Weak labeling for cropland mapping in Africa

Gilles Hacheme[1], Girmaw Abebe Tadesse[1], Caleb Robinson[2], Rahul Dodhia[2], and Juan Lavista Ferres[2]

[1] Microsoft, AI for Good Lab, Nairobi, Kenya[2] and Microsoft, AI for Good Lab, Redmond, USA[3]

{ghacheme, gtadesse, caleb.robinson, rahul.dodhia, jlavista}@microsoft.com

Abstract

- Cropland mapping is crucial for environmental, agricultural, and food security policies in Africa.
- Unfortunately, most previous cropland maps for Africa are low or medium-resolution.
- Building higher-resolution cropland mapping typically requires extensive human labeling, which is a significant bottleneck for scalability.
- To address this challenge, we propose a method that improves existing weak labels using K-means clustering to train higher-resolution cropland mapping models.
- Our results demonstrate the added value of this approach for large-scale cropland mapping.

Introduction

- Accurate cropland mapping through high-resolution satellite imagery is crucial for environmental and food security policies in Africa.
- Existing methods offer low to medium-sized resolution and require extensive human labeling, which hinders large-scale cropland mapping.
- We propose a method that leverages K-means clustering to improve weak labels and train higher-resolution deep semantic segmentation cropland mapping models.

Problem Statement

Consider an area of interest (AOI) represented by a $k \times k$ dimensional matrix A where a_{ij} is the pixel from A located at (i, j). We assume that we have a corresponding mask M with the same dimensions, where each pixel from M, $m_{ij} \in \{0, 1, 2\}$, and where $0 = unknown$, $1 = non-cropland$, $2 = cropland$. In a cropland mapping semantic segmentation problem, let's consider the distribution of M's pixels in different scenarios:

1. Complete ground truth (CGT): we would like the distribution of m_{ij} to be $\{0 : 0, 1 : 1-p, 2 : p\}$, where $p \in [0, 1]$ is the ground truth proportion of pixels covering croplands.
2. Complete labeling (CL): in this case, we may want to label through a labeling tool every single pixel a_{ij}. In this case, we can reasonably assume that the distribution of m_{ij} is $\{0 : 0, 1 : 1-(p-\epsilon), 2 : p-\epsilon\}$ where $(p-\epsilon) \in [0, 1]$ and $\epsilon \in [0, 1]$ is the level of noise (mislabels) introduced by the labeling.
3. **Partial labeling** (PL): in this case, labeling is not performed on the whole AOI, which is often the case in real-life settings. So, the distribution of m_{ij} is $\{0 : q, 1 : 1-(p-\epsilon+\alpha q), 2 : p-(\epsilon+(1-\alpha)q)\}$ where $(p-(\epsilon+(1-\alpha)q)) \in [0, 1]$, and $q \in [0, 1]$ is the share of unknown/unlabeled pixels.

Using K-means to Strengthen Weak Labels

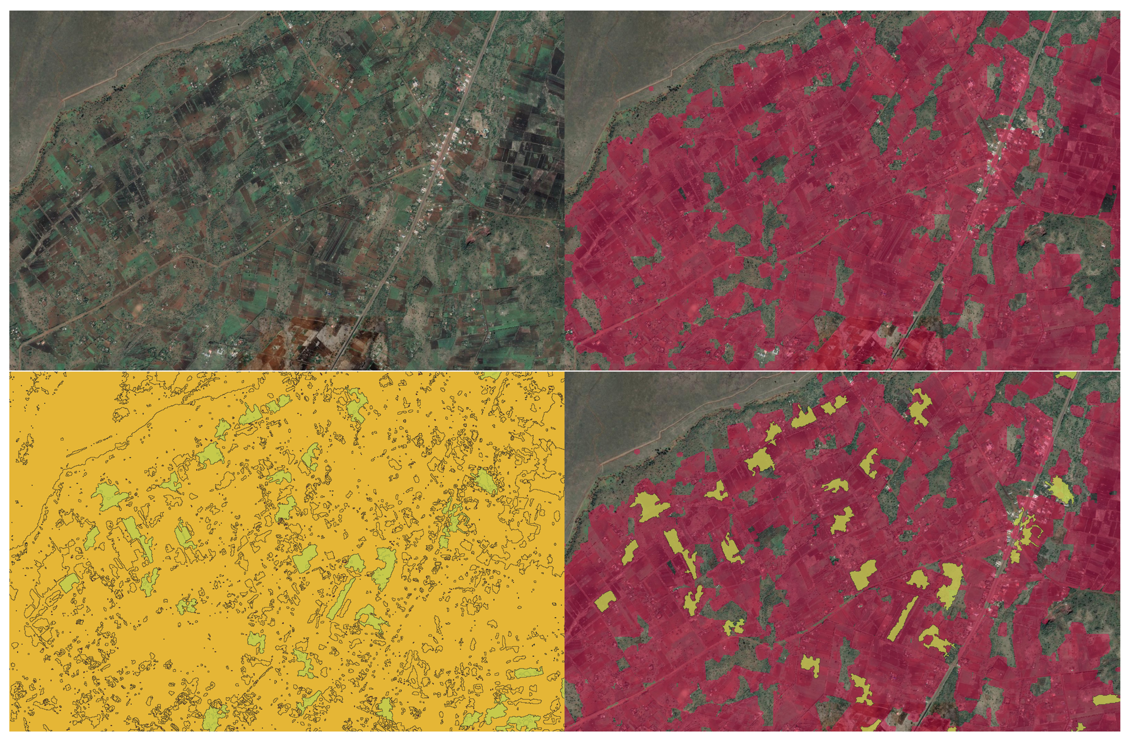

K-means clustering to improve weak labels

- We use weak labels obtained from The Nature Conservancy (TNC) and Planetscope Basemap imagery provided by Norway's International Climate and Forests Initiative (NICFI).
- K-means clustering is used to improve existing weak labels and create better boundaries.

Simulation Experiments

1. Human labels: we train the model on the AOI with the complete set of human labels, and we evaluate on the exact same AOI.
2. Human mined labels: we train the model on the AOI with the complete set of human labels and improved weak (mined) labels.
3. Human mined negative labels: we train the model on the AOI with the complete set of human labels and improved weak (mined) negative labels (i.e., Non-cropland labels only).
4. Human mined positive labels: we train the model on the AOI with the complete set of human labels and improved weak (mined) positive labels.
5. Human TNC labels: we train the model on the AOI with the complete set of human labels and TNC's raw positive weak labels.
6. Human TNC mined negative labels: we train the model on the AOI with the complete set of human labels, TNC's raw positive weak labels, and the improved weak (mined) negative labels.
7. **Half human labels [mined [negative/positive] labels]**: we conduct the same experiments as previously but with only half human labels. This case is for simulating more realistic real-world scenarios where we only have a fraction of the whole data labeled by humans.

Conclusion

- Our proposed approach addresses the human labeling bottleneck for scaling cropland mapping.
- By improving existing weak labels, we can train higher-resolution cropland mapping models.
- Our approach could be an essential tool for large-scale cropland mapping in Africa.

Experiment	Label	%Pixels	Human Labels	Mined Labels	Human Area	Mined Area	F1 score	Precision	Recall
Human labels	Cropland	4.056	67	0	7.093	0.000	0.980	1.000	0.960
	Non-Cropland	95.944	26	0	167.764	0.000	0.991	1.000	0.982
Human mined labels	Cropland	4.056	67	606	7.093	11.016	0.979	0.999	0.960
	Non-Cropland	95.944	26	369	167.764	6.702	0.991	1.000	0.982

Experiment	Label	%Pixels	Human Labels	Mined Labels	Human Area	Mined Area	F1 score	Precision	Recall
Half human labels	Cropland	4.056	34	0	3.439	0.000	0.533	0.408	0.767
	Non-Cropland	95.944	13	0	84.290	0.000	0.962	0.991	0.935
Half human mined labels	Cropland	4.056	34	606	3.439	11.016	0.694	0.553	0.931
	Non-Cropland	95.944	13	369	84.290	6.702	0.974	0.999	0.950
Half human mined negative labels	Cropland	4.056	34	606	3.439	11.016	0.841	0.916	0.777
	Non-Cropland	95.944	13	369	84.290	6.702	0.985	0.992	0.979
Half human mined positive labels	Cropland	4.056	34	606	3.439	11.016	0.324	0.196	0.929
	Non-Cropland	95.944	13	369	84.290	6.702	0.901	0.998	0.821
Half human TNC labels	Cropland	4.056	34	0	3.439	0.000	0.289	0.170	0.960
	Non-Cropland	95.944	13	0	84.290	0.000	0.880	1.000	0.785
Half human TNC mined negative labels	Cropland	4.056	34	606	3.439	11.016	0.581	0.417	0.959
	Non-Cropland	95.944	13	369	84.290	6.702	0.961	1.000	0.925

Note: Some experiment rows have been removed from the table for conciseness. Areas are in km^2.

Deep Learning Indaba Conference 2023, Accra — Ghana