche Zeng. Logistic regression in rare events data. *Political Analysis*, 9(2):137–163, 2001. ISSN 15487660. doi:10.18637/jss.v008.i02.

Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pages 1–15, 2015.

Christopher Small. Grand Challenges in Remote Sensing Image Analysis and Classification. *Frontiers in Remote Sensing*, 1(April):1–4, 2021. doi:10.3389/frsen.2020.605220.