Remote Sensing of Environment 216 (2018) 183–200

Contents lists available at ScienceDirect

Remote Sensing of Environment

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

Multi-decade, multi-sensor time-series modelling—based on geostatistical concepts—to predict broad groups of crops

Matthew J. Pringle[a,b,*], Michael Schmidt[a,b], Daniel R. Tindall[a,b]

[a] Department of Environment and Science, GPO Box 2454, Brisbane, Queensland 4001, Australia
[b] Joint Remote Sensing Research Program, School of Earth and Environmental Sciences, University of Queensland, St. Lucia, QLD 4072, Australia

ARTICLE INFO

Keywords:
Phenology
Landsat
Sentinel-2
Variogram
Kriging
GEOBIA

ABSTRACT

We have mapped the broad groups of crops grown each summer and winter, from 1987 to 2017, for a 300,000-km^2 region of Queensland, Australia. These maps are part of a legislated decision-making process for the protection of prime agricultural land. For summer, the two groups of crops are 'Coarse-grain & Pulse' and 'Cotton'. For winter, the two groups of crops are 'Cereal' and 'Pulse'. Non-crop groups, present in both summer and winter, are 'Bare soil' and 'Other' (comprising pastures, woody vegetation, and crop residues). The foundation of the maps is time-series modelling—specifically, applying the concepts of geostatistics in the temporal domain—to model the variation in land-surface phenology within a growing season. The time-series model is flexible, robust, parsimonious, parallelisable, and able to deal with irregular observations. We combined satellite imagery from the Landsat sensors, as well as, when available, Sentinel-2A and MODIS (with the last two reprojected to the 30-m grid of Landsat). We applied the time-series model pixel-wise across the study region, to three variables derived from satellite imagery gathered for an individual growing season: enhanced vegetation index, and the sub-pixel proportions of bare-ground and non-photosynthetic vegetation. Weekly-averaged predicted phenological metrics then served as explanatory variables in a tiered, tree-based classification model, for the prediction of the groups. The classification model comprised two expert rules and two random forests. Prior to fitting the classification model, geospatial object-based image analysis was used to change the scale of analysis from individual pixels to (approximately) field-based segments. From the perspective of a map-user, in any given growing season we predicted 'Coarse-grain & Pulse' correctly in 79% of cases; the values for 'Cotton', 'Cereal', and 'Pulse' were 90%, 84%, and 73%, respectively; 'Bare soil' was 72% in summer, and 88% in winter. 'Other' was the most accurately mapped group (98% correct in summer, and 99% correct in winter).

1. Introduction

With Earth's human population currently increasing by approximately 83-million per year (UNDESA, 2015), food-supply and the management of limited terrestrial resources are critical issues for the coming decades (Godfray et al., 2010). As part of its Sustainable Development Goals, the United Nations wishes to 'end hunger, achieve food security and improved nutrition and promote sustainable agriculture' by 2030 (United Nations, 2015). To achieve this goal, it will become increasingly important for policy-makers to know where and when particular crops are grown. This knowledge will be able to inform diverse subjects such as: mechanistic models of water quality (Carroll et al., 2012); debate about competing land-uses (Lambin and Meyfroidt, 2011); understanding the potential for biofuel production (Beringer et al., 2011); the logistics of commodity transportation (Gurning and Cahoon, 2011); and the management of financial risk (Bokusheva et al., 2016).

It has long been recognised that the spectral characteristics of vegetation, particularly actively growing crops, can be detected by Earth-observing satellites. For example, Tucker and Sellers (1986) promoted the analysis of multiple satellite images per growing season, to reliably characterise the temporal evolution of primary production. Early work with optical satellites in the 1980s and 1990s was often limited by the 1.1-km spatial resolution of the daily images scanned by Advanced Very High Resolution Radiometer (AVHRR) instruments (Cracknell, 1997), or by the cost of obtaining cloud-free Landsat images (30-m spatial resolution, 16-day temporal resolution) (Wulder et al., 2012).

In the 2000s two events broke the constraints to operational crop-mapping, using data from Earth-observing satellites. First, the Moderate Resolution Imaging Spectroradiometer (MODIS) satellites were launched. These deliver free imagery that is useful for vegetation monitoring, and do so at a finer spatial resolution (250–500 m) than

* Corresponding author at: Department of Environment and Science, GPO Box 2454, Brisbane, Queensland 4001, Australia.
E-mail address: matthew.pringle@qld.gov.au (M.J. Pringle).

https://doi.org/10.1016/j.rse.2018.06.046
Received 30 November 2017; Received in revised form 26 June 2018; Accepted 27 June 2018

AVHRR. To obtain cloud-free data, contemporary studies would apply a standardising atmospheric correction to AVHRR or MODIS imagery, followed by compositing over a particular time interval (Van Leeuwen et al., 1999). The resulting stack of composite images enabled large-scale crop-mapping studies (Becker-Reshef et al., 2010; Estel et al., 2016). The second event was the opening of the vast Landsat archive by the United States Geological Survey in 2008 (Wulder et al., 2012), which enabled research at a spatial resolution more compatible with crop management, delving as far back as the 1970s. The recently launched Sentinel-2 satellites, which are spectrally similar to Landsat, will be increasingly used for vegetation monitoring. One of the challenges for the remote-sensing community is to combine multi-sensor imagery into a single coherent product, to create what Wulder et al. (2015) termed a 'virtual constellation'. For example, a method exists to adjust MODIS reflectance so that it estimates Landsat reflectance (Gao et al., 2006), while Flood (2017) demonstrated how Sentinel-2A reflectance can be adjusted to estimate that of Landsat.

Lhermitte et al. (2011) proposed that the temporal variation in a stack of satellite imagery—or some derivation thereof, such as a vegetation index—can be summarised in three ways: (i) distance measures on the original data (e.g. correlation coefficients); (ii) transformations of the original data to reduce dimensionality (e.g. principal component analysis); or (iii) metric-based approaches, which summarise temporal variation as a set of parameters. Missing data, which can be common in a stack of satellite imagery, will create problems for (i) and (ii), although Yan and Roy (2015) presented a promising solution. Applications based on (iii)—e.g. the HANTS (Verhoef et al., 1996) and TIMESAT (Jönsson and Eklundh, 2004) algorithms—have been tailored to deal with missing data, while simultaneously enabling a link between time-series modelling and land-surface phenology. Further research into this kind of time-series modelling is justified, because the most appropriate model may be context-specific (Atkinson et al., 2012). We contend that an appropriate time-series model for crop-mapping will be flexible, robust, parsimonious, parallelisable, and hold no assumption of regular observations (Table 1).

Geostatistical concepts, known for their spatial applications (Van der Meer, 2012), offer an alternative to HANTS and TIMESAT for time-series modelling of the data from Earth-observing satellites. When applied pixel by pixel within an image stack, a geostatistics-based time-series model can satisfy all five aspects of Table 1. The fundamentals of geostatistics (Matheron, 1963)—autocorrelation, variograms, and kriging—are methods familiar enough to have been thoroughly described in textbooks, e.g. Webster and Oliver (2001). Some ideas for applying geostatistics temporally, in the context of phenology, were explored by Pringle (2013).

As well as summarising temporal variation, there are two further important considerations for crop-mapping. The first is the classification model used to predict the location of actively growing crops. The classification model takes as input the phenological metrics and returns a prediction of the most-probable group of vegetative cover (if any) at a particular time and place. An interesting recent trend is the use of tiered models for classification. For example, in the model of Massey et al. (2017), a first tier distinguished areas of 'Fallow' from 'Cropland', a second tier split 'Cropland' into either 'Alfalfa' or 'Cotton, Rice', and a third tier split 'Cotton, Rice' into its constituents. Similar ideas have been presented for crop-mapping by Bellón et al. (2017) and Lebourgeois et al. (2017). An advantage of this structure is that the top tiers might take the form of simple expert-elicited rules that filter easy-to-identify cases, leaving an automated classifier to work on the subset of more-challenging patterns. One of the pit-falls of a tiered model is that the outcomes of a particular tier must be conditional on the preceding tier(s).

The other consideration is the spatial scale at which we create and apply the classification model. Landholders conduct cropping activities within homogeneous management units, i.e. at the spatial scale of a field. It therefore makes sense for the classification model to be fitted and applied at field-scale, rather than pixel-scale. During the last decade there has been an increasing use of geospatial object-based image analysis (GEOBIA; Blaschke, 2010) as a means of changing the scale of spatial data; it has the added advantage of reducing data volumes. In the context of crop-mapping, GEOBIA represents a way to allocate a group of contiguous pixels into an approximation of a field. Notable recent crop-mapping studies to have used GEOBIA are Peña-Barragán et al. (2011), Maxwell and Sylvester (2012), and Schmidt et al. (2016).

The above discussion elucidates some of our tenets for crop-mapping: combining observations from different sensors; appropriate time-series modelling; and using a tiered classification model to predict the occurrence of different groups of crops at the spatial scale of a field (as distinct from pixel-scale). To be ultimately useful for government policy, these tenets need to be realised within an operational framework, i.e. be applicable in an on-going way, over a large area, with known accuracy. Operational crop-mapping has been successfully demonstrated with MODIS imagery (Becker-Reshef et al., 2010; Estel et al., 2016; Massey et al., 2017), but this is more difficult with Landsat, due to the vastly increased data loads, and the complex landscape variability these satellites reveal. While the annual delivery of the Cropland Data Layer in the United States is an exemplar—a product that combines imagery from a variety of sources, at a Landsat-like spatial resolution (Johnson and Mueller, 2010)—many Landsat-based crop-mapping studies have ignored operational considerations. There are studies on relatively small areas over few years (Vieira et al., 2012; Li et al., 2015; Wang et al., 2017), studies on large areas over few years (Matton et al., 2015; Song et al., 2017), or small areas over many years (Maxwell and Sylvester, 2012). Our earlier Landsat-based attempt at operational crop-mapping (Schmidt et al., 2016) created field-scale predictions over a large area and over many years, but was limited by the binary 'Crop' versus 'Other' classification model.

The principal aim of this study was to produce biannual maps of broad groups of crops (Table 2), over a 30-year period—maps that can ultimately serve as instruments of government policy. To achieve this, we posited four minor aims: (1) to summarise growing-season

Table 1
Desirable aspects for a time-series model that characterises land-surface phenology.

Aspect	Comment
Flexible	There are potentially billions of different cases, not just the bell-shaped fluctuations of a crop, but also, conceivably, negligible variation.
Robust	We must minimise the influence of outliers, which arise due to factors independent of the land-surface (e.g. imperfectly masked cloud-shadows or topographic effects).
Parsimonious	Each parameter of the model must be justified.
Parallelisable	The model must be fitted within a reasonable amount of computing time.
Irregular	The model must hold no assumption of regular observations.

Table 2
The groups (crop and non-crop) to be mapped; the cropping-phase when they occur ('S' = Summer, 'W' = Winter); prior probability of occurrence; and their major constituents (NA = not applicable).

Group	Phase	Prior probability	Constituents
Coarse-grain & Pulse crops	S	0.072	Sorghum, maize, mungbean, soybean
Cotton crop	S	0.018	Cotton
Cereal crop	W	0.077	Wheat, barley, oats
Pulse crop	W	0.013	Chickpea
Bare soil	S,W	0.030	NA
Other	S,W	0.880	Pastures, woody vegetation, crop residues