Remote Sensing of Environment xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Remote Sensing of Environment

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

Towards fine resolution global maps of crop yields: Testing multiple methods and satellites in three countries

George Azzari *, Meha Jain, David B. Lobell

Department of Earth System Science, Center on Food Security and the Environment, Stanford University, Stanford, CA 94305, USA

ARTICLE INFO

Article history:
Received 3 July 2016
Received in revised form 31 March 2017
Accepted 12 April 2017
Available online xxxx

Keywords:
Agriculture
Landsat
MODIS
Google Earth Engine
Crop modeling
Big data
Remote sensing
Yield

ABSTRACT

One of the greatest challenges in monitoring food security is to provide reliable crop yield information that is temporally consistent and spatially scalable. An ideal yield dataset would not only extend globally and across multiple years, but would also have enough spatial granularity to characterize productivity at the field and sub-field level. Rapid increases in satellite data acquisition and platforms such as Google Earth Engine that can efficiently access and process vast archives of new and historical data offer an opportunity to map yields globally, but require efficient and robust algorithms to combine various data streams into yield estimates. We recently introduced a Scalable satellite-based Crop Yield Mapper (SCYM) that combines crop models simulations with imagery and weather data to generate 30 m resolution yield estimates without the need for ground calibration. In this study, we tested new large-scale implementations of SCYM, focusing on three regions with varying crops, field sizes and landscape heterogeneity: maize in the U.S. corn belt (390,000 km^2), maize in Southern Zambia (86,000 km^2), and wheat in northern India (450,000 km^2). As a benchmark, we also tested a simpler empirical approach (PEAKVI) that relates yield to the peak value of a time series of spatially aggregated vegetation indices, similar to methods used in current operational monitoring. Both SCYM and PEAKVI were applied to data from all Landsat's sensors and MODIS for more than a decade in each region, and evaluated against ground-based estimates at the finest available administrative level (e.g., counties in the U.S.). We found consistently high correlations ($R^2 \geq 0.5$) between the spatial pattern of ground- and satellite-based estimates in both U.S. maize and India wheat, with small differences between methods and source of satellite data. In the U.S., SCYM outperformed PEAKVI in tracking temporal yield variations, likely owing to its explicit consideration of weather. In India, both methods failed to track temporal yield changes, with various possible explanations discussed. In Zambia, the PEAKVI approach applied to MODIS tracked yield variations much better ($R^2 > 0.5$) than any other yield estimate, likely because the frequent cloud cover in this region confounds the other approaches. Overall, this study demonstrates successful approaches to yield estimation in each region, and illustrates the importance of distinguishing between accuracy for spatial and temporal variation. The 30 m resolution of Landsat-based SCYM does not appear to offer large benefits for tracking aggregate yields, but enables finer scale analyses than possible with the other approaches.

1. Introduction

The productivity of major crops is a key characteristic of any agricultural region, as evidenced by the substantial resources that governments and others devote to collecting information on productivity each year. Crop yield, defined as the ratio of total mass of harvested product (e.g., grain) to cropped area, is one of the most basic and widely sought measures of productivity (Carletto et al., 2015). The utility of yield data depends on the timeliness and spatial scale of estimates relative to the needs of a particular application. For example, a common desire is to have in-season forecasts of pending yields over large geographic regions, such as states or entire countries, in order to inform decisions about trade or potential government assistance. At the same time, yield data for individual fields is useful for a wide range of questions concerning field management and crop yield gaps (Lobell, 2013), even if the data are not available until after harvest.

Satellite remote sensing offers promise as a tool to assess crop yields, with many studies demonstrating high correlations between satellite-based estimates and traditional sources in specific case studies (e.g., Clevers, 1997; Shanahan et al., 2001). Although satellites still play a limited role in most operational efforts to monitor yields, several recent developments have enabled progress towards more routine use of satellites for yield assessment.

First, access to satellite data has become significantly easier in recent years. This trend results from three main driving factors: a) additional

* Corresponding author.
E-mail address: gazzari@stanford.edu (G. Azzari).

http://dx.doi.org/10.1016/j.rse.2017.04.014

satellites have been launched into orbit, including public satellites such as Landsat 8 (Roy et al., 2014) and the Sentinel constellation (Drusch et al., 2012), and private satellites (Belward and Skøien, 2015; Hand, 2015), b) all archives of satellite images from public institutions, such as NASA and ESA, and some private archives have been made available for free to the public, and c) new cloud-based platforms such as Google Earth Engine have made it easier to access large volumes of data without the need to download them onto local computers. Second, access to the computational power needed to process large volumes of satellite data has also dramatically increased, as the data processing platforms utilize parallel computing resources well beyond the capacity of most individual research groups.

Third, new algorithms have been developed that provide a more generalizable way to estimate crop yields from satellite measurements. Whereas earlier approaches tended to be specific to a given crop or region, often with year and site-specific calibration, newer algorithms offer a more scalable approach. In this paper we focus on two approaches. The first uses fine temporal resolution, coarse spatial resolution sensors such as the Moderate Resolution Imaging Spectroradiometer (MODIS) to track variation in vegetation indices (VIs) such as the normalized difference vegetation index (NDVI) over cropped pixels, and then relates the peak value of spatially averaged NDVI to crop yields (Becker-Reshef et al., 2010; Franch et al., 2015). This approach (referred to hereafter as "PEAKVI") has been developed as part of the Group on Earth Observations Global Agricultural Monitoring Program (GEOGLAM; http://geoglam-crop-monitor.org/), which provides in-season crop condition ratings for major grain producing regions around the world. The PEAKVI approach is used as one of the lines of evidence in the crop monitoring effort, often along with in situ data sources reported by local experts. In some measure, the PEAKVI approach is similar to other methods that predict yields using VIs on specific dates near the middle of the growing season (e.g. Johnson, 2014, who uses NDVI as well as land surface temperature data).

A second approach, the scalable crop yield mapper (hereafter "SCYM"), was developed to use coarser temporal resolution but finer spatial resolution Landsat data to estimate yields at 30 m resolution (Lobell et al., 2015). In order to bypass the need for ground calibration data, SCYM employs simulations from crop models to train a regression that relates final crop yield to observed values of vegetation indices for available images during the growing season.

The PEAKVI and SCYM approaches have both been tested in a few locations (Becker-Reshef et al., 2010; Farmaha et al., 2016; Franch et al., 2015; Lobell et al., 2015), and both offer promise for more widespread application. Although PEAKVI has been developed with MODIS imagery, and SCYM with Landsat imagery, both methods are in principle applicable to any source of data. A relative strength of the PEAKVI approach is its simplicity, as it relies solely on a VI dataset and crop mask. In contrast, the use of simulation models in SCYM provides the ability to use temporally sparse observations and incorporate weather effects not captured in the VI data. In addition, SCYM provides estimates at the native resolution of the imagery (e.g., 30 m for Landsat), whereas PEAKVI provides estimates only at aggregated scales (e.g., for counties or states).

With the goal of advancing prospects for routine yield assessment from satellites across large areas, this study tests the two methods for multiple years in three different cases: maize in the United States, wheat in India, and maize in Zambia. These systems provide a contrast in field size, crop type, management intensity, landscape heterogeneity, cloud cover, and other factors that could affect the performance of each method. Given a lack of field-scale yield measurements, we focus the analysis on the ability of each method to track spatial variation at the finest aggregation level for publicly available data, and temporal variation across the study time period (14 years in U.S., 13 years in India and Zambia). We also test each method using both Landsat and MODIS data, which provide a contrast in temporal and spatial resolution. In summary, the three primary goals of the study were to (i) quantify the ability of easily scalable satellite-based yield estimates to track variations in yields as reported in ground-based datasets across different regions, (ii) compare performance for a method that utilizes crop model simulations to produce estimates at native resolution of the imagery (SCYM) with a simpler benchmark that uses spatial aggregates of VI (PEAKVI), and (iii) compare performance when using Landsat vs. MODIS.

We utilized the Google Earth Engine platform throughout the study, in part because addressing these questions required analysis of large volumes of data – more than a decade of Landsat and MODIS data for significant fractions of each country. Specifically, the total area covered by the study was roughly 1 million km^2 (see below). An individual Landsat 8 tile covers roughly 30,000 km^2, and is 2 GB in size (http://landsat.usgs.gov/landsat8.php). If each location is imaged an average of 10 times by Landsat during each year for 14 years, this translates to over 10 TB of Landsat data alone. Incorporating MODIS data, weather grids, and crop masks adds substantially to this total. In addition to the data volumes, our study also required an ability to perform several operations such as cloud masking and image compositing at the pixel level, while also dealing with reprojection and resampling issues on the fly. In short, dealing with "big remotely sensed data" requires an ability to seamlessly extract significant, higher-level information out of vast amount of raw data without heavy user interaction, and Google Earth Engine provided that capability for this study.

2. Methods

2.1. Study regions

We investigated three distinct agricultural regions, each corresponding to a key growing region within their respective countries. For the U.S., we focus on maize in the states of Iowa, Illinois, and Indiana (Fig. 5). These are three leading producers of maize in the country, comprising well over one-third of total national production (NASS, 2015). Although the methods employed in this study are easily extended to other states, we restricted the comparisons in this study to these three states since they are the only Corn Belt states for which USDA Cropland Data Layer (CDL) maps used to identify maize fields exist back to 2001 (see below). The vast majority of maize in these three states are grown without irrigation.

Wheat is one of the main two grain crops grown in India, the other being rice, and it is important both nationally and globally for food security. We decided to focus on wheat instead of rice for several reasons. First, there is sufficient subnational data to test the performance of our estimates. Second, there are reliable crop model simulations for wheat in this region. Third, rice is primarily grown during the monsoon season in India, when cloud cover is very high and it may have been difficult to obtain high quality imagery throughout the monsoon season for each year considered in our study. Wheat, on the other hand, is grown during the dry winter season, when there is more image availability. Finally, we think that demonstrating an approach on more than one crop is helpful for demonstrating its generality. In India, we focus on the four main wheat-producing states of the Indo-Gangetic Plains: Punjab, Haryana, Uttar Pradesh, and Bihar (Fig. 8). These four states contribute well over 70% of national production of wheat, which is second only to rice as the main staple crop in India. Most wheat fields in this region are irrigated, although some fields in Uttar Pradesh and Bihar receive limited amounts of water. Field sizes are an order of magnitude smaller than for U.S. maize (~2 ha vs. ~20 ha), with field size generally decreasing from west to east across the wheat-growing region.

In Zambia, we focus on maize in the Southern Province, one of the main maize growing regions within the country. Maize is the primary staple in Zambia, and is grown under rainfed conditions on smallholder fields, with 42% of farmers cultivating fields below 1 ha, and 97% of fields below 5 ha (Mason et al., 2011).

Overall, the study covers nearly 1 million km^2 of land area – ~ 390,000 km^2 in the U.S., ~450,000 km^2 in India, and ~86,000 km^2 in

Please cite this article as: Azzari, G., et al., Towards fine resolution global maps of crop yields: Testing multiple methods and satellites in three countries, Remote Sensing of Environment (2017), http://dx.doi.org/10.1016/j.rse.2017.04.014