Remote Sensing of Environment xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

Landsat-based classification in the cloud: An opportunity for a paradigm shift in land cover monitoring

G. Azzari*, D.B. Lobell

Center on Food Security and the Environment, Department of Earth System Science, Stanford University, Y2E2 Building, 473 Via Ortega, Stanford, CA 94305, United States

ARTICLE INFO

Article history:
Received 3 July 2016
Received in revised form 25 April 2017
Accepted 20 May 2017
Available online xxxx

Keywords:
Classification
Cover map
Landsat
Satellite
Big data
Google Earth Engine
Zambia
Land cover
Machine learning

ABSTRACT

Satellite-derived land cover maps play an important role in many applications, including monitoring of smallholder-dominated agricultural landscapes. New cloud-based computing platforms and satellite sensors offer opportunities for generating land cover maps designed to meet the spatial and temporal requirements of specific applications. Such maps can be a significant improvement compared to existing products, which tend to be coarser than 300 m, are often not representative of areas with fast-paced land use change, and have a fixed set of cover classes. Here, we present two approaches for land cover classification using the Landsat archive within Google Earth Engine. Random forest classification was performed with (1) season-based composites, where median values of individual bands and vegetation indices were generated from four years for each of four seasons, and (2) metric-based composites, where different quantiles were computed for the entire four-year period. These approaches were tested for six land cover types spanning over 18,000 locations in Zambia, with ground "truth" determined by visual inspection of high-resolution imagery from Google Earth. The methods were trained on 30% of these points and tested on the remaining 70%, and results were also compared with existing land cover products. Overall accuracies of about 89% were achieved for the season- and metric-based approaches for individual classes, with 93% and 94% accuracy for distinguishing cropland from non-cropland. For the latter task, the existing Globeland30 dataset based on Landsat had much lower accuracies (around 77% on average), as did existing cover maps at coarser resolutions. Overall, the results support the use of either season or metric-based classification approaches. Both produce better results than those obtained from previous classifiers, which supports a general paradigm shift away from dependence on standard static products and towards custom generation of on-demand cover maps designed to fulfill the needs of each specific application.

1. Introduction

Accurate maps of land cover are fundamental to many applications in land management and environmental monitoring. In response to the large demand for these maps, a variety of products are available, most of which rely extensively on classification of satellite data. Classification approaches attempt to use spectral differences between land cover types. In order to maximize classification accuracy, input images must have minimal contamination from clouds, haze, shadow, or other disturbances. Such images can be obtained by compositing together, according to specific criteria, large sets of observations recorded across a certain time span (Lück and van Niekerk, 2016). Because resolving fine spectral differences across similar classes using a single image can be challenging, often more than one composite is generated to characterize phenological variation, quantified either as explicit change across seasons or within the growing season, or as other annual and inter-annual temporal metrics (Zhong et al., 2011; Simonetti et al., 2015; Chang et al., 2007; Hansen et al., 2011, 2014).

Depending on the overall number of observations involved, generating such composites and running the classification requires significant data storage capacity, high computational power, and the ability to distribute non-trivial algorithms across multiple machines. Until very recently, such requirements were the prerogative of few institutions and very specialized individuals, who invested significant time and resources in generating global cover maps applicable to the widest range of applications. A few of these global products were produced during the last fifteen years (Waldner et al., 2015). Some of them became standard tools used widely by the remote sensing community and have played an invaluable role in the advancement of several fields.

* Corresponding author.
E-mail addresses: gazzari@stanford.edu (G. Azzari), dlobell@stanford.edu (D. Lobell).

http://dx.doi.org/10.1016/j.rse.2017.05.025
0034-4257/

Most global land cover products tend to have a resolution coarser than the scale of true heterogeneity in many places. For example, the Global Land Cover database for the year 2000 (GLC2000), a SPOT4-based map produced by the European Commission's Joint Research Center (JRC), has a resolution of 1/112 degree per pixel (about 1 km at the equator, larger at higher latitudes) (Fritz et al., 2003; Bartholomé and Belward, 2005). Globcover, produced by the European Space Agency with ENVISAT-MERIS data, has a 300 m resolution (Arino et al., 2007; Bicheron et al., 2008), and the Moderate-resolution Imaging Spectroradiometer (MODIS) Collection 5 Land Cover Product has a 500 m resolution (Friedl et al., 2010). The Global Land Cover Characterization database (GLCC), resulting from a concerted effort between the U.S. Geological Survey (USGS), University of Nebraska Lincoln (UNL) and the JRC, is based on the Advanced Very High Resolution Radiometer (AVHRR) and has a nominal resolution of 1 km (Loveland et al., 2010). The Global Food Security-support Analysis Data Product (GFSAD1000) has a nominal scale of 1 km (Thenkabail et al., 2012). The global Spatial Production Allocation Model (SPAM) dataset, generated using an entropy-based method that utilizes several datasets and ancillary information, has a resolution of 5 arcmin(about 10 km at the equator). The International Institute for Applied Systems Analysis-International Food Policy Research Institute (IIASA-IFPRI) cropland product, a cropland percentage hybrid map generated using some of the datasets mentioned above, has a resolution of 1 km (Fritz et al., 2015). A more complete and detailed review of available crop cover maps, which is beyond the scope of this study, can be found in Waldner et al. (2015). Although undoubtedly useful for many applications, cover maps with resolutions of 300 m or coarser are of limited use in the complex, smallholder-dominated agricultural landscapes of many developing countries (Anderson et al., 2015see also Fig. 9 for a visual comparison). Moreover, recent studies have shown how some of these datasets are often not reliable over cropland areas, as they show significant disagreement with each other and with national statistics (Waldner et al., 2015; Fritz et al., 2011; Ramankutty et al., 2008).

A much smaller number of higher resolution static cover maps based on Landsat data have been generated and made publicly available. For example, Hansen et al. (2014, 2011) used Landsat composites at various timescales to estimate temporal metrics that capture growing-season phenology and generated a CONUS land cover map. Their work, however, was mostly aimed at forestry applications and did not include a cropland class. More recently, the National Geomatic Center of China (NGCC) produced a 30 m global land cover map. Their work was in part based on classification of Landsat and China Environmental Disaster Alleviation Satellite (HJ-1) images. In addition, GNCC used a large amount of ancillary data, including pre-existing cover maps, thematic data, and topographic information (Chen et al., 2015).

Regardless of their spatial resolution, all current global land cover maps have two main limitations: 1) the lack of temporal updates, and 2) their fixed number and type of classes, which may fit well in global studies but are of limited use in finer scale applications. The recent implementation of new powerful cloud-based computational frameworks, along with the growing availability of imagery resulting from the Landsat Global Archive Consolidation (LGAC) initiative (Wulder et al., 2016), is making custom Landsat-based classification more accessible, and may help overcome the limitations of existing products. Particularly, Google Earth Engine (GEE) has emerged as an invaluable tool by offering a vast data pool of satellite imagery and access to advanced algorithms that are highly parallelized behind the scenes. In fact, over the last two years several applications have been developed that aim at some form of classification using Google Earth Engine at regional (e.g. Dong et al., 2016; Patel et al., 2015; Simonetti et al., 2015; Miettinen et al., 2016; Padarian et al., 2015) to global scales (e.g. Hansen et al., 2013 ; Pekel et al., 2016). Similarly, collecting ground points for training and validation has become easier thanks to: (i) easy access to global (and occasionally multi-temporal) high-resolution images through Google Earth (Yu and Gong, 2012), (ii) tools for easy KML creation and editing, (iii) open-access GIS data (e.g. Open Street Map), and (iv) crowd-sourcing-based data collection approaches such as Geo-Wiki.org (Fritz et al., 2009; See et al., 2015).

The objective of the current study is to assess whether these new tools are capable of generating custom cover maps that are at least as good for detecting crops as standard land cover products. To investigate this question, we use Google Earth Engine to classify agricultural landscapes in Zambia, and resolve between different natural vegetation covers, urban areas, rainfed crops and irrigated crops. In particular, we train a random forest classifier on two different types of Landsat composites: a phenological composite and a metrics composite. We then quantify the accuracy of the resulting cover map in resolving these cover types and compare results with existing products in distinguishing crop from non-crop observations. We demonstrate that both of our GEE-generated maps produce higher accuracies than existing products, while also allowing more flexibility in terms of cover class selection and reference years. This result demonstrates how frameworks such as Earth Engine, by simplifying access and processing of large amount of satellite data, are changing the paradigm in land cover monitoring from a static, product-based approach into a more dynamic and application-specific one without any loss of accuracy.

2. Materials and methods

2.1. Research area

We focused this case study in the Republic of Zambia (Fig. 1). This 752,617 km^2 region well represents the complex, smallholders-based agricultural landscape of Sub-Saharan Africa, and also contains relatively large areas of industry-driven irrigated cropland and urban settlements. According to CountrySTAT (FAO, 2016) the largest fraction of the total crop area is utilized for maize (average of 76% between 2011 and 2015), followed by groundnuts (7%), seed cotton (4%), sunflower seeds and sorghum (3%), sweet potatoes and soybeans (2%), and other minor crops covering less than 1% each. Zambia has also a diverse set of natural vegetation cover types, including forests, shrubland, savannas, grassland, and swamps (Ellenbroek, 1987).

Climatologically, Zambia provides an interesting research area, as it is characterized by three distinct seasons (Ellenbroek, 1987): 1) a warm, rainy season from November to April, 2) a cool, dry season from April to August, and 3) a hot, dry season from August to November. Sowing dates for maize are around the beginning of November, while peak season is between the end of January and mid-February. Based on the climatological and agricultural calendar, we defined four phenological seasons: S0) from November 1 to January 31, S1) from February 1 to April 30, S2) from May 1 to July 31, and S3) from August 1 to October 31. For practical reasons, we also define a "seasonal year" X as the range of dates starting from November 1 ($X-1$) and ending October 31 X. For example, seasonal year 2015 starts November 1, 2014 and ends October 31, 2015.

2.2. Landsat composites

Landsat-based composites are particularly challenging to obtain because, despite Landsat's 16-day nominal revisit time, its actual data availability is heavily limited by cloud cover, technical issues, data acquisition strategies, downlink capability, and changes in mission management the program underwent over the years (Ju and Roy, 2008; Whitcraft et al., 2015; Yu et al., 2015; Roy et al., 2010b; Wulder et al., 2016). In addition, Landsat 7 ETM+ data acquisition is