Remote Sensing of Environment 179 (2016) 210–221

Contents lists available at ScienceDirect

Remote Sensing of Environment

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

A generalized computer vision approach to mapping crop fields in heterogeneous agricultural landscapes

Stephanie R. Debats [a,*], Dee Luo [b], Lyndon D. Estes [a], Thomas J. Fuchs [c], Kelly K. Caylor [a]

[a] Department of Civil & Environmental Engineering, Princeton University, Princeton, NJ 08544, USA
[b] Department of Operations Research & Financial Engineering, Princeton University, Princeton, NJ 08544, USA
[c] NASA Jet Propulsion Laboratory, Pasadena, CA 91109, USA

ARTICLE INFO

Article history:
Received 14 September 2015
Received in revised form 27 February 2016
Accepted 10 March 2016
Available online 11 April 2016

Keywords:
Land cover
Agriculture
Sub-Saharan Africa
Computer vision
Machine learning

ABSTRACT

Smallholder farms dominate in many parts of the world, particularly Sub-Saharan Africa. These systems are characterized by small, heterogeneous, and often indistinct field patterns, requiring a specialized methodology to map agricultural land cover. Using a variety of sites in South Africa, we present a new approach to mapping agricultural fields, based on efficient extraction of a vast set of simple, highly correlated, and interdependent features, followed by a random forest classifier. We achieved similar high performance across agricultural types, including the spectrally indistinct smallholder fields as well as the more easily distinguishable commercial fields, and demonstrated the ability to generalize performance across large geographic areas. In sensitivity analyses, we determined multi-temporal information provided greater gains in performance than the addition of multi-spectral bands available in DigitalGlobe Worldview-2 imagery.

1. Introduction

Improving the capacity to monitor the spatial distribution of agriculture, particularly among smallholder farmers, is critical to increasing agricultural productivity and food security in many parts of the world. Smallholder agriculture, which dominates in Sub-Saharan Africa and other world regions, features rainfed production for household consumption and use of family labor and minimal technology (Altieri & Koohafkan, 2008; Eastwood, Lipton, & Newell, 2010; Gollin, 2014; Morton, 2007). These systems are also characterized by small, heterogeneous, and often indistinct field patterns (Estes et al., 2016; Fritz & See, 2008; Lobell, 2013; See et al., 2015). The prevalence of smallholder farms highlights the need for a specialized methodology to monitor agriculture across farming types, including both smallholder and commercial.

Of the regions where smallholder agriculture dominates, Sub-Saharan Africa is the most important, due to its geographic size and status as a potential center of agricultural growth in the coming decades. Of all farms in Sub-Saharan Africa, 80% are <2 ha and the mean farm size of 1.6 ha is significantly smaller than most world regions (Table 1) (von Braun, 2004; FAO, 1997; IFAD and UNEP, 2013; Lowder, Skoet, & Singh, 2014; Nagayets, 2005). Increasing agricultural productivity is crucial in Sub-Saharan Africa, because its population is expected to double by 2050 (Haub & Kaneda, 2013). The population remains predominantly rural despite recent urbanization (Masters et al., 2013), and 60% of the workforce is employed in agriculture (AfDB et al., 2014). As the population grows, agricultural field sizes are driven down and farmers are pushed onto marginal lands. As a result, growing climate variability, characterized by less frequent and more intense rain events, increases farmers' vulnerability to food insecurity (Davidson et al., 2003; Jayne, Chamberlin, & Headey, 2014; Masters et al., 2013; Oba, Post, & Stenseth, 2001; Thornton, Jones, Alagarswamy, & Andresen, 2009; Thornton, Jones, Ericksen, & Challinor, 2011; World Bank, 2013). As smallholder field sizes decrease, large-scale international land acquisitions and government-sponsored agricultural growth corridors promote consolidation of remaining farmland into commercial enterprises (Cotula & Vermeulen, 2009; Davis, D'Odorico, & Rulli, 2014; Gollin, 2014; Nogales, 2014; Rulli & D'Odorico, 2014; Rulli, Saviori, & D'Odorico, 2013). Efforts to monitor agricultural change on the ground are confounded by a widespread shortage of agricultural data in Sub-Saharan Africa, in part due to limited government capacity (Alliance for a Green Revolution in Africa (AGRA), 2013; Carletto, Jolliffe, & Banerjee, 2013; Glassman et al., 2014; IFAD and UNEP, 2013; World Bank, 2013). Therefore, an accurate accounting of agricultural land cover, across both smallholder and commercial farming, is needed to track and promote food security in Sub-Saharan Africa.

In data-sparse regions, satellite imagery provides alternative means to monitor agriculture. However, land cover data sets derived from remote sensing contain large uncertainties regarding the total area as well as the spatial distribution of agriculture (Estes et al., 2016; Fritz &

* Corresponding author.
E-mail address: sdebats@princeton.edu (S.R. Debats).

http://dx.doi.org/10.1016/j.rse.2016.03.010
0034-4257/

Table 1
Estimated farm sizes of selected world regions (von Braun, 2004; FAO, 1997; Nagayets, 2005).

World region	Average farm size (hectares)
Africa	1.6
Asia	1.6
Western Europe	27.0
Latin America and Caribbean	67.0
North America	121.0

See, 2008; Fritz et al., 2011; Siebert, Portmann, & Döll, 2010). Moreover, readily available data sets, such as MODIS (250 m) and Landsat (30 m), lack sufficient spatial resolution to study smallholder fields, which are often smaller than 1 ha (100 m × 100 m) (Estes et al., 2016; Jain, Mondal, DeFries, Small, & Galford, 2013; Lobell, 2013). For example, a previous study using Landsat imagery to identify agricultural fields in the United States struggled with small, irregular fields of <1.5 ha, though these conditions represented only a minority of studied fields (Yan & Roy, 2014).

High-resolution satellite imagery (≤2 m) provides the necessary detail to observe smallholder agriculture and is becoming increasingly available and affordable (Estes et al., 2016; Fritz et al., 2015; Hayes, Miller, & Murphy, 2014; Jain et al., 2013; Lobell, 2013; See et al., 2015). As coverage improves, automated classification algorithms are used to extract actionable data from images. Prior efforts to automatically classify agricultural fields using high-resolution imagery have struggled due to (1) the nature and appearance of smallholder agriculture, (2) the properties of high-resolution imagery itself, and (3) the design of classification algorithms.

First, the high spatial variability in land cover poses a classification challenge, particularly among smallholder agricultural fields, which are small and irregularly shaped (Palm et al., 2010). Smallholder fields are less visually defined, exhibiting indistinct boundaries between neighboring fields as well as ambiguity between fields and natural vegetation (Estes et al., 2016; Fritz & See, 2008; Lobell, 2013; See et al., 2015). They also exhibit more variability in spectral and phenological signatures, due to rainfed farming, sub-optimal management, low cropping intensity, fallowing, abandonment, and the inclusion of large trees within fields (Jain et al., 2013; Mayes, Mustard, & Melillo, 2015; Siebert et al., 2010; Vintrou et al., 2012) (Fig. 1). Thus, these conditions highlight the need for a specialized methodology for smallholder agriculture.

Second, the level of detail increases in high-resolution imagery, which raw spectral values or simple features have difficulty describing. In high-resolution imagery, land cover classes have lower inter-class and higher intra-class spectral variability, creating ambiguities in classification (Lu & Weng, 2007; Tokarczyk, Wegner, Walk, & Schindler, 2013, 2015). Variability also increases with image mosaicking, which is often necessary due to lower collection frequency of high-resolution imagery (Estes et al., 2016; Hayes et al., 2014). Previous efforts to handle increased detail in high-resolution imagery have focused on expanding the feature space by manually handcrafting higher order features suitable to a specific application. Features that capture textural and contextual information (e.g. grey-level co-occurrence matrices (GLCM), filter bank responses, and textons) have been found to improve classification accuracy over spectral information alone, but their use has been limited by their complexity and high computational cost (Butusov, 2003; Kurosu, Yokoyama, Fujita, & Chiba, 2001; Leung & Malik, 2001; Lu & Weng, 2007; Podest & Saatchi, 2002; Rao et al., 2002; Schmid, 2001; Shaban & Dikshit, 2001; Shotton, Johnson, & Cipolla, 2008; Tokarczyk et al., 2013, 2015). Incorporating additional multi-spectral bands in features, either as raw inputs or calculated image transformations like vegetation indices, may also increase accuracy, but without selection of the most discriminative bands, this approach is limited by high band correlation (Lu & Weng, 2007; Mausel, Kramber, & Lee, 1990; Thenkabail, Enclona, Ashton, Legg, & De Dieu, 2004). Multi-temporal imagery has also been found to increase accuracy, especially for agriculture, by capturing details about phenological profiles and filling in missing data (Duveiller & Defourny, 2010; Guerschman, Paruelo, Bella, Giallorenzi, & Pacin, 2003; Liu, Takamura, Takeuchi, & Shao, 2002; Lu & Weng, 2007; Oetter, Cohen, Berterretche, Maiersperger, & Kennedy, 2001). Overall, extracting more features derives greater information from high-resolution imagery. However, classification algorithms struggle with the expanded feature space, necessitating the use of feature selection to determine the most useful ones, often in a separate stage prior to classification (Hughes, 1968; Lu & Weng, 2007; Price, Guo, & Stiles, 2002).

Third, a classification algorithm must efficiently handle the vast feature spaces of highly correlated and interdependent features required to adequately describe smallholder agriculture in high-resolution imagery (Hughes, 1968; Lu & Weng, 2007; Price et al., 2002). An algorithm must also resist overfitting to training data and have high generalization performance to classify large expanses of new imagery (Mascaro et al., 2014). Previous efforts have found that tracking land cover changes over time lends itself to a supervised classification approach, in which a classifier is initially trained to identify prescribed classes with labelled data and then repeatedly deployed on a time series of images (McIver & Friedl, 2001). Yet supervised classification is often deemed a local approach for small areas, with a reputation of being difficult to repeat over large areas (Mayes et al., 2015). Within supervised classification, non-parametric classifiers are increasingly preferred, due to their ability to handle extremely large feature spaces and data sets, robustness to outliers and noise, and outperformance of parametric classifiers in complex landscapes (Adam, Mutanga, Odindi, & Abdel-Rahman, 2014; De Fries, Hansen, Townshend, & Sohlberg, 1998; Friedl, Brodley, & Strahler, 1999; Gislason, Benediktsson, & Sveinsson, 2006; Gopal, Woodcock, & Strahler, 1999; Ham, Chen, Crawford, & Ghosh, 2005; Hayes et al., 2014; Lu & Weng, 2007; McIver & Friedl, 2001; Paola & Schowengerdt, 1995; Rodriguez-Galiano et al., 2011). Furthermore, classifiers with probabilistic output, as opposed to hard classifications, are gaining popularity in remote sensing for highlighting the spatial variation in classification quality and confidence, which is crucial for utilizing results in decision making (Liu, Gopal, & Woodcock, 2004; Lu & Weng, 2007; McIver & Friedl, 2001). Probabilistic output can also be post-processed in a variety of ways, ranging from simple thresholding for obtaining hard classifications to image segmentation and object detection.

This study develops a methodology to differentiate heterogeneous agricultural land cover in high-resolution imagery of Sub-Saharan Africa that is effective across a range of agricultural types, including the small, irregular fields of the dominant smallholder class. In a supervised classification approach, we utilize techniques from computer vision and machine learning, two fields that are increasingly used in remote sensing analyses as high-resolution imagery becomes readily available. Computer vision provides the means to detect, recognize, and track complex objects in images, whereas machine learning enables assimilation of vast quantities of data (Szeliski, 2010), a powerful combination for deriving actionable data from high-resolution imagery that generalizes across large areas.

First, to address the problems of adequately describing the small, irregular fields of smallholder agriculture and handling the great level of detail in high-resolution imagery, we efficiently extract a unique set of many simple, highly correlated, and interdependent features. Second, we utilize a supervised, non-parametric machine learning classifier to handle the vast feature space and provide probabilistic output of the likelihood that pixels belong to agricultural fields. We assess our classification accuracy using the metrics of receiver operating characteristic (ROC) curves and the related area under the ROC curve (AUC), finding similar high performance (AUC > 0.90) across agriculture types, including smallholders. We also investigate the relative roles of multi-temporal and multi-spectral information, given the data limitations of high-resolution satellite imagery in Sub-Saharan Africa, and find that multi-temporal imagery contributes to greater performance gains.