EGU24-11571, updated on 22 Nov 2024
https://doi.org/10.5194/egusphere-egu24-11571
EGU General Assembly 2024

Simple temporal domain adaptation techniques for mapping the inter-annual dynamics of smallholder-dominated croplands over large extents

Lyndon Estes[1], Sam Khallaghi[1], Rahebe Abedi[1], Mary Asipunu[2], Nguyen Ha[1], Boka Luo[1], Cat Mai[1], Amos Wussah[2], Sitian Xiong[1], and Yao-Ting Yao[1]
[1]Graduate School of Geography, Clark University, Worcester, MA, USA
[2]Farmerline Ltd, Kumasi, Ghana

Tracking how agricultural systems are changing is critical to answering important questions related to socioeconomic (e.g. food security) and environmental sustainability (e.g. carbon emissions), particularly in rapidly changing regions such as Africa. Monitoring agricultural dynamics requires satellite-based approaches that can accurately map individual fields at frequent (e.g. annual) intervals over national to regional extents, yet mapping Africa's smallholder-dominated agricultural systems is difficult, as the small and indistinct nature of fields promotes mapping error, while frequent cloud cover leads to coverage gaps. Fortunately, the increasing availability of high spatio-temporal resolution imagery and the growing capabilities of deep learning models now make it possible to accurately map crop fields over large extents. However, the ability to make consistently reliable maps for more than one time point remains difficult, given the substantial domain shift between images collected in different seasons or years, which arises from variations in atmospheric and land surface conditions, and results in less accurate maps for times beyond those for which the model was trained. To cope with this domain shift, a model's parameters can be adjusted through fine-tuning on training data from the target time period, but collecting such data typically requires manual annotation of images, which is expensive and often impractical. Alternatively, the approach used to develop the model can be adjusted to improve its overall generalizability. Here we show how combining several fairly standard architectural and input techniques, including careful selection of the image normalization method, increasing the model's width, adding regularization techniques, using modern optimizers, and choosing an appropriate loss function, can significantly enhance the ability of a convolutional neural network to generalize across time, while eliminating the need to collect additional labels. A key component of this approach is the use of Monte Carlo dropout, a regularization technique applied during inference that provides a measure of model uncertainty while producing more robust predictions. We demonstrate this procedure by training an adapted U-Net, a widely used encoder-decoder architecture, with a relatively small number of labels (~5,000 224X224 image chips) collected from 3 countries on 3.7 m PlanetScope composite imagery collected primarily in 2018, and use the model, without fine-tuning, to make reliable maps of Ghana's (240,000 km^2) annual croplands for the years 2018-2023 on 4.8 m Planet basemap mosaics. We further show how this approach helps

to track agricultural dynamics by providing a country-wide overview of cropping frequency, while highlighting hotspots of cropland expansion and intensification during the 6-year time period (2018-2023).