
Annual and in-season mapping of cropland at field scale with sparse labels

Gabriel Tseng *

gabrieltseg95@gmail.com
NASA Harvest

Hannah Kerner *

hkerner@umd.edu
University of Maryland, College Park

Catherine Nakalembe

cnakalem@umd.edu

University of Maryland, College Park

Inbal Becker-Reshef

ireshef@umd.edu

University of Maryland, College Park

Abstract

Spatial information about where crops are being grown, known as cropland maps, are critical inputs for analyses and decision-making related to food security and climate change. Despite a widespread need for readily-updated annual and in-season cropland maps at the management (field) scale, these maps are unavailable for most regions at risk of food insecurity. This is largely due to lack of in-situ labels for training and validating machine learning classifiers. Previously, we developed a method for binary classification of cropland that learns from sparse local labels and abundant global labels using a multi-headed LSTM and time-series multispectral satellite inputs over one year. In this work, we present a new method that uses an autoregressive LSTM to classify cropland during the growing season (i.e., partially-observed time series). We used these methods to produce publicly-available 10m-resolution cropland maps in Kenya for the 2019-2020 and 2020-2021 growing seasons. These are the highest-resolution and most recent cropland maps publicly available for Kenya. These methods and associated maps are critical for scientific studies and decision-making at the intersection of food security and climate change.

1 Introduction

Agriculture and climate change are entangled in a negative feedback loop, as agriculture is both a driver of climate change and one of the most vulnerable sectors impacted by it. Agriculture land use contributes to over 20% of global CO₂ emissions and over half of CO₂ emissions in Africa [3]. Still, this land use is critical for food security and economic development, yet global crop production is projected to decrease by 2-6% every decade due to climate change [24]. In a 2019 report, annual per capita maize production in Kenya was found to have declined from ~ 140 kg per person in the early 1980s to 70 kg in 2010–2017 [4]. Hence, more timely and accurate mechanisms as well as investments that support the monitoring and forecasting of food security have become even more critical [1, 16]. Climate-smart policies and practices that increase agricultural productivity and climate resilience without sacrificing carbon sinks such as forests are needed to reverse these trends. Analyses such as grain yield and production estimates are needed to help design future climate-smart initiatives, but these analyses require accurate, up-to-date maps of where and which crops are being grown at the field scale. These datasets do not exist or are not publicly accessible for most of Africa.

*Both authors contributed equally to this research.




Figure 1: Models for post-season (a) and in-season (b) crop classification

Machine learning techniques that use Earth observation (EO) data are a promising approach to generate cropland maps, but existing approaches often cover small (e.g., [11, 17, 9, 18, 26, 15]) or homogeneous (e.g., [12, 27, 23]) areas, require large labeled datasets for training (e.g [20, 2, 10]), and/or lack sufficient resolution for identify smallholder farms (e.g., [21, 28]). In addition, these classifiers typically require at least a full year of observations and are produced months to years after the map year, which limits their use for informing early warning decisions.

In this paper, we present a new method for in-season classification of smallholder croplands that uses an autoregressive LSTM for time series forecasting of satellite observations. We used this method to predict the multispectral satellite time series for the remainder of the 2020-2021 growing season to produce an open-access 2020 cropland map for Busia county, Kenya. We also present a new nation-wide 10m resolution cropland map for 2019 in Kenya using our multi-headed LSTM approach [13]. These are the highest resolution and most up-to-date cropland maps publicly available for Kenya. All source code², photo-interpreted labels³, and maps³ from this study are openly available to promote operational uptake and future research.

2 Data

Labels. We compiled labeled samples within Kenya from three different sources: (i) 8,449 plots in Busia county (2017-2019) collected by Plant Village, (ii) 2,207 maize plots in Western and Central Kenya collected by One Acre Fund, and (iii) 2,969 hand labelled non-crop points collected using high-resolution Google Earth Imagery from 2019. As in [13], these labels were supplemented with the GeoWiki cropland dataset, a global dataset of 35,866 crop/non-crop labels [22].

EO Data. We used Sentinel-2 top of atmosphere reflectance (Level 1C) observations overlapping with label locations as input. We exported 160×160 m patches for each labeled pixel using Google Earth Engine (GEE) [7], then extracted the closest pixel within the patch to the label location. We prepared monthly, cloud-free multi-spectral time series inputs with 10m spatial resolution following the same procedure described in [13]. All input data were normalized to have a band-wise mean of 0 and standard deviation of 1 (calculated from the training and validation datasets). We used observations acquired April 2017-April 2018 for GeoWiki labels and April 2019-April 2020 otherwise.

3 Methods

3.1 Crop classification model

We trained a one layer long short term memory (LSTM) network [8] to classify pixels as crop/non-crop. The model input is the 12 month time-series for a single pixel. To allow the model to focus on the local region of interest, Kenya, while also learning from diverse global examples of cropland, we used a multi-headed LSTM (Figure 1a) as described in [13]. The LSTM base learns an embedding from the input multispectral time series. The output of the LSTM base is passed to a separate model that classifies examples from the Kenya and GeoWiki datasets separately. Each classifier consists of

²<https://github.com/nasaharvest/kenya-crop-mask>

³<https://doi.org/10.5281/zenodo.4271143>

linear layers followed by a sigmoid activation. All layers are trained end-to-end with a combined loss function. A single batch could contain both Kenya and GeoWiki labels. We combined the losses from both classifiers into the following loss function to train the models:

$$\mathcal{L}_{\text{classifier}} = \frac{W}{\alpha} \mathcal{L}_{\text{GeoWiki}} + \mathcal{L}_{\text{Kenya}} \quad (1)$$

$\mathcal{L}_{\text{GeoWiki}}$ and $\mathcal{L}_{\text{Kenya}}$ are binary cross-entropy losses for the GeoWiki and Kenya classification layers, and α is a weighting parameter to encourage the model to learn features more relevant for Kenya. W is the ratio of global and local instances used to weight the contribution from each loss in the batch ($W = \frac{\text{Number of GeoWiki instances in batch}}{\text{Number of Kenya instances in batch}}$).

3.2 Forecasting model

We trained an autoregressive LSTM to predict the timesteps required to fill in the complete 12-month time series input. This enables the multi-headed LSTM classifier to be used during the growing season. Specifically, if we have N months of observations and require a total of T months, then we need to predict $P = T - N$ timesteps. In this study, we had 5 months of data available (April-August) and predicted the 7 remaining timesteps (September-March) for a total of $T = 12$ timesteps. The autoregressive LSTM predicts the 2^{nd} to T^{th} timesteps. Up to the N^{th} timestep, observed input data is fed to the model. Past the N^{th} timestep, the model is given own predictions to predict the next timestep. This process is shown in Figure 1b. We trained the model using a smooth-l1 loss [6]. The loss for a single training sample is given by:

$$\mathcal{L}_{\text{encoder}}(\hat{y}, y) = \frac{1}{T-1} \sum_{i=2}^T \begin{cases} 0.5(\hat{y}_i - y_i)^2, & \text{if } |\hat{y}_i - y_i| < 1 \\ |\hat{y}_i - y_i| - 0.5, & \text{otherwise} \end{cases} \quad (2)$$

We applied variational dropout [5] with dropout probability 0.2 between timesteps in the LSTM. In addition, we applied a linear layer to the hidden output of the LSTM so that each timestep's output has the same shape as the input, independent of hidden vector size.

3.3 End to end training

We trained the forecasting and crop classification model end to end, combining the losses of both models. The crop mapping model is trained on both the forecasted data, $\mathcal{X}_{\text{forecasted}}$, and the original complete input $\mathcal{X}_{\text{observed}}$. We concatenated the observed N months with the forecasted P months of multispectral observations to form the T -month time series input. Therefore, for observed observations x_i and forecasted observations o_i , $\mathcal{X}_{\text{forecasted}} = \{x_0, x_1, \dots, x_N, o_{N+1}, \dots, o_T\}$. The forecaster and cropland classification models are trained end-to-end with a combined loss function:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{encoder}} + \frac{1}{2} (\mathcal{L}_{\text{classifier}}(\mathcal{X}_{\text{observed}}) + \mathcal{L}_{\text{classifier}}(\mathcal{X}_{\text{forecasted}})) \quad (3)$$

We trained the model using an Adam optimizer [14] and early stopping with a patience of 10. A validation set consisting of 10% of the data was used for early stopping, with the validation loss used to determine the early stopping point. To improve model generalization, we applied Gaussian noise with standard deviation 0.1 to the observed inputs when training the classifier.

4 Results

We evaluated the models on a test set consisting of 10% of the total dataset. We compared the performance of the in-season classifier when (i) the complete timeseries is forecasted and (ii) when the classifier receives the partial timeseries (no forecasting). Figure 2a shows the AUC score for each model as more time steps are observed. Forecasting future timesteps boosts performance when few timesteps have been observed. As the number of observed timesteps approaches the length of the full time series, the forecasted and partial model scores converge to the performance of the model given the full time series. The normalized NDVI predictions of the forecaster trained with 5 months of input data are plotted in Figure 2b for both a crop and non crop test instance.




Figure 2: (a) MAE of predicted timesteps and AUC score for classifier with combined observed + forecasted timesteps as more timesteps are observed, compared to using the partial time series. Dotted horizontal line shows AUC using fully-observed time series. (b) A plot of forecasted NDVI values compared to true values for a crop and non crop test instance. The dotted vertical line indicates where the model begins receiving its predictions as input.




Figure 3: Cropland probability maps

We used the trained LSTM for crop classification (no forecasting) to produce a 10m cropland map for Kenya for the 2019-2020 season (Figure 3a). Following [19], we labeled an independent random stratified test sample to evaluate the map accuracy. We randomly sampled 250 points classified as crop and non-crop each and labeled them based on visual interpretation of 3m resolution PlanetScope quarterly mosaics [25]. Of the 500 points, we labeled 183 as crop and 272 as non-crop; we omitted 45 points that we could not confidently determine labels for using available imagery. We determined that our Kenya cropland map has overall accuracy 0.86, precision 0.77, recall 0.92, and F1 score 0.84. In addition, we used the trained LSTM *with* forecasting to produce an in-season cropland map for Busia county (an important county for agriculture in Kenya) for the 2020-2021 season. Using the same technique as for Kenya, we obtained a test sample of 230 crop and 270 non-crop points. We determined that our Busia in-season cropland map has overall accuracy 0.79, precision 0.75, recall 0.81 and F1 score 0.78.

5 Conclusion

Up-to-date cropland maps that capture smallholder farms at the field scale are critical inputs for assessments of agricultural productivity in relation to climate and land management. We present a new method for binary classification of cropland that can be used to produce in-season cropland maps. This method uses an autoregressive LSTM to forecast multispectral satellite observations and is trained end-to-end with a multi-headed LSTM cropland classifier. Using these methods, we produced high-resolution (10m) cropland maps of the entirety of Kenya for the 2019-2020 growing seasons and of Busia county, Kenya for the 2020-2021 growing seasons.

References

- [1] W. Chivasa, O. Mutanga, and C. Biradar. Application of remote sensing in estimating maize grain yield in heterogeneous african agricultural landscapes: A review. *Int. J. Remote Sens.*, 38(23):6816–6845, dec 2017.
- [2] P. Defourny, S. Bontemps, N. Bellemans, C. Cara, G. Dedieu, E. Guzzonato, O. Hagolle, J. Ingla, L. Nicola, T. Rabaute, et al. Near real-time agriculture monitoring at national scale at parcel resolution: Performance assessment of the sen2-agri automated system in various cropping systems around the world. *Remote sensing of environment*, 221:551–568, 2019.
- [3] FAO. The contribution of agriculture to greenhouse gas emissions, 2020.
- [4] C. Funk, F. Davenport, G. Eilerts, N. Nourey, and G. Galu. Contrasting Kenyan Resilience to Drought: 2011 and 2017. Technical report, 2018.
- [5] Y. Gal and Z. Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1019–1027. Curran Associates, Inc., 2016.
- [6] R. Girshick. Fast R-CNN. *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [7] N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202:18 – 27, 2017.
- [8] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [9] S. Ji, C. Zhang, A. Xu, Y. Shi, and Y. Duan. 3D convolutional neural networks for crop classification with multi-temporal remote sensing images. *Remote Sensing*, 10(1):75, 2018.
- [10] Z. Jin, G. Azzari, C. You, S. Di Tommaso, S. Aston, M. Burke, and D. B. Lobell. Smallholder maize area and yield mapping at national scales with Google Earth Engine. *Remote Sensing of Environment*, 228:115–128, 2019.
- [11] H. Kerner, C. Nakalembe, and I. Becker-Reshef. Field-Level Crop Type Classification with k Nearest Neighbors: A Baseline for a New Kenya Smallholder Dataset. *arxiv.org*, 2020.
- [12] H. Kerner, R. Sahajpal, S. Skakun, I. Becker-Reshef, B. Barker, M. Hosseini, E. Puricelli, and P. Gray. Resilient in-season crop type classification in multispectral satellite observations using growth stage normalization. 2020.
- [13] H. Kerner, G. Tseng, I. Becker-Reshef, C. Nakalembe, B. Barker, B. Munshell, M. Paliyam, and M. Hosseini. Rapid response crop maps in data sparse regions. In *ACM SIGKDD Conference on Data Mining and Knowledge Discovery Workshops*, 2020.
- [14] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] M.-J. Lambert, P. C. S. Traoré, X. Blaes, P. Baret, and P. Defourny. Estimating smallholder crops production at village level from sentinel-2 time series in mali’s cotton belt. *Remote Sensing of Environment*, 216:647–657, 2018.
- [16] C. Nakalembe. Urgent and critical that Sub-Saharan African Countries invest in Earth Observations (EO) based agriculture early warning and monitoring systems. *Manuscr. Reveiw*, pages 1–3, 2020.
- [17] E. Ndikumana, N. Baghdadi, D. Courault, L. Hossard, and D. Ho Tong Minh. Deep recurrent neural network for agricultural classification using multitemporal sar sentinel-1 for camargue, france. *Remote Sensing*, 10:1217, 08 2018.

- [18] Neetu and S. S. Ray. Exploring machine learning classification algorithms for crop classification using Sentinel-2 data. In *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences (ISPRS)*, volume XLII-3/W6, pages 573–578, 2019.
- [19] P. Olofsson, G. M. Foody, M. Herold, S. V. Stehman, C. E. Woodcock, and M. A. Wulder. Good practices for estimating area and assessing accuracy of land change. *Remote Sensing of Environment*, 148:42–57, 2014.
- [20] M. Rußwurm, C. Pelletier, M. Zollner, S. Lefèvre, and M. Körner. BreizhCrops: A time series dataset for crop type mapping. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences ISPRS (2020)*, 2020.
- [21] K. Samasse, N. P. Hanan, J. Y. Anchang, and Y. Diallo. A High-Resolution Cropland Map for the West African Sahel Based on High-Density Training Data, Google Earth Engine, and Locally Optimized Machine Learning. *Remote Sens. 2020, Vol. 12, Page 1436*, 12(9):1436, may 2020.
- [22] L. See. A global reference database of crowdsourced cropland data collected using the Geo-Wiki platform, 2017.
- [23] A. Shelestov, M. Lavreniuk, V. Vasiliev, L. Shumilo, A. Kolotii, B. Yailymov, N. Kussul, and H. Yailymova. Cloud approach to automated crop classification using Sentinel-1 imagery. *IEEE Transactions on Big Data*, 2019.
- [24] K. R. Smith, A. Woodward, D. Campbell-Lendrum, D. D. Chadee, Y. Honda, J. M. Olwoch, B. Revich, and R. Sauerborn. Human health: impacts, adaptation, and co-benefits.
- [25] P. Team. Planet application program interface: In space for life on earth, 2018–.
- [26] M. Wang, Z. Liu, M. H. A. Baig, Y. Wang, Y. Li, and Y. Chen. Mapping sugarcane in complex landscapes by integrating multi-temporal Sentinel-2 images and machine learning algorithms. *Land Use Policy*, 88:104190, 2019.
- [27] S. Wang, S. Di Tommaso, J. M. Deines, and D. B. Lobell. Mapping twenty years of corn and soybean across the us midwest using the landsat archive. *Scientific Data*, 7(1):1–14, 2020.
- [28] J. Xiong, P. S. Thenkabail, M. K. Gumma, P. Teluguntla, J. Poehnelt, R. G. Congalton, K. Yadav, and D. Thau. Automated cropland mapping of continental africa using google earth engine cloud computing. *ISPRS Journal of Photogrammetry and Remote Sensing*, 126:225 – 244, 2017.