Exploiting the multi-angularity of the MODIS temporal signal to identify spatially homogeneous vegetation cover: A demonstration for agricultural monitoring applications

Grégory Duveiller, Raul Lopez-Lozano, Alessandro Cescatti

Remote Sensing of Environment 166 (2015) 61–77

Contents lists available at ScienceDirect

Remote Sensing of Environment

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

Exploiting the multi-angularity of the MODIS temporal signal to identify spatially homogeneous vegetation cover: A demonstration for agricultural monitoring applications

Gregory Duveiller *, Raul Lopez-Lozano, Alessandro Cescatti

European Commission Joint Research Centre, Via Enrico Fermi 2749, I-21027 Ispra (VA), Italy

ARTICLE INFO

Article history:
Received 16 September 2014
Received in revised form 24 April 2015
Accepted 1 June 2015
Available online 12 June 2015

Keywords:
MODIS
Pixel purity
Signal-to-noise ratio (SNR)
Crop monitoring
Homogeneous vegetation cover

ABSTRACT

MODIS has been providing daily imagery for retrieving land surface properties with a spatial resolution 250 m since the year 2000. In many places, this pixel size is closer to that of individual landscape elements, such as managed forest stands or crop fields, than those of time series more typically used in vegetation analyses (with pixel sizes ranging from about 1 km to 8 km). With such spatial resolution, combined with its increasingly long archive, MODIS data offers great potential for vegetation monitoring applications in general, and crop growth monitoring in particular. However, due to its whiskbroom design, the observation geometry of the MODIS instrument, combined with the spatial uncertainty in the registration of the images, can result in different (albeit neighbouring) physical areas being mapped onto the same pixel depending on the view zenith angle (which varies from day to the next). Rather than considering this as an inconvenience, a method is here proposed to exploit this peculiarity to identify pixels corresponding to a homogeneous plant cover, in order to retrieve surface specific time series of satellite products. This method is based on quantifying the temporal signal-to-noise ratio (SNR, hereafter) of the daily MODIS NDVI time series, defined as the variance of smoothed temporal signal over the variance of the residues. If consecutive observations of the same pixel (which have thus sampled the spatial vicinity of that pixel) provide similar NDVI values, the resulting temporal signal is relatively stable and the SNR is high. Such stability can indicate that the signal comes from a spatially homogeneous surface, such as a single large field covered by the same crop with similar agro-management practices. On the contrary, a noisy time series indicates instead a transition zone between different land uses or between fields with different management practices. SNR maps therefore serve as a proxy for sub-pixel homogeneity from which time series originating from a specific land cover or land use can be retrieved. This approach is demonstrated over 12 contrasting agricultural landscapes across the globe from which clearly distinctive crop specific signals are extracted. Exploiting the full MODIS archive to derive surface specific information in this way should open new avenues for regional to global agricultural monitoring applications. Expanding this method to derive satellite products for specific land cover classes could also be useful for many other applications linked to dynamics of land cover and land use change.

1. Introduction

The synoptic measuring capacities of satellite remote sensing provide a fundamental tool to understand biogeochemical and biophysical land surface processes at local to global scales. Land surface properties such as leaf area index (Myneni et al., 2002; Baret et al., 2013), land surface temperature (Wan, 2008) or evapotranspiration (Mu, Zhao, & Running, 2011) have been derived at a global scale from remote sensing measurements of electromagnetic radiation. Beyond their value as monitoring tools by themselves, the biophysical variables obtained from satellite also serve to calibrate, validate or parametrise models, such as land surface models and crop growth models, which are needed to tackle challenges such as climate change and food security.

Many applications require a combination of high revisit frequency measurements over large spatial extents and with a long temporal data record. Satellite systems providing such datasets have a spatial resolution which is often coarser than the targeted landscape elements (such as managed forest stands or crop specific fields). Moreover, the actual footprint of the satellite observation is generally larger than the rectangular ground projection of the pixel (Cracknell, 1998). Uncertainty in the geometric correction of the imagery combined with gridding artefacts (Tan et al., 2006; Gómez-Chova et al., 2011) further complicates the problem of matching the observations to the targets. Since landscapes are often spatially heterogeneous, the adequacy of observation and target will vary in space, leading to distinctions

* Corresponding author.
E-mail address: gregory.duveiller@jrc.ec.europa.eu (G. Duveiller).

http://dx.doi.org/10.1016/j.rse.2015.06.001

between what is common to refer to as "mixed" and "pure" pixels. This spatial heterogeneity must be taken into account explicitly in algorithms for biophysical variable retrievals because of the non-linear nature of the physical relationships linking what is observed from space to land surface properties (Chen, 1999; Garrigues, Allard, Baret, & Weiss, 2006; Jacob & Weiss, 2014). It can also be added that the spatial homogeneity of a given land cover depends on the scale at which it is considered: a land cover class can be regarded as homogeneous in its heterogeneity, or be composed of a mosaic of different land cover classes.

A method capable of *a priori* quantifying pixel purity with respect to the target landscape element can have many uses in Earth Observation sciences. Remote sensing products often rely on existing networks of ground measurement sites (*e.g.* Fluxnet, Aeronet, *etc.*) for validation. Pixel purity could be used to assess the representativity of ground stations and measurements in terms of spatial homogeneity with respect to land use or plant functional type in order to better compare with satellite observations (*e.g.* Cescatti et al., 2012). Pixel purity could further be used to design new sampling schemes, or complete existing ones, by organizing measurements or setting up flux towers at pre-defined areas that are more adequate to compare with coarse satellite products. Other applications include: (1) extracting pure time series for specific plant functional types to calibrate, validate or parametrise dynamic global vegetation models (DGVMs), which typically run at a much coarser scale; (2) assisting the incorporation of mixed pixels in various stages of supervised classifications (*e.g.* Foody, 1996); (3) assisting and validating unmixing approaches to infer subpixel land cover specific biophysical properties (*e.g.* Kuusinen, Tomppo, & Berninger, 2013); (4) monitoring vegetation phenology (*e.g.* Zhang et al., 2003) while avoiding mixed pixels; (5) constructing global databases of crop calendars (*e.g.* Whitcraft, Becker-Reshef, & Justice, 2014) based on pure crop specific pixels; or (6) to use such crop specific pixels to monitor crop growth at a regional scale (*e.g.* Duveiller, Baret, & Defourny, 2012).

Assessing pixel purity, or characterising landscape spatial heterogeneity in general, typically requires fine (decametric) spatial resolution data or vectorial maps. The proportion of surface elements falling within coarse (hectometric or kilometric) pixels can be calculated from their representation at finer scales, ideally taking into account the spatial response, or point spread function (PSF), of the instrument (Duveiller, Baret, & Defourny, 2011). Landscape ecologists have developed a large array of metrics to characterise landscape fragmentation from finer categorical map patterns (McGarigal, Cushman, & Ene, 2012). Other methods designed to capture spatial heterogeneity, based on local moving windows (Woodcock & Strahler, 1987) or variograms (Garrigues et al., 2006), also require concurrent fine spatial resolution imagery. Yet, handling decametric spatial resolution imagery at global scale is a daunting task, and observations are not available with high temporal frequency for the past. This has led some effort to derive information on spatial heterogeneity directly from coarse spatial resolution by exploiting the wealth of information that such data has in the temporal (de Bie, Nguyen, Ali, Scarrott, & Skidmore, 2012; Ali, de Bie, Skidmore, Scarrott, & Lymberakis, 2014) or angular (Chen, Menges, & Leblanc, 2005; Pinty, Widlowski, Gobron, Verstraete, & Diner, 2002) domain.

The objective of this paper is to describe and test a methodology capable of identifying areas with homogeneous vegetation cover from coarse spatial resolution time series without requiring ancillary data such as a land cover map. The idea is that such selection of time series can serve for applications requiring a spatialized sample of landscapes dynamics, and not a spatially exhaustive coverage that would forcefully include time series obtained from both spatially homogeneous and heterogenous areas. The method is designed for the Moderate-resolution Imaging Spectroradiometer (MODIS), one of the most widely used satellite remote sensing instruments for global land surface application, and it exploits both the high temporal and angular resolution that MODIS offers. The approach is here demonstrated for the particular application of crop growth monitoring, for which there is currently a strong demand for crop specific information at global scale (Justice & Becker-Reshef, 2007; Atzberger, 2013). Some early exploratory research on a single site in North American Great Plains showed some promising preliminary results (Duveiller & Lopez-Lozano, 2013). In this paper, the method has therefore been consolidated and improved using both simulated and real datasets, before applying it to various sites across the world to demonstrate its potential.

2. Particularities of the MODIS spatial support

The design of the MODIS instrument has important consequences on its spatial response, *i.e.* on the spatial support over which an observation is made. Like with other remote sensing instruments, the observation footprint is not equivalent to the ground projection of the squared pixel, or grid cell, in which the data is delivered to the user (Cracknell, 1998). This footprint is firstly defined by the instrument's net PSF, which depends on the optics, the detectors and the electronics involved in encoding the signal (Schowengerdt, 2007). During the processing of the data, the measurements are forced into a grid, such as the Level 2 Grid (L2G) in the case of MODIS (Wolfe, Roy, & Vermote, 1998), thereby causing gridding artefacts reported by Tan et al. (2006), and resulting in a mismatch between the observation support and the grid cells (see Fig. 1a).

MODIS is a whiskbroom instrument, which means it scans the Earth along a scan line perpedicular to the track and then uses the forward motion of the satellite to move from one scan line to the next. This scanning system has an across-track field of view of 110°, resulting in a wide swath that enables the observation of the entire equator every 2 days and to have a full daily global coverage above approximately 30° latitude (Wolfe et al., 1998). However, a consequence is that the view zenith angles (VZA) of daily observations can reach up to 65° at the end of a scan line when accounting for the Earth's curvature (Tan et al., 2006). As represented schematically in Fig. 1, the size of the observation footprints increases considerably with higher scan angles, compounding the mismatch between observation support and grid cell. Daily MODIS reflectances are provided with a layer indicating the observation coverage, or *obscov* for every pixel, defined as a ratio between: (1) the intersection area between the nominal observation and the grid cell; and (2) the nominal area of the observation. While *obscov* provides some valuable information of the adequacy between the observation support and the pixel, it is rarely used explicitly in remote sensing studies.

The critical point of interest for this paper is that MODIS retrievals attributed to a given pixel will have different footprints from one day to the next. This results from 1) artefacts in image registration and 2) from the 16-day revisit cycle of the platform, meaning that a different VZA is used in every day of that cycle. The premise of this work is that time series in which each consecutive observation falls within a spatially homogeneous surface can be expected to have a more temporally consistent signal that those covering a heterogeneous target, such as a mixture of adjacent fields covered by different crops (see Fig. 2). This study is based on a vegetation index, the NDVI, that is designed (partly) to reduce the effects due to differences in the bidirectional reflectance distribution function (BRDF) on the time series consistency. The working assumption is that differences in NDVI measured over the same crop type but with different angles are smaller than those obtained from different crops with the same angle. This also requires an assumption that the sampled area is relatively flat not to include a signal alteration due to topographical effects.

3. Material and methods

3.1. Exercise on simulated datasets

A set of simulation exercises is proposed to illustrate the expected behaviour of the MODIS temporal signal when it is composed of distinct land covers sampled with different observation footprints. Simulated datasets are created by randomly generating their spatial patterns using the Modified Random Cluster algorithm proposed by Saura

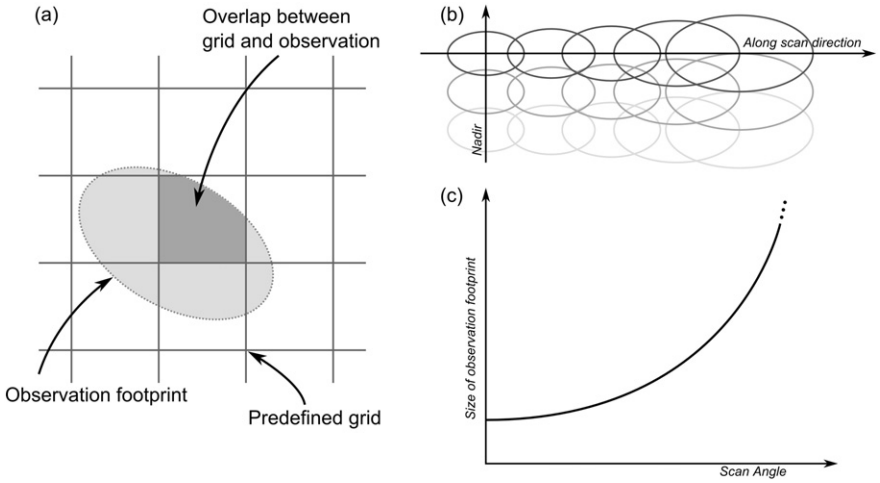

Fig. 1. Schematic representation of how the MODIS observation footprint is (a) different from the grid cell in which it is encoded in and (b & c) how it changes according to the scan angle. Figure adapted from (Duveiller et al., 2011).

& Martinez-Millan (2000) and implemented as the randomHabitat function in the R package secr (Efford, 2015); and then generating the temporal signal for a selection of land cover types based on prescribed NDVI and knowledge of the changing observation support of MODIS.

3.1.1. Generating landscape patterns

The first generated landscape is binary and serves to study interactions between two individual land covers exclusively. Four generic land covers are considered: deciduous broadleaf forest (DBF), winter crops (WCR), summer crops (SCR) and grasslands (GRA). In a series of analyses, each land cover will be assigned to one class in the binary mask while the other class will change amongst the other three. The mask is defined based on three parameters: *P*, the degree of fragmentation or aggregation of the patches; *A*, the proportion occupied by the patches; and *mp*, a minimal patch size in pixels. The values used here are $p = 0.5$, $A = 0.5$ and $mp = 15$, resulting in the mask shown on Fig. 3a. Since this binary case involves a strong oversimplification of agricultural systems, a more complex landscape including all four land cover types (with different shape types and spatial proportions) is generated (see Fig. 3b) using the parameterization in Table 1. Because the function only generates spatial patterns in binary mode, the multi-class landscape is constructed by first considering that it is completely covered by GRA, and then sequentially overlaying the SCR, WCR and DBF random patterns. In this way, the largest patches are mostly constituted of DBF, WCR patches are smaller, and SCR and GRA are even smaller, representing somehow a typical North European agricultural landscape. Both binary and multi-class landscapes have a spatial resolution of 10 m and cover respectively an area of 10 by 10 km and 20 by 20 km.

3.1.2. Generating temporal patterns

The temporal behaviour of NDVI is prescribed for each land cover type. It is simulated by combining a base NDVI value of *a*, a logistic function representing growth and a second logistic function representing senescence, as shown in the following formula:

$$NDVI(t) = a + \frac{K_g - a}{1 + exp(-B_g(t - M_g))} - \frac{K_s}{1 + exp(-B_s(t - M_s))} \quad (1)$$

where *K*, *B* and *M* define the shapes of the logistic functions for either growth (g) or senescence (s). Parameter values are chosen for each land cover (reported in Table 1) with the objective to represent their expected phenological behaviour (see Fig. 3c). The NDVI values are calculated for 250 consecutive days and assigned to the respective land cover to form an NDVI data block with 250 layers.

3.1.3. Simulating the MODIS observations

The next step consist in simulating how the MODIS instrument would see the artificial landscapes taking into account its 16-day revisit orbit and the particularities of its PSF. To do so, a specific PSF model is constructed for each of the 16 viewing geometries of the MODIS instrument at a given latitude. Details on how such PSF model are constructed and applied can be found in Duveiller et al. (2011) and Duveiller (2012). When convolving such a model on a single NDVI image that has a finer spatial resolution than MODIS (such as those generated above), the result is a spatialized field with the same fine spatial resolution but indicating what NDVI value a MODIS observation would contain if its centroid would fall at a given place. From this field, one can extract

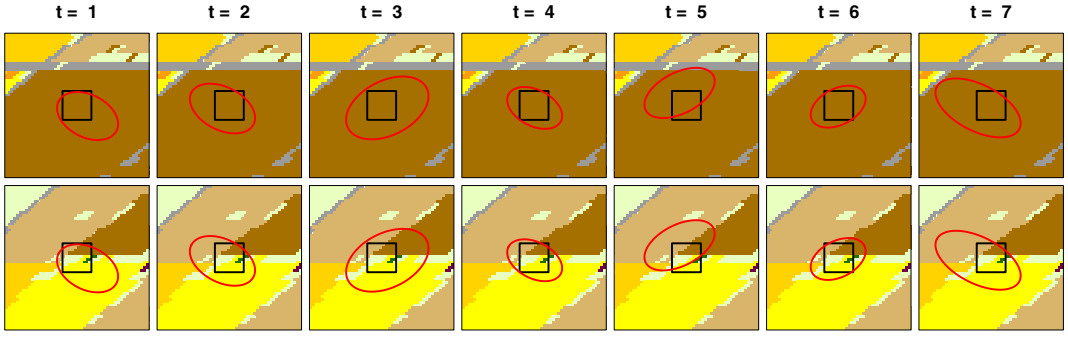

Fig. 2. Schematic representation of how the MODIS observation footprint (in red) can change in time for a fixed MODIS L2G grid cell (in black) and the repercussion this may have regarding what type of land surface is sampled. The top row illustrate a situation in which a pure crop signal is obtained, while in the lower row the signal comes from a mixture of 4 different crop types.

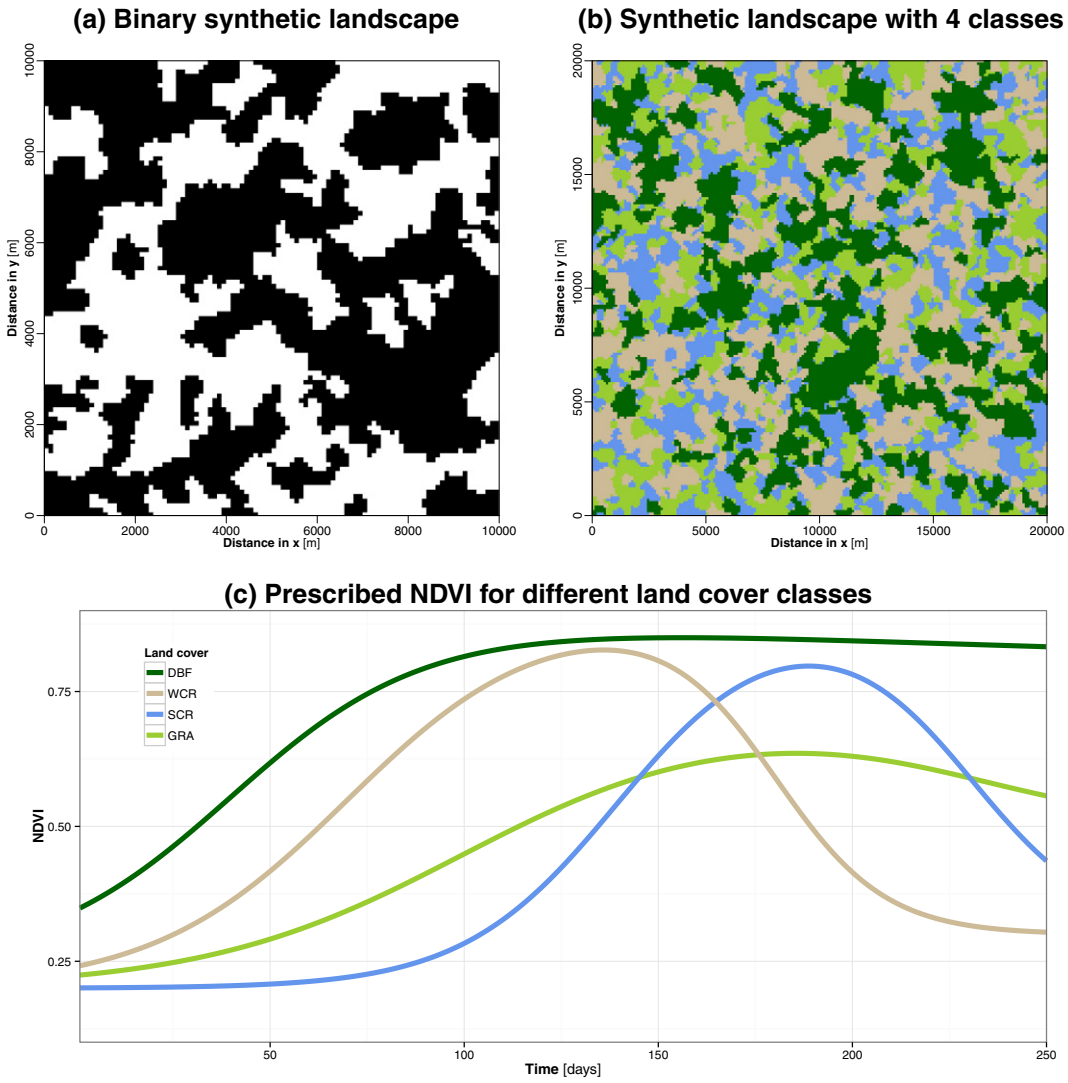

Fig. 3. Set-up for the exercise on simulated datasets: (a) the binary mask; (b) the multi-class map; and (c) the prescribed NDVI for each land cover class considered: deciduous broadleaf forest (DBF), winter crops (WCR), summer crops (SCR) and grasslands (GRA). Note that (b) is four times larger in area the (a).

values using a grid corresponding to all MODIS observation centroids and generate a MODIS image in its L2G grid. When doing so, a random positional error is considered following a gaussian distribution with $\sigma = 50$ m, corresponding to the geolocation error estimated for MODIS by Wolfe et al. (2002). Each of the 16 PSF models provides a different value according to the view zenith angle at which the observation is taken in that orbit. Knowing the daily sequence of MODIS orbits, the corresponding PSF model is applied to each consecutive layer in the NDVI data block to produce a simulated block of MODIS NDVI observations.

When convolving the PSF models on binary maps instead of NDVI images, the result is no longer an integrated observation but the contribution of the land cover delineated by the mask to the observation, otherwise referred to as pixel purity. Therefore, for each estimation of MODIS NDVI observations, the corresponding pixel purity for each class is also calculated (integrating also the same $\sigma = 50$ m positional error).

Two more steps are taken to render the simulated MODIS time series more realistic. A gaussian white noise ($\sigma = 0.025$) is added to all NDVI simulated values to represent the cumulated uncertainty from various effects including: intra-patch variability due to soil differences and potential species composition, inter-patch variability due to differences in land management, BRDF effects and residual noise from the atmospheric correction. The second step is to randomly remove 2/3 of the layers to take into account days when observations are not available due to cloud cover.

Table 1
Parameter values used for the generation of the simulated datasets. The parameters used to generate the spatial landscape patterns are: the degree of fragmentation or aggregation of the patches (p), the proportion occupied by the patches (A) and the minimal patch size (mp). The parameters used to generate the prescribed temporal NDVI behaviour are a base NDVI value (a) and two sets of three parameters (K, B and M) defining logistic functions of growth (g) and senescence (s) respectively.

Land cover	Space						Time			
	p	A	mp	a	K_g	B_g	M_g	K_s	B_s	M_s
Deciduous broadleaf forest (DBF)	0.50	0.3	20	0.25	0.9	0.04	40	0.1	0.01	180
Winter crops (WCR)	0.45	0.4	15	0.20	0.9	0.04	70	0.6	0.07	180
Summer crops (SCR)	0.40	0.5	10	0.20	0.9	0.05	140	0.6	0.06	230
Grasslands (GRA)	NA	NA	NA	0.20	0.7	0.03	100	0.2	0.04	230

3.1.4. Computation of the temporal SNR

The metric proposed to assess the temporal consistency of the daily time series and thus serve as a proxy for pixel purity is based on the notion of temporal signal-to-noise ratio (SNR). SNR is here defined as

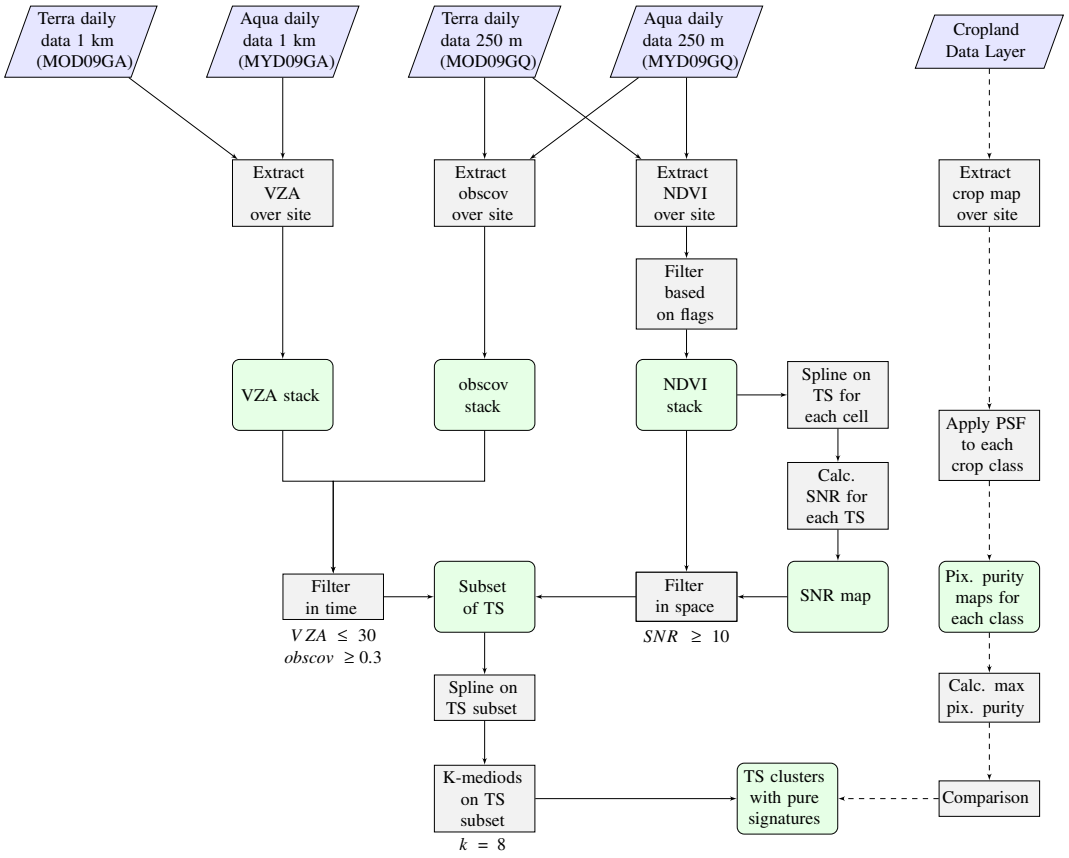

Fig. 4. Location of test sites overlaid on the global cropland extend estimated by (Pittman, Hansen, Becker-Reshef, Potapov, & Justice, 2010). Orange stars indicate the site is also a JECAM site.

the variance of the smoothed temporal signal over the variance of the residues during the period of interest (Fleming, 2010):

$$SNR = \frac{\sigma^2_{signal}}{\sigma^2_{noise}} \quad (2)$$

where σ^2 represents the variance of either the signal or the noise in the time series. To estimate the signal, spline smoothing is applied to each time series individually (cubic splines with 8 degrees of freedom are used). This method has been selected as it can be easily applied to smooth time-series with a wide range of signal-to-noise ratio, but other alternative smoothing algorithms that may be deemed more

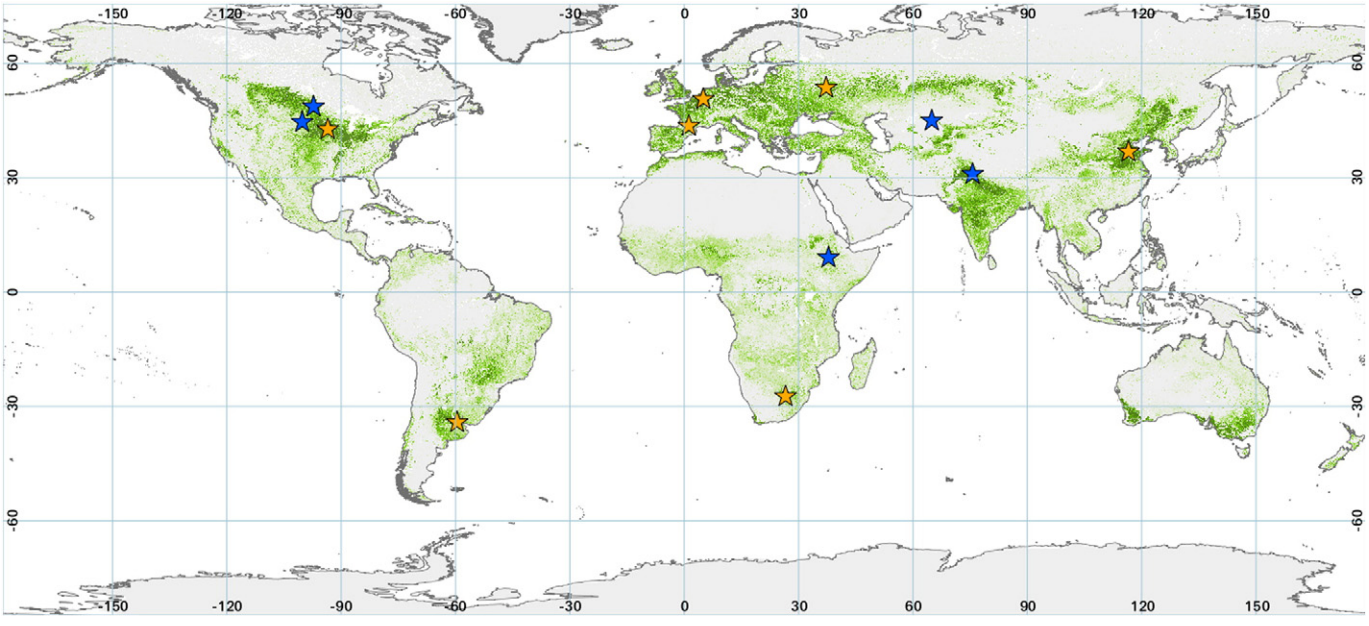

Fig. 5. Flowchart representing the methodological steps taken in this study for each test site. The dashed line represents the pathway for inter-comparison with the Cropland Data Layer, which is only applied to a subset of sites. TS stands for 'time series'.

Table 2
List of sites used in this paper. The coordinates represent the centre of the site. The start and end time represent the period during which the MODIS observations are collected. The source of the high spatial resolution reference image is also provided, in which all Spot4 images come from the Spot4/Take5 programme and CDL stands for the USDA Cropland Data Layer.

Site ID	Country (region)	Lat	Lon	Start date	End date	Reference
ARG	Argentina (B.A. province)	−34.20	−59.58	2012-Aug-01	2013-Apr-30	Spot4 2013-Mar-05
BEL	Belgium (Hesbaye)	50.65	5.00	2013-Feb-01	2013-Sep-30	Spot4 2013-Apr-22
FRA	France (Midi-Pyrenees)	43.61	1.20	2013-Feb-01	2013-Sep-30	Spot4 2013-Apr-13/17
SAF	South Africa (Free State)	−27.37	26.58	2012-Oct-01	2013-Apr-30	Spot4 2013-Mar-17
CHN	China (Shandong)	36.83	116.57	2013-Mar-01	2013-Jun-30	Spot4 2013-Apr-30
ETH	Ethiopia (West Shewa)	9.05	37.85	2013-Mar-01	2013-Nov-30	Spot4 2013-May-31
KHA	Kazakhstan (Kyzyl-Orda)	44.99	64.94	2011-Apr-01	2011-Oct-31	Landsat5 2011-May-21
RUS	Russia (Tula)	53.63	37.23	2012-Apr-01	2012-Aug-31	Landsat7 2012-May-02
IND	India (Punjab)	30.98	75.64	2012-Oct-01	2013-May-31	Landsat7 2013-Feb-28
US1	U.S.A. (South Dakota)	44.60	−100.22	2007-Mar-01	2007-Oct-31	CDL 2007
US2	U.S.A. (Iowa)	42.74	−93.50	2007-Mar-01	2007-Oct-31	CDL 2007
US3	U.S.A. (North Dakota)	48.68	−97.25	2013-Mar-01	2013-Oct-31	CDL 2013

appropriate can be used. The variance of the signal is calculated based on the values estimated by the spline interpolation at the time corresponding to every observation, while the variance of the noise is calculated based on the estimated values minus the observed values. For the simulated datasets, the SNR can be then compared to the pixel purity calculated using the PSF models.

3.2. Exercise on real landscapes

To show the applicability of the proposed approach beyond the simulated exercise described above, the methodology is also applied to real agricultural landscapes as summarised in the flowchart in Fig. 4.

3.2.1. Study site

Twelve sites of approximately 50 by 50 km have been selected across various agro-ecological landscapes (see Table 2 and Fig. 5). The choice of the sites was reasoned to provide a diversity of field spatial patterns in different climates and for different crops. The sites were also chosen to represent important crop producing regions including the US Corn belt, the north China plains, Punjab (commonly considered to be the breadbasket of India), northwestern Europe, the province of Buenos Aires in Argentina, the Ethiopian highlands and irrigated systems in Central Asia. Where possible, the sites were deliberately chosen to match those of the Joint Experiment of Crop Assessment and Monitoring (JECAM, http://www.jecam.org/) of the Group of Earth Observations (GEO). An additional requirement in the selection of study sites was the availability of reference images or classifications at a decametric spatial resolution that can provide a visual representation of the spatial structure of the landscape within the crop season of interest. That requirement was obtained from 3 different sources: (a) crop maps from the USDA National Agricultural Statistics Service Cropland Data Layer (CDL henceforth) and available at http://nassgeodata.gmu.edu/CropScape; (b) images from the SPOT4 (Take5) experiment lead by CESBIO and available at http://www.cesbio.ups-tlse.fr/multitemp/?page_id=406; or (c) Landsat imagery from USGS downloaded from http://landsat.usgs.gov/Landsat_Search_and_Download.php.

3.2.2. Real MODIS data

MODIS collection 5 daily surface reflectance data was retrieved for all sites from the NASA LPDAAC data servers accessible at http://reverb.echo.nasa.gov/reverb but in this paper downloaded via the dedicated R package called MODIS (Mattiuzzi, 2014). To avoid any unnecessary reprojection of the MODIS data, the actual boundaries of each site are defined based on the MODIS L2G grid by selecting a box of 201 by 201 pixels at 250 m resolution with the central pixel covering the coordinates listed in Table 2. Note that since the MODIS L2G is in a sinusoidal projection centred on Greenwich, the sites have an increasingly sheered shape when they are located further away from that meridian. Similarly, landscape spatial patterns in this projection have the same sheered effect when shown in this projection.

For both Terra and Aqua platforms the products at 250 m spatial resolution (MOD09GQ and MYD09GQ) are needed to obtain the red band (b01), the near-infrared band (b02), the *obscov* and the 250 m quality flags. The 1 km spatial resolution products (MOD09GA and MYD09GA) are also required to obtain the VZA and quality state flags (indicating cloud coverage). For each site, data is collected for a different period covering the main crop growing season (see "Start date" and "End date" columns in Table 2). The data are summarised into 3 stacks per site: one for VZA, one for *obscov*, and one for NDVI (calculated from the red and near-infrared bands at 250 m spatial resolution). A time vector to indicate when images are available is generated, incorporating the differences in overpass between the Terra (in the morning) and Aqua (in the afternoon). The MODIS quality flags are used to block out the observations that are either contaminated by clouds or of limited quality.

3.2.3. Selection of purer time series based on SNR

SNR is calculated for every time series in the MODIS NDVI stacks just like it was done for the simulated datasets (Section 3.1.4). The SNR is spatialised into a map that can serve to identify pixels with a radiometrically homogeneous behaviour along the period of interest. This can serve to isolate the time series that can serve to represent the crop specific temporal dynamic in the study zone. A spatial subset is done by filtering out all time series that do not satisfy the criteria: $SNR \geq 10$. This value has been selected empirically seeking a compromise between signal purity and spatial representativity of the selected sample set, with the added intention of keeping a single threshold for all test sites in order to compare performance. A temporal filtering is then applied to remove all individual observations which do not satisfy the following condition: $VZA \leq 30 \cap obscov \geq 0.3$. This follows the recommendations from a previous study (Duveiller et al., 2011) that showed how such filtering improved the temporal quality of MODIS time series. Although these observations are necessary for the calculation of the SNR (they contribute considerably to sampling the vicinity of the grid cell), they are prone to have lower quality for properly characterising the temporal signal than near-nadir observations with a smaller footprint.

3.2.4. Cluster classification

The final step then consists in characterising the crop specific temporal dynamic in the study zone. A new spline curve is fitted on the filtered observations to produce a subset of smooth NDVI time series with a regular daily time step. These are then clustered together based on the shape of their temporal behaviour using the PAM (Partitioning Around Mediods), or k-mediods, clustering algorithm as implemented in the R package cluster (Maechler, Rousseeuw, Struyf, Hubert, & Hornik, 2014). This algorithm is a robust version of the more typical k-means algorithm that can be used along with the silhouette technique to obtain an optimal number of clusters (Rousseeuw, 1987). Here the initial number of seeds is set to $k = 8$ and is then reduced to a smaller number (3 to 4) automatically according to the similarity of the initial groups.

The final result is the separate clusters representing the differentiated plant covers present in the site.

3.2.5. Comparison with a high resolution crop classification

For the 3 North American sites (US1, US2 and US3), the results can be confronted with an independent product: the crop specific classification provided by the CDL maps. To quantify the adequacy between the MODIS observation footprints and the CDL, the actual crop specific pixel purity maps are generated by convolving PSF models on crop specific masks in the same way as described in Section 3.1.3. In this case, only one PSF model is used corresponding to the worst observation scan angle that can be encountered (55°). The calculation of pixel purity maps is done separately for every class in the CDL that covers more than 5% of the pixels in the study area (see Fig. 6). The information contained in these various maps is then merged in a pair of layers, the first containing the maximum crop specific pixel purity obtained for a given MODIS grid cell, and the second indicating for every cell what is the dominant crop (*i.e.* which crop mask generated the highest purity for that cell). This information can then be crossed with the results obtained from the clusters of time series.

4. Results

4.1. Theoretical relationship between SNR and pixel purity

The simulation exercise provides the first evidence of how the temporal coherence of the MODIS signal, as expressed by the SNR, can serve to estimate pixel purity. Fig. 7 resumes the information from the simulation exercise using the binary land cover maps. Because each time series has various pixel purity values (depending on the viewing angle and the geopositioning error), the mean purity divided by the standard deviation of the purity (μ/σ) is used in the plot abscissae in order to have a single value per time series. Because the landscapes are binary, only points corresponding to time series with a mean pixel purity above 0.5 are shown. This does not pose a problem as values with lower purities have a purity above 0.5 for the reciprocal land cover, and are thus shown in the other colour.

Fig. 7 illustrates how a clear positive relationship is shown between pixel purity and SNR for all but one case. However, the relationship is different depending both on what land cover is targeted, and with what it is mixed with, even thought the exact same mask has been used. The r^2 informs on the percentage of the variance in the purity

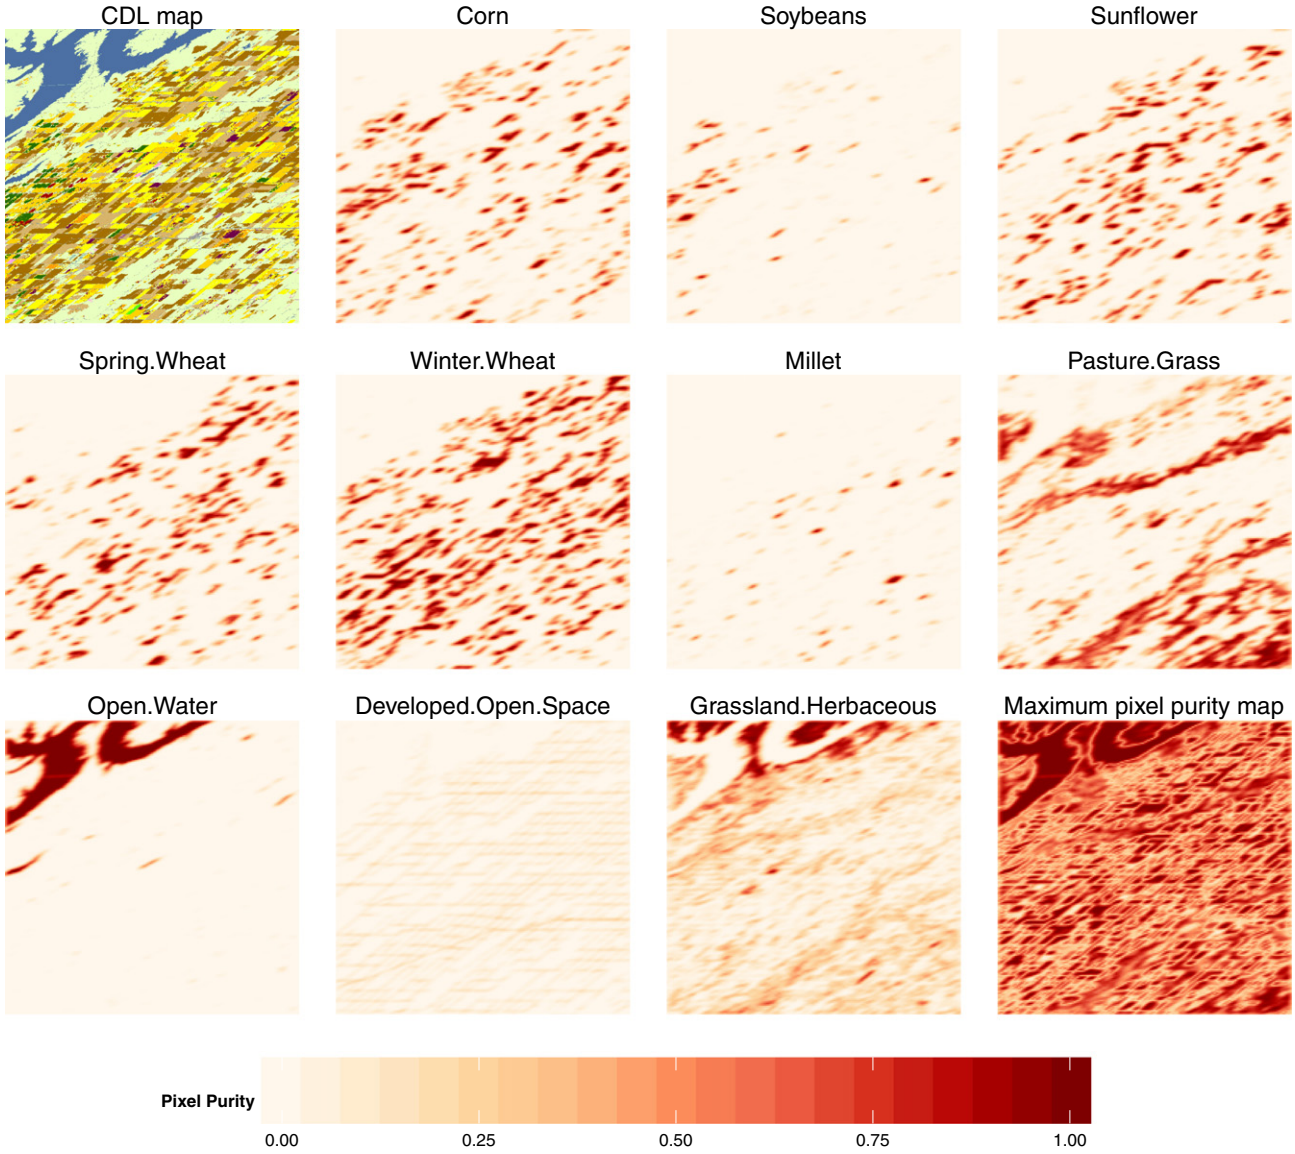

Fig. 6. Cropland Data Layer map for US1 site and the derived pixel purity maps obtained for each land use type covering at least 5% of the area. The last map is the maximum pixel purity map where every pixel contains the maximum possible purity independent of the land use it covers.

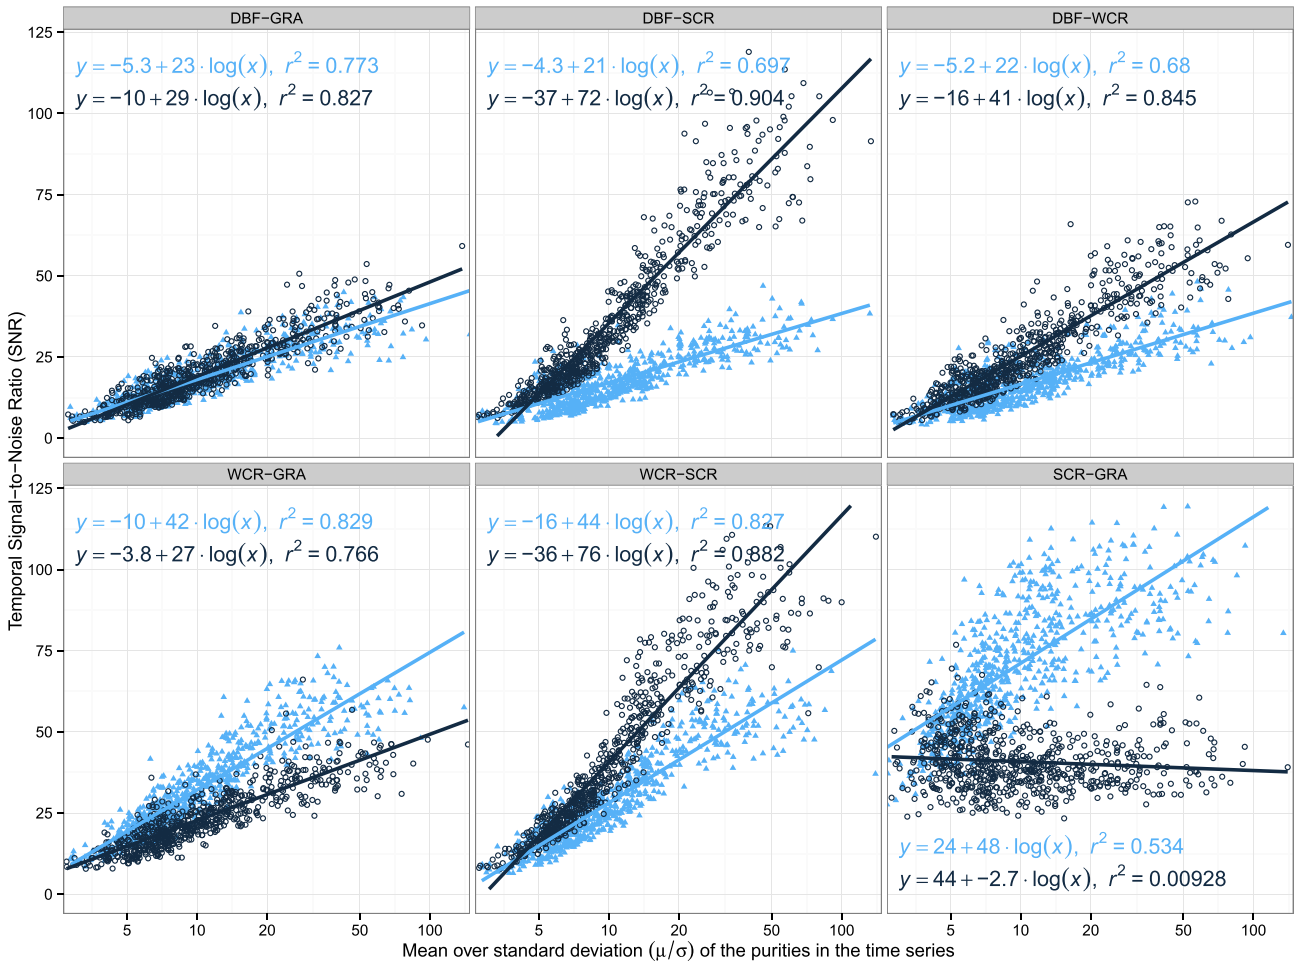

Fig. 7. Results of the simulated exercise based on the binary mask. The light blue triangles and the dark circles respectively represent the first and second land cover in each land cover pair. The x-axis uses a logarithmic scale, and plots the mean purity divided by its standard deviation (μ/σ) for each time series. For the sake of clarity, only points corresponding to time series with a mean pixel purity above 0.5 are shown, and the regression lines are also based only on these subsets.

that can be explained by the SNR metric. The steepness of the slopes (using a logarithmic axis) indicates the sensitivity of the SNR to changes in purity. The slope is steeper for the land covers with a more marked phenology (typical agricultural classes such as SCR and WCR) because these will have a higher σ_{signal}, and thus a higher SNR, for a given noise level σ_{noise}. This property has the consequence that if a single SNR threshold is selected, this will favour catching high pixel purities for agricultural land covers, which is a desirable property for this demonstration targeting agricultural applications. Together, both r^2 and slope indicators suggests that land cover pairs having more contrasting spectro-temporal properties (such as WCR-SCR or DBF-SCR) will be easier to discriminate than those with more similar curves (*e.g.* SCR-GRA). The relationship is flat in the case of grasslands mixed with summer crops (but not inversely), indicating that under these specific circumstances, grassland cannot be isolated from summer crops using the SNR.

When all four classes are considered together, the SNR also relates with pixel purity (as shown on Fig. 8), generally retaining similar slopes but with lower capacity to explain the variance in the purity (as expressed by lower r^2 values). As expected, the classes with patches of higher median surfaces (DBF > WCR > SCR > GRA) have time series with the highest purity values, but they are not necessarily the ones with the steeper slopes or higher r^2, partly because of what has been mentioned before: the SNR favours times series with marked phenology. Overall, the results from both binary and multi-class experiments suggest that SNR is generally a good proxy for pixel purity, but the strength of this relationship varies depending on the composition of land cover in the observation footprint.

4.2. Application to landscapes with reference crop masks

To illustrate the performance of the proposed methodology in real conditions, a series of contrasting situations (with respect to SNR and land use) is selected within a single test site (US1). From this selection (presented in Fig. 9), it is clear that the SNR maps adequately reflect the spatial patterns within the landscape. MODIS grid cells straddling over the border between fields (Fig. 9b & d) covered by different crops typically have much noisier temporal signatures (and thus lower SNR) than those falling inside a homogeneous field (Fig. 9a, c & e). The latter clearly shows distinctive temporal signatures depending on whether they fall over a summer crop (such as corn on Fig. 9c) or a winter crop (such as spring wheat on Fig. 9a or winter wheat on Fig. 9e). It can also be noted that when SNR is high, there is a good complementarity between Aqua and Terra data, and the time series have little residual noise even with high viewing angles.

The capacity of the SNR to relate to the radiometric homogeneity (or heterogeneity) of the landscape elements is quantified using the maximum pixel purity map generated for the landscapes US1, US2 and US3. Fig. 10 presents, for classes of increasing SNR value, the proportion of each of them that is composed of different maximum pixel purity classes according to the CDL. This figure shows how higher SNR

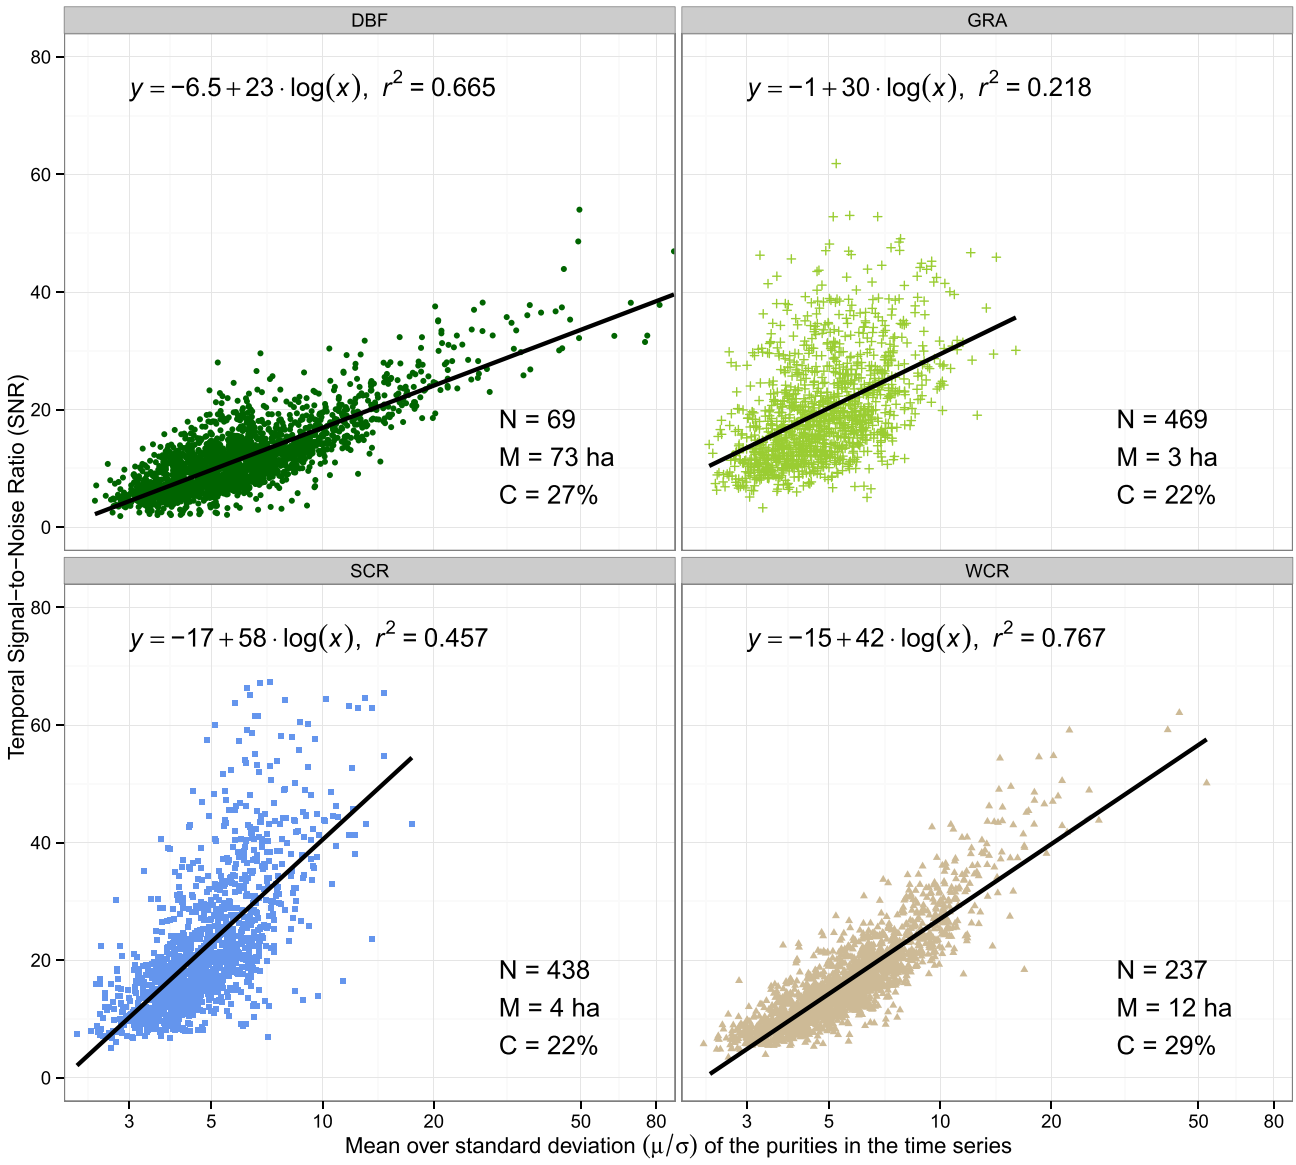

Fig. 8. Results of the simulated exercise based on the four-class map. The x-axis uses a logarithmic scale, and plots the mean purity divided by its standard deviation (μ/σ) for each time series. Each subplot also indicates the number of patches (N), the median patch size in hectares (M) and the percentage of the map covered by a given land cover (C).

values correspond to purer pixels for all 3 landscapes, albeit showing different patterns in the increase of purity with respect to SNR values. US2 still contains a higher proportion of mixed pixels at high SNR values because most consist of a mixture of corn and soybeans with very similar phenology. The proportion of higher maximum pixel values is actually underestimated in all cases because the PSF model used is conservative. Indeed, the inclined PSF model is calculating the pixel purity at a scan angle of 55°: the least favourable situation to obtain pure pixels. This only occurs for certain orbits and over certain parts of the swath. The pixel purity calculated with PSF with smaller angles (which are more frequent) are much higher (results not shown here). Furthermore, the pixel purity estimations may also have a certain degree of error depending on the quality of the CDL classification.

The second result of the comparison exercise provides the detailed composition of the identified clusters according to the CDL maps. Both the initial sets of 8 clusters and the sets of aggregated clusters by the silhouette method are shown for the three North American landscapes in Fig. 11. For US1 (Fig. 11a), several pre-aggregation clusters are almost exclusively composed of a single crop: winter wheat, corn, spring wheat, pasture or even sunflower; and only one cluster contains a substantial mixture of classes. After aggregation, the three resulting groups separate crops between winter wheat, summer crops and a third groups with a mixture of spring wheat and pastures. The result in US2 (Fig. 11b) shows that corn is identifiable using the $SNR \geq 10$ threshold if a small amount of soybean commission is tolerated, but the opposite is not the case. For the very heterogeneous landscape of US3 (Fig. 11c), the 8 primary clusters do contain some almost pure groups of spring wheat and soybeans, while the 3 aggregated clusters seem to separate spring wheat from summer crops, leaving a third highly heterogeneous cluster containing wetlands, amongst other classes.

4.3. Application to other contrasting agricultural landscapes

The results showing the SNR maps and the identified temporal cluster for all 12 sites are shown in Fig 12 and Fig. 13. The SNR maps are also available as downloadable KMZ files in the online supplementary material of this paper. Overall, the SNR maps reflect well the spatial structure of most landscapes. An exception is US2, which shows little correlation

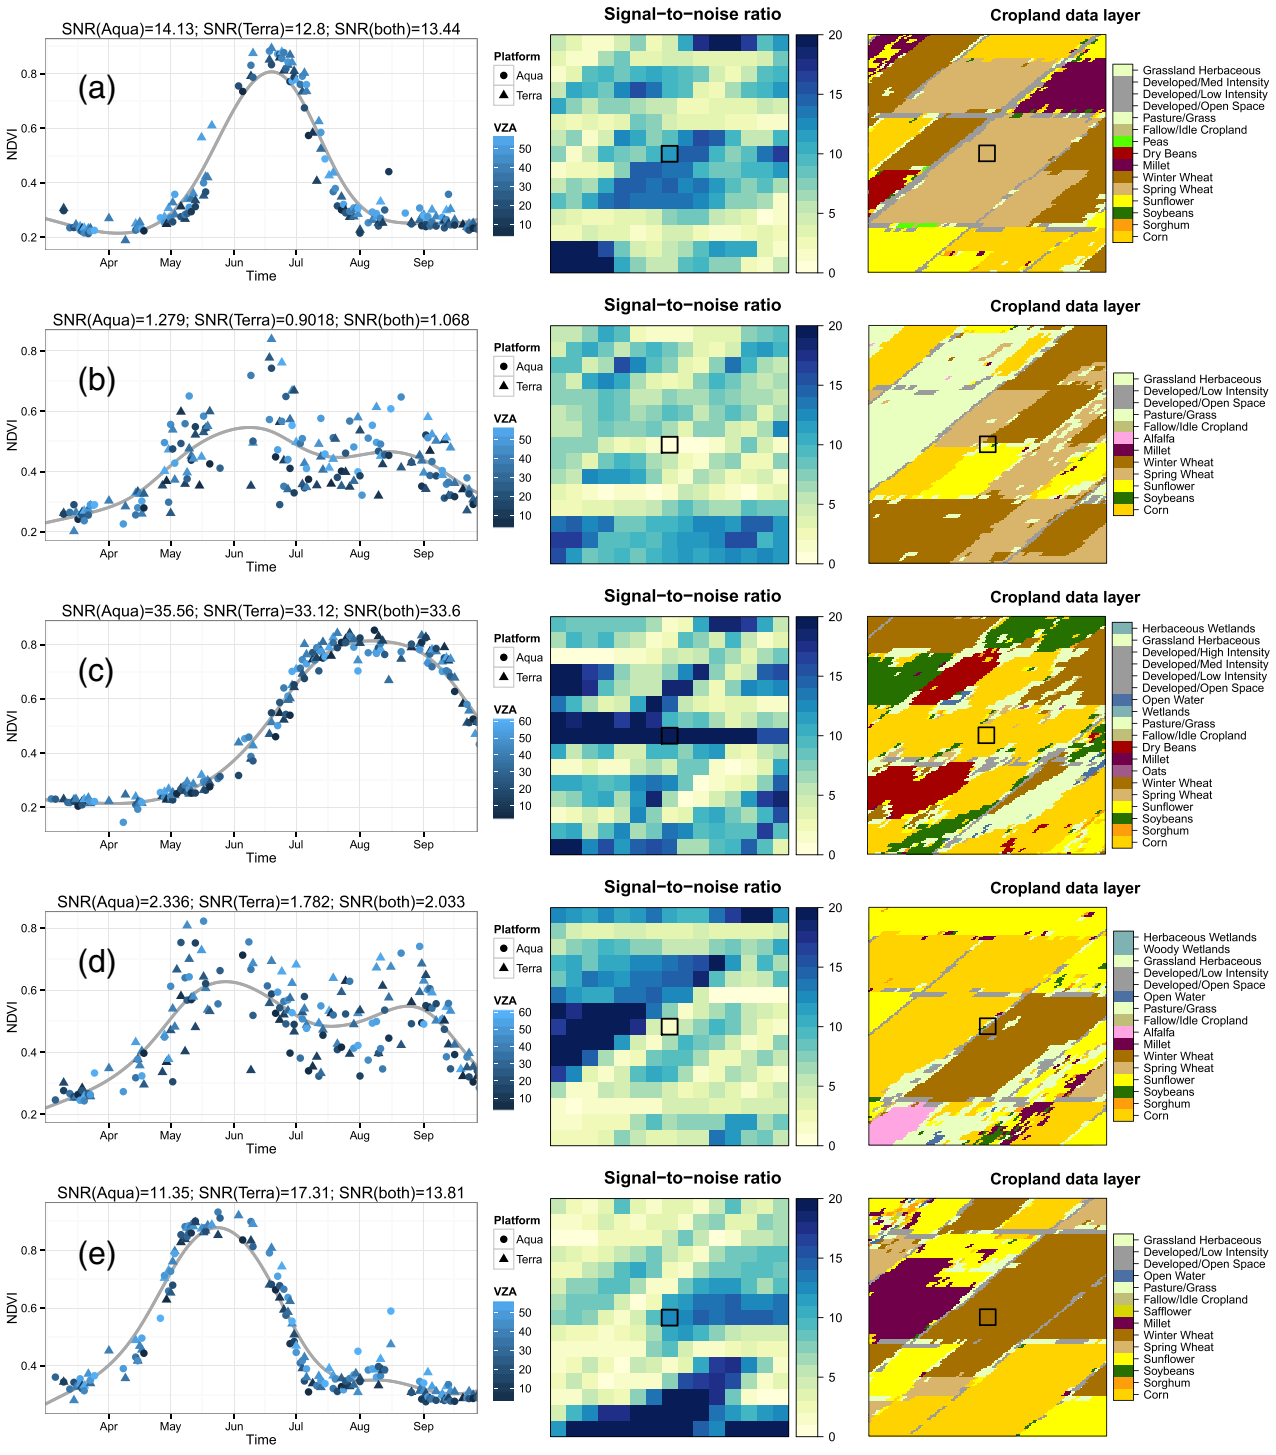

Fig. 9. For a selection of points in the US1 demonstration site, the temporal behaviour of NDVI is shown (left column) next to the corresponding signal-to-noise (SNR) maps of the vicinity of the target pixel (middle column) along with the land use classes reported by the USDA-NASS Cropland Data Layer (right column). The grey line in the time series represents the smoothing based on a spline over all data points.

between SNR and the reference classification. This is mainly because the landscape is practically a binary mosaic of soybean and corn, two crops which in Iowa have a very synchronised phenology and therefore appear as a single, non-fragmented spatial entity. IND and CHN also show very little heterogeneity over croplands. For these cases the reason lies in the fact that essentially a single crop is grown over the area during the studied period. However, in all these three landscapes, crops can clearly be separated from non-crop landscape features such as urban areas and rivers using the SNR maps. Towns and rivers have low SNR values since these landscape feature typically have low NDVI phenology (hence low σ_{signal}) and are susceptible to have high radiometric variations with different viewing configurations (hence high σ_{noise}).

The methodology results in clusters with distinctively different temporal behaviours in various landscapes. In FRA and BEL, the 3 resulting clusters have the temporal NDVI patterns expected for the 3 dominant

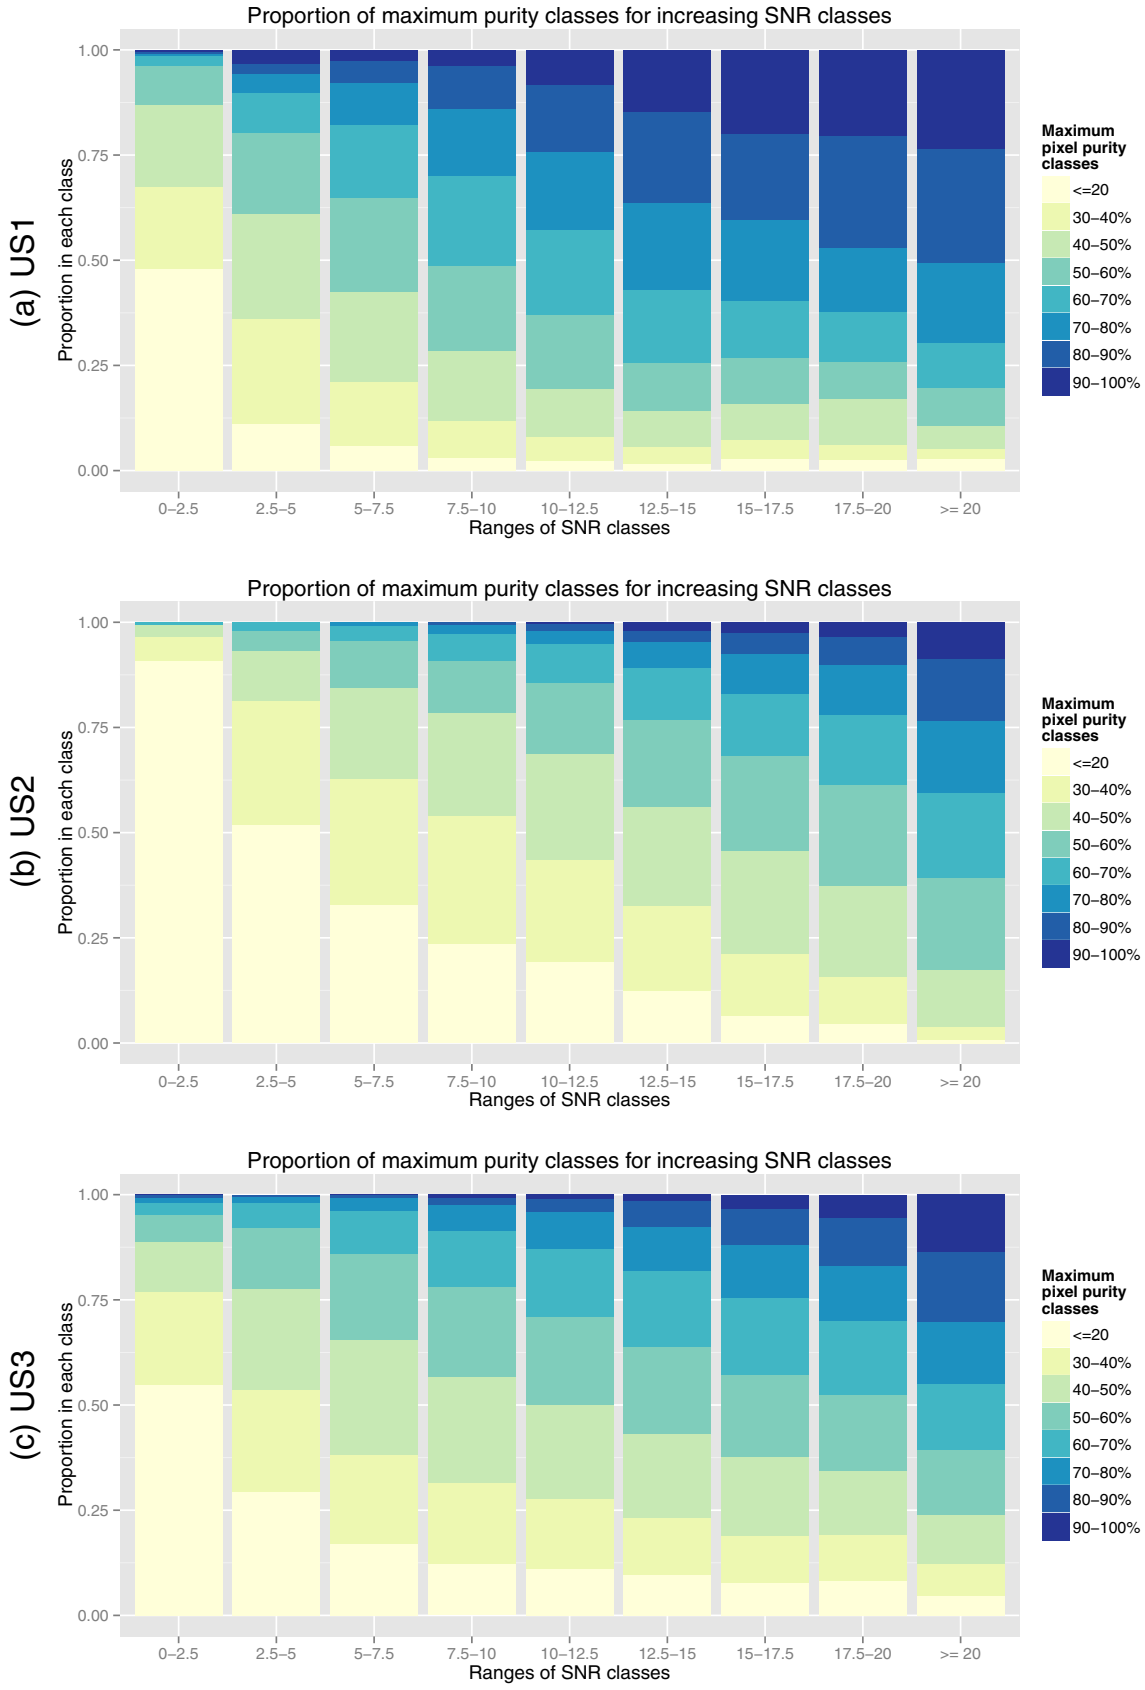

Fig. 10. Proportion of maximum pixel purity values (calculated based on the Cropland Data Layer) obtained for different classes of SNR.

land uses: winter crops, summer crops and deciduous forest. In ARG, two single season and one double season cropping patterns are isolated. In US1, the separation appears to be between winter wheat, summer crops and a mixture of spring wheat and pastures. Several landscapes also show that identified time series are spatially clustered together in big fields (see RUS, KHA, SAF and ARG), adding confidence to the

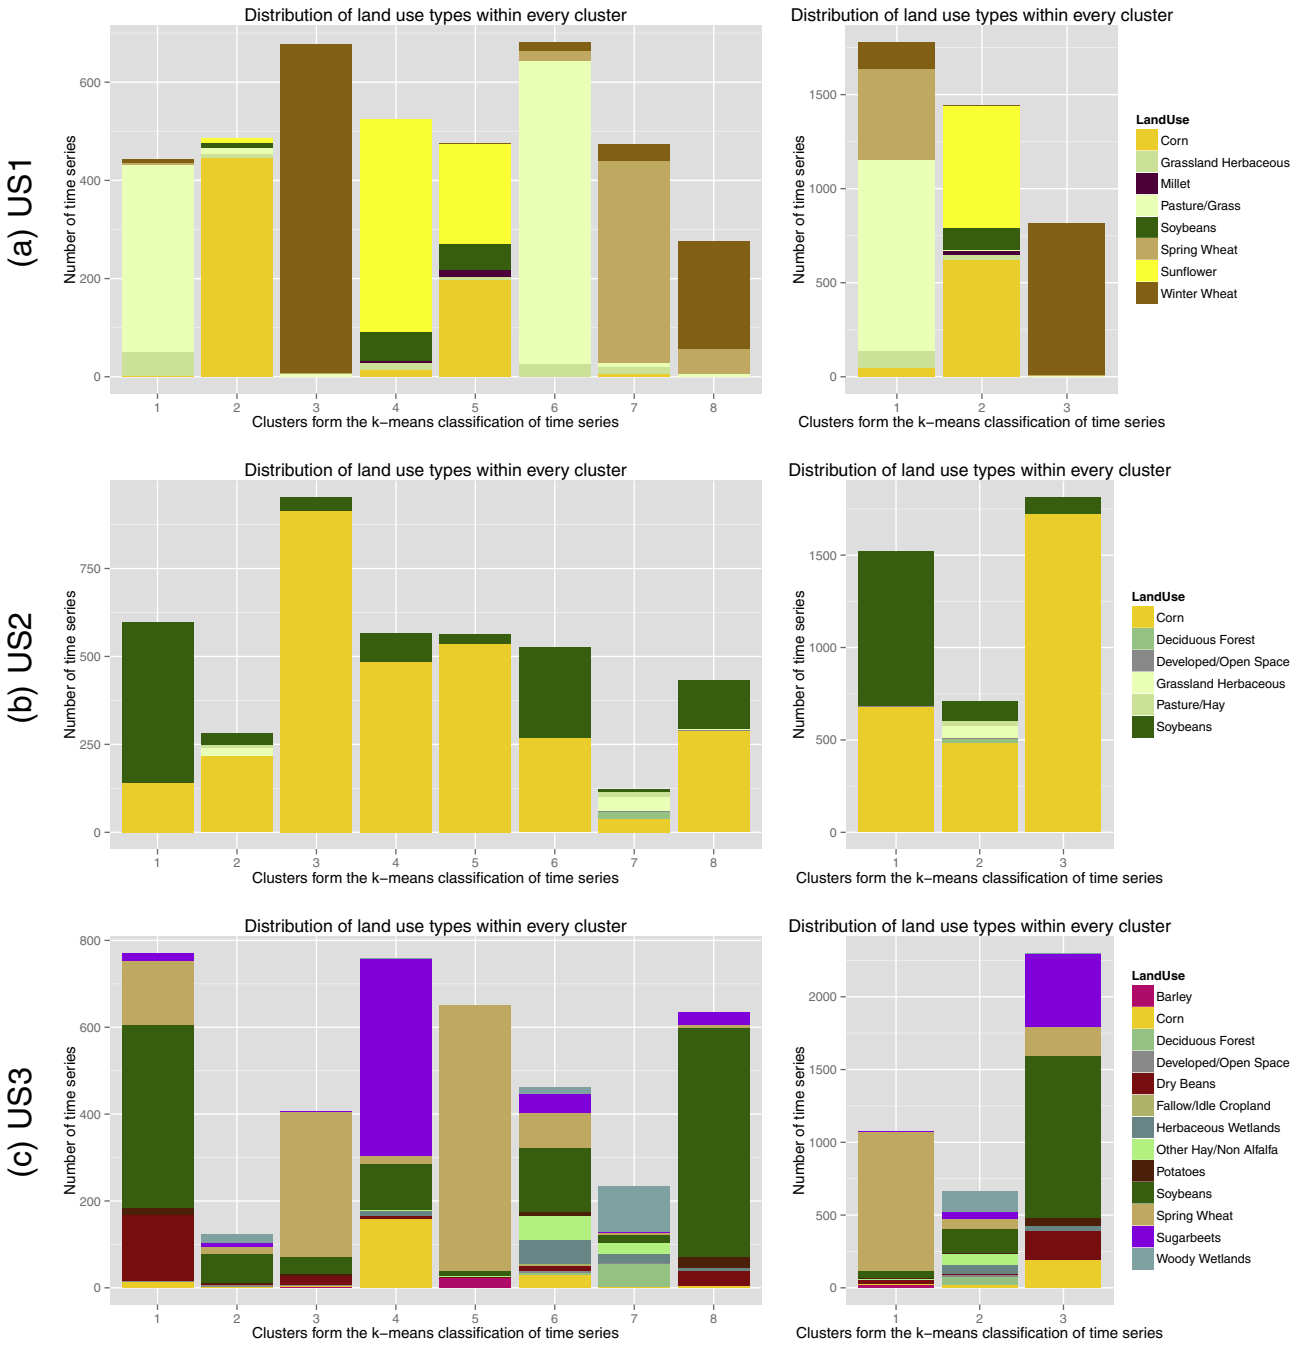

Fig. 11. For each of the 3 north American landscapes, the bars represent the distribution of land use types that each cluster is composed of for either (left column) the simple k-mediods clustering or (right column) the same clusters after silhouette aggregation.

validity of the temporal clustering approach since it uses no spatial information. Only in ETH does the method seem unconvincing in extracting clear temporal signals. This is probably a question of scale: field sizes in Ethiopia are much smaller than the 250 m resolution of MODIS.

5. Discussion

The overarching idea behind this paper is that the angular and temporal resolution of daily MODIS time series can be exploited to identify the time series of surface properties originating from a homogeneous area. The simulation exercise demonstrates how the indicator selected to condense this information is, in theory, strongly related to pixel purity, but also that this strength depends on the composition of land cover types within the observation footprint. In practice, the results of the exercise on real landscapes clearly show that the approach has a great potential for crop specific agricultural monitoring. Indeed, in most of the cases analysed here, coherent samples of summer crops time series are effectively separated from those of winter crops, a feat that is seldom achieved with pixels of this size. Beyond this demonstration exercise, there seem to be no obstacles to use the proposed approach for various other applications that could benefit from either isolating subsamples of time series covering homogeneous vegetation types or knowing *a priori* where mixed pixels are.

The approach based on the SNR to resume the angular and temporal information contained in the MODIS signal is modular. It can be fine-tuned according to the desired application. By comparing the noise to the amplitude of the signal using the SNR metric, the method was

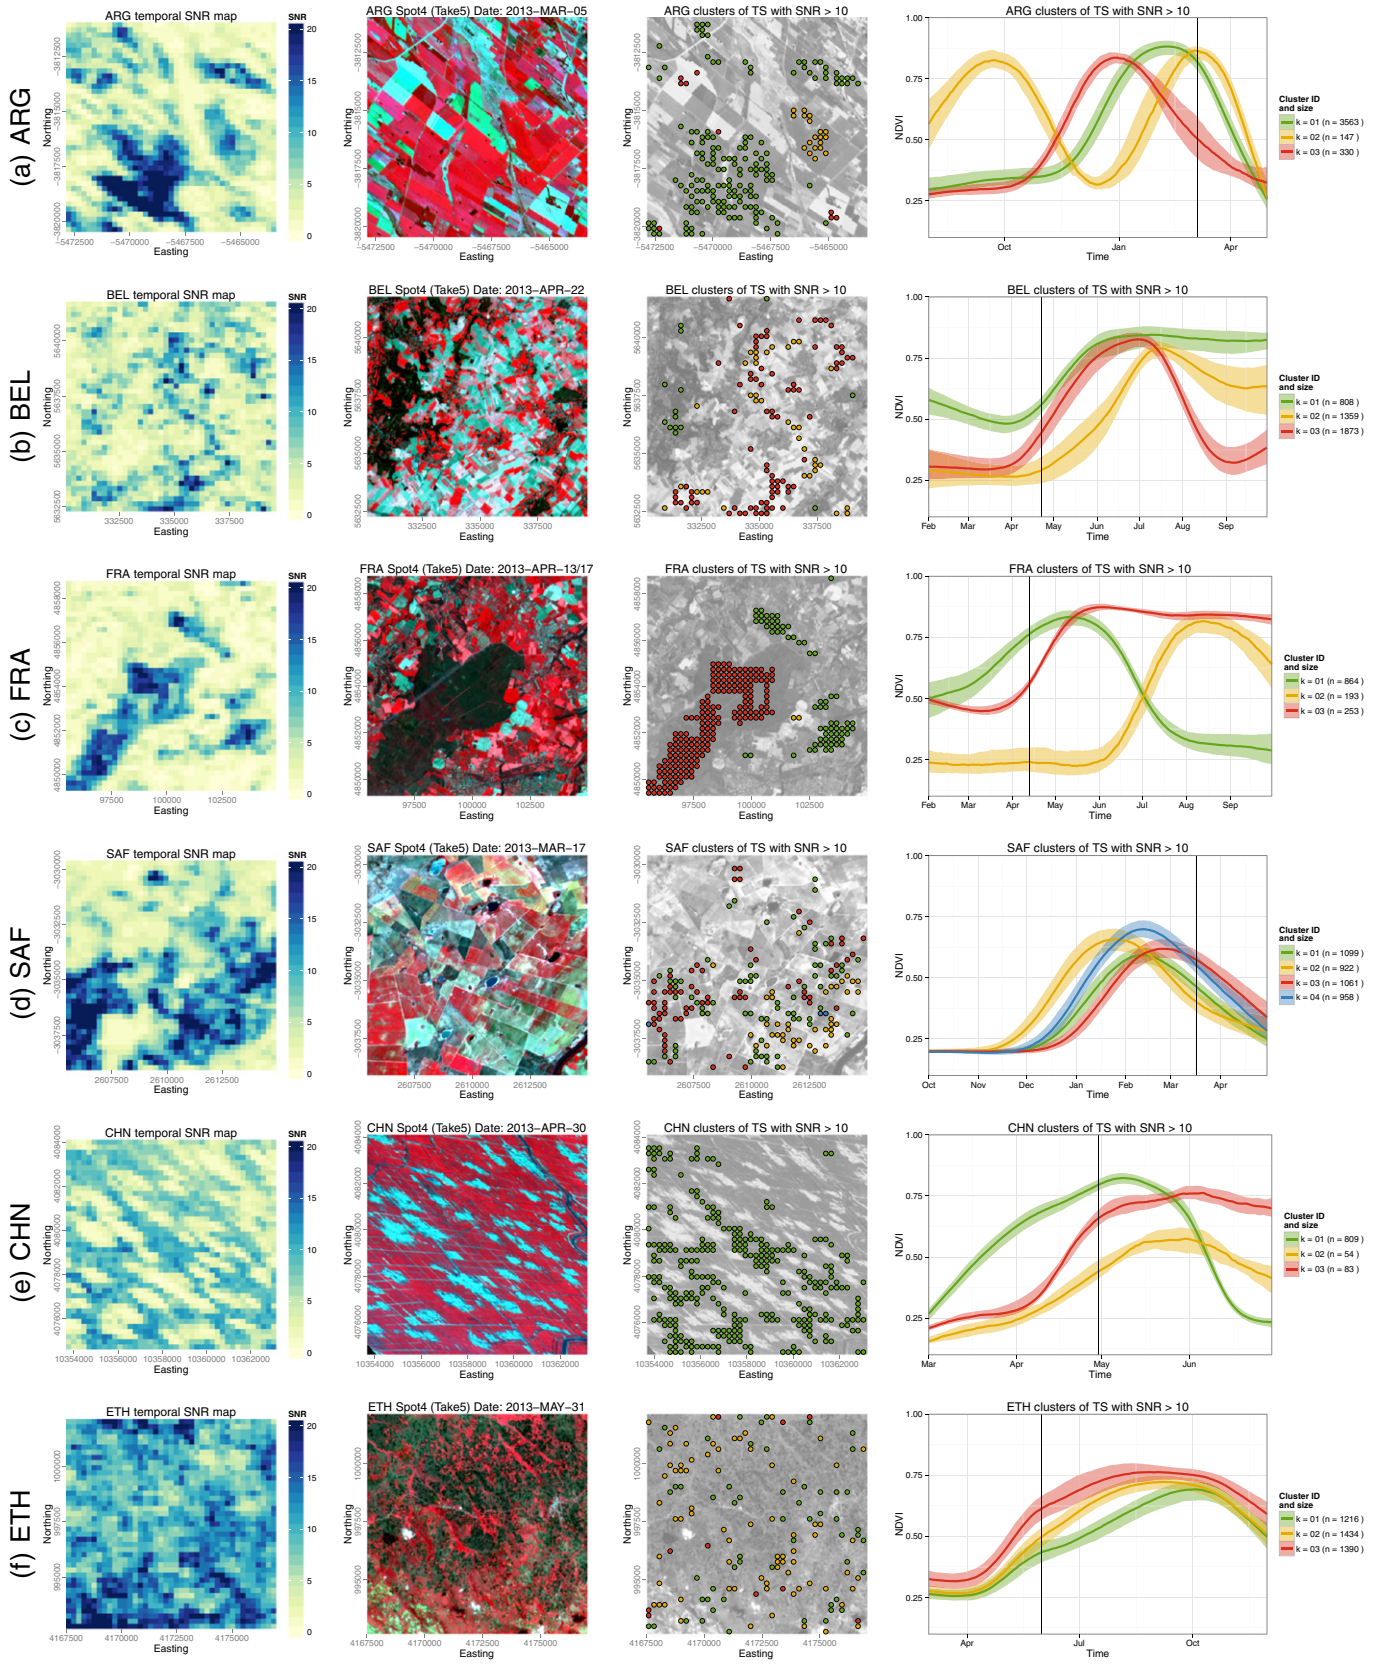

Fig. 12. For the first 6 landscapes listed in Table 2, each row shows (from left to right): the signal-to-noise ratio (SNR) map indicating the spatial heterogeneity of the landscape; the reference image; the spatial selection of identified crop specific clusters; and their median temporal behaviour overlaid over their 5–95 percentile range of variability. For the sake of clarity, only an ~10 by 10 km extract of the test sites are shown in the first 3 columns, but temporal clusters refer to all time series in each ~50 by 50 km area. The solid line in the cluster graphs indicates the timing of the reference image.

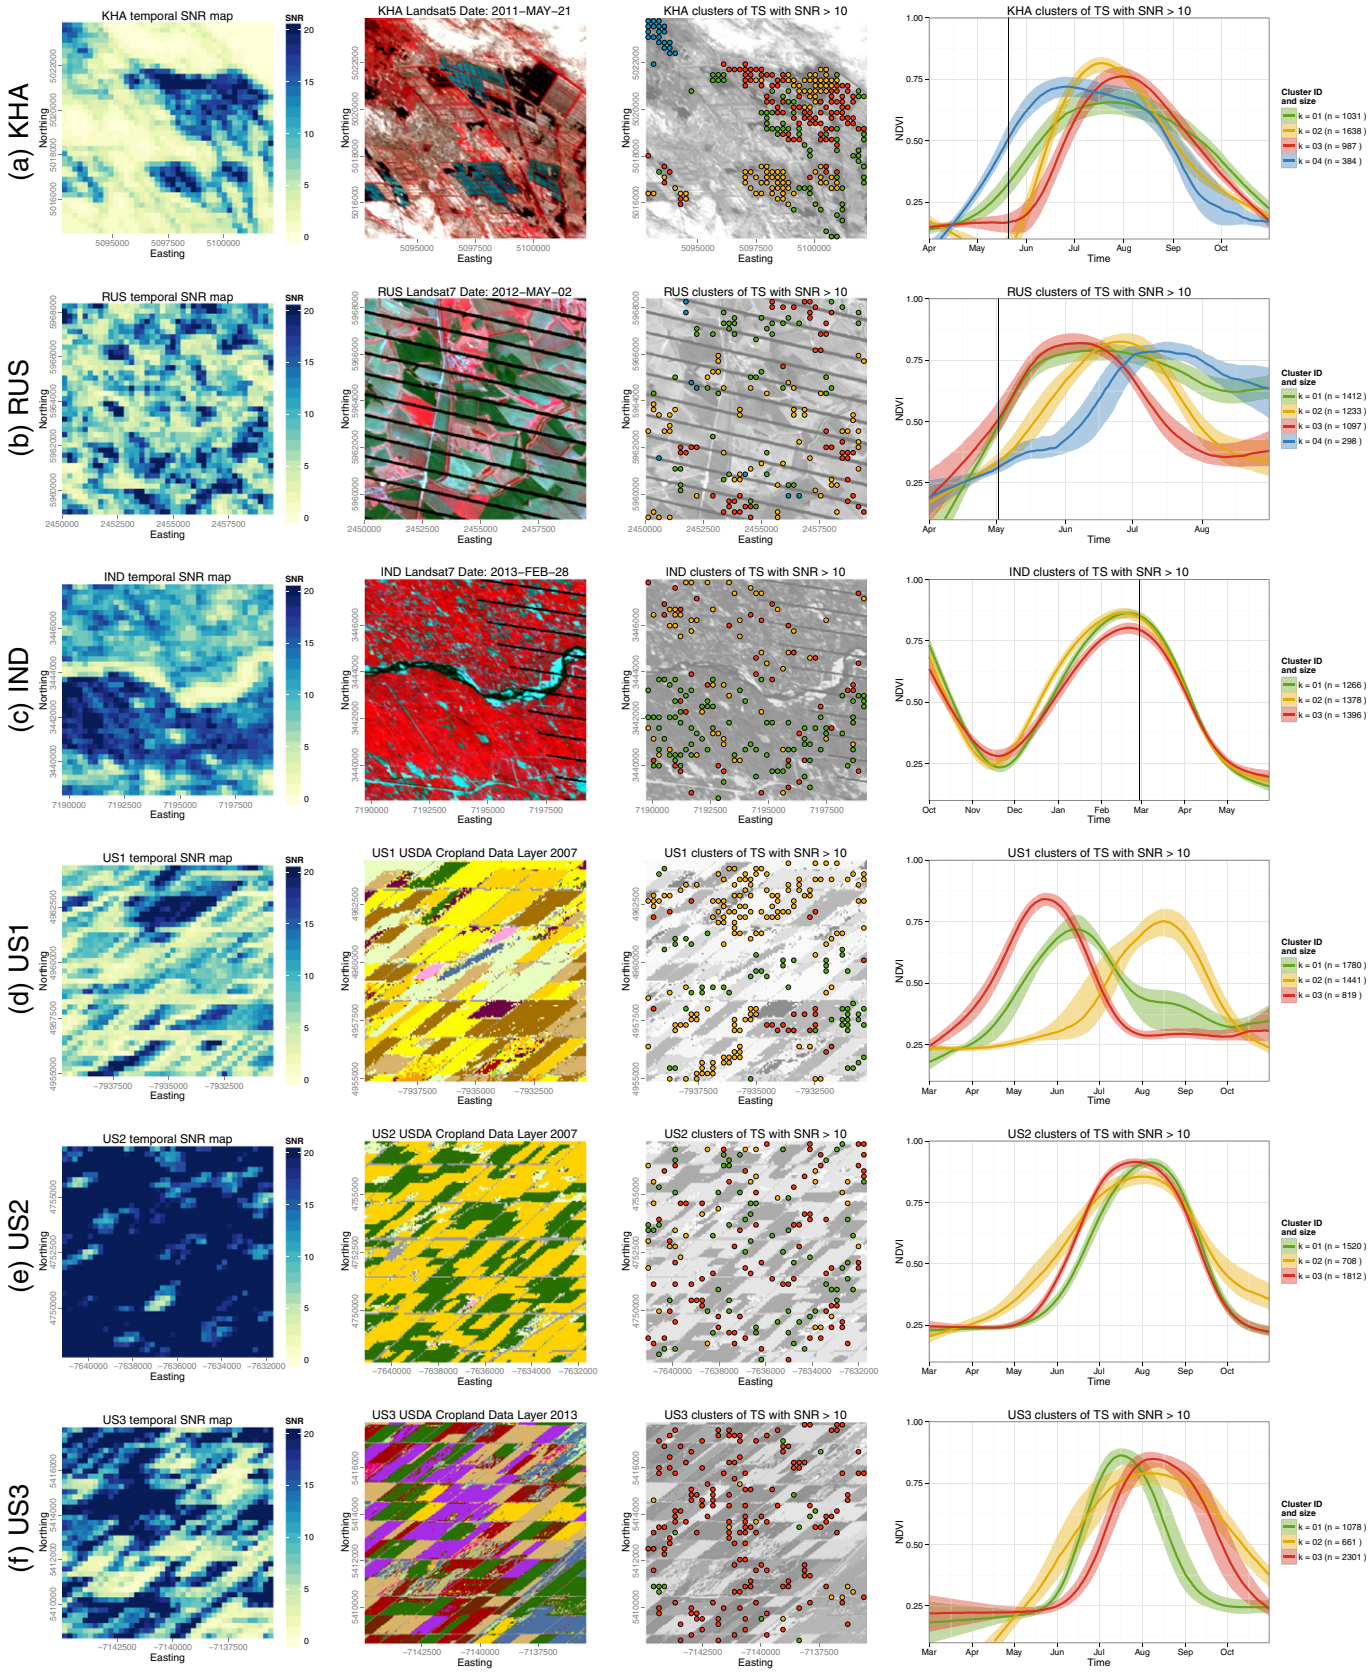

Fig. 13. Same as 12 but for the last 6 landscapes in Table 2.

purposely geared towards facilitating the identification of crops, which are characterised by strong seasonality and (generally) low intra-field variability. To improve the method, or to extend it beyond croplands, it may be advisable to use another metric quantifying the noise independently from the amplitude of the signal. It may also be interesting to find a metric less dependent on the types of land covers encountered

(as shown by the differences in slope and r^2 for different land cover types in the simulation exercise). The approach can also be adapted with respect to the length and timing of the period over which it is applied. In this paper, the temporal window was focused on the growing period for every site. This allows one to maximize the signal of the vegetation growth and avoid the potential noise from fallow periods between two successive cropping seasons in which the surface would be expected to be more heterogeneous. But the SNR, or any other metric, could be applied on a moving window, on a yearly basis, or even on a multi-year period covering all the MODIS archive depending on the target application. In this respect, the framework behind the simulation exercise presented here could serve as a seedbed to explore performances of different metrics and on their application, such as under various scenarios of cloud coverage for example.

While pixel purity information tailored to specific applications is possible, various potential users could prefer a global product systematically available as ancillary data to other MODIS products. Such a product should probably be done using a moving window (compatible with that of other MODIS products such as albedo (Schaaf et al., 2002) enabling users to determine when the signal is pure for a given area. Seasonal and inter-annual patterns of pixel purity could be characterised in this way. This product should ideally be applied on a BRDF-corrected daily signal, perhaps using a pragmatic method such as that proposed by Vermote, Justice, and Bréon (2009)), and with a metric describing noise independently of the value of the signal (*i.e.* not SNR). For higher latitudes, it may also be worthwhile to use the non-gridded MODIS observations to increase the density of observations in the time series (when orbits overlap, only one value is added to the L2G grid). A yearly product using SNR could be envisaged to identify, at 250 m pixel level, areas with strong seasonality that could serve to improve current phenology products.

Regarding its use for agricultural monitoring, a couple of improvements or extensions to the present demonstration exercise could be foreseen. The SNR threshold to select spatially the time series to classify was fixed to 10 to allow a comparison across landscapes. Instead of fixing it (rather arbitrarily), it probably should be adjusted on a per landscape basis, or even avoided altogether by selecting only the 10% or 25% time series with the highest SNR values (with a minimum acceptable SNR to avoid collecting samples in landscapes with very mixed time series). The selection of the time period over which the SNR is calculated could be refined. Studies (*e.g.* Löw & Duveiller, 2014) have shown how the effectiveness of crop identification evolves along the season. This could be evaluated with SNR calculated over increasingly larger temporal windows to find the optimum. The number of clusters could also be changed per landscape, perhaps increased for some with higher landscape diversity such as North Dakota (US3) to avoid merging groups unnecessarily with the silhouette approach. The classification itself may be done differently, steered for instance by an *a priori* knowledge of the curve, using ancillary spectral information, or by focusing on the separability criteria to some critical moments in the time profile. This may aid in separated soybean from maize in Iowa (US2), for instance.

A general caveat with the method applied to crop monitoring is the limit imposed by the spatial resolution of the MODIS instruments. As shown by the field size maps of (Fritz et al., 2015), there are many parts of the world in which fields are small to very small, which the authors define as those that would require Landsat and very high-resolution imagery to be monitored. The method proposed here probably ensures MODIS can be used not only to large fields, as suggested by (Fritz et al., 2015), but also to their medium-sized fields. Major grain producing zones are often covered by such large or medium-sized fields. Notable exceptions are India and China where fields are very small, but the Indian and Chinese sites chosen in the present study suggest that adjacent small fields, with the same crops and with similar agro-management practices, may very well act as a larger field in with pure MODIS pixels can be found. Still, global agriculture monitoring should also target regions with small fields characterised with strong inter-field heterogeneity for which no pure MODIS pixel can be found. The method proposed here may be useful to delineate these areas, in order to prioritize where to focus efforts of using higher spatial resolution. Finally, there is also the case of landscapes containing both small and large fields of the same crops. In this case, it would be wise to analyse if no bias is committed when using only the purer pixels falling in the larger fields as an estimation of the average crop behaviour at regional level.

New observation systems will provide decametric spatial resolution with a high (*circa* 5 day) temporal revisit frequency globally. Such instruments (*e.g.* Sentinel 2 and Landsat 8) should allow to partially overcome many of the issues associated with pixel purity, but they will not have a long archive comparable to what MODIS can offer for years to come. Having comparable data for the past is a prerequisite for many applications such as detecting anomalies (Meroni, Verstraete, Rembold, Urbano, & Kayitakire, 2014), coupling remote sensing with crop models (Duchemin, Maisongrande, Boulet, & Benhadj, 2008; de Wit, Duveiller, & Defourny, 2012), statistically deriving yield estimates from remote-sensing indicators (Becker-Reshef, Vermote, Lindeman, & Justice, 2010; Kouadio et al., 2012; López-Lozano et al., 2015) or evaluating the impact of climate variability and trends on crop phenology and yield (Brown, de Beurs, & Marshall, 2012; Vrieling, de Leeuw, & Said, 2013; Lobell, Ortiz-Monasterio, Sibley, & Sohu, 2013). The proposed method allows the exploitation of the MODIS archive that started in 2000 and will probably still continue for several years more.

Can the method be applied to other satellite instruments than MODIS? The prerequisite to the approach is the possibility of sampling the local vicinity of an observation. The whiskbroom design of the MODIS sensor, combined with its large VZA and high temporal revisit capacity, helps sampling effectively that region. Both the precursor and successor of MODIS, namely AVHRR and VIIRS, have a similar design that should allow one to use the proposed approach as is. For pushbroom scanners like SPOT-VEGETATION, MERIS, PROBA-V and Sentinel3-OLCI, in which a line of sensors arranged perpendicular to the track direction scan the surface as the platform moves along its orbit, the footprint of consecutive observations over the same area change less than with whiskbroom scanners. However, as shown for MERIS by Gómez-Chova et al. (2011), the non-overlap between consecutive observations may still be large due to gridding artefacts and limits in geometric correction performance. This may be enough to exploit this "problem" to assess the pixel purity from pushbroom instruments using the approach presented here. In last recourse, it is also possible to use the 3 by 3 vicinity around every pixel to sample the vicinity, but this neighbourhood might be too large to cover homogeneous landscape elements for some applications.

6. Conclusion

Two main conclusions can be drawn from this work. First, that innovative use of currently available remote sensing data can provide much more information than what is often considered. The daily variability of the MODIS observation footprint, which is often either ignored or considered a drawback, is exploited to provide a clear indicator of the adequacy between the MODIS time series and a spatially homogeneous target. This is possible by using daily observation rather than composites, which in turn is becoming increasingly easier given the constantly increasing capacities in terms of computing power and storage space.

The second conclusion is that extracting crop specific signals at regional scale from MODIS is possible for many landscapes without requiring ancillary spatial data, such as crop maps, by exploiting pixel purity maps derived from the proposed method. Exploiting the full MODIS archive to derive crop specific information in this way should open new avenues for regional to global agricultural monitoring applications. Furthermore, the possible interest in this method applied to the MODIS archive goes beyond just crop monitoring, and should benefit many applications based on satellite Earth Observation.

Acknowledgement

The research in this paper was funded by the Joint Research Centre of the European Commission under the institutional projects AGRI4CAST (N° 1016) and CLIMEcoS (N° 995). The authors would like to thank the anonymous reviewers for their contribution to improving the manuscript. The authors also thank those responsible for the efforts on providing the free tools (R, Q-GIS) and datasets used in this work. The R packages were obtained from the Comprehensive R Archive Network (http://cran.r-project.org/) or R-forge (https://r-forge.r-project.org/). The MODIS and Landsat data were obtained through the online Data Pool at the NASA Land Processes Distributed Active Archive Center (LP DAAC), USGS/Earth Resources Observation and Science (EROS) Center, Sioux Falls, South Dakota (https://lpdaac.usgs.gov/get_data). The SPOT4 imagery was obtained under the SPOT4/Take5 programme. Imagery is copyrighted to CNES under the mention: "© CNES 2013, all rights reserved. Commercial use of the product prohibited". The Cropland Data Layer is made available by the USDA National Agricultural Statistics Service at http://nassgeodata.gmu.edu/CropScape/.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version, at http://dx.doi.org/10.1016/j.rse.2015.06.001. These data include Google maps of the most important areas described in this article.

References

Ali, A., de Bie, C., Skidmore, A., Scarrott, R., & Lymberakis, P. (2014). Mapping the heterogeneity of natural and semi-natural landscapes. *International Journal of Applied Earth Observation and Geoinformation*, *26*, 176–183. http://dx.doi.org/10.1016/j.jag.2013.06.007.

Atzberger, C. (2013). Advances in remote sensing of agriculture: Context description, existing operational monitoring systems and major information needs. *Remote Sensing*, *5*, 949–981. http://dx.doi.org/10.3390/rs5020949.

Baret, F., Weiss, M., Lacaze, R., Camacho, F., Makhmara, H., Pacholcyzk, P., et al. (2013). GEOV1: LAI and FAPAR essential climate variables and FCOVER global time series capitalizing over existing products. Part1: Principles of development and production. *Remote Sensing of Environment*, *137*, 299–309. http://dx.doi.org/10.1016/j.rse.2012.12.027.

Becker-Reshef, I., Vermote, E., Lindeman, M., & Justice, C. (2010). A generalized regression-based model for forecasting winter wheat yields in Kansas and Ukraine using MODIS data. *Remote Sensing of Environment*, *114*, 1312–1323. http://dx.doi.org/10.1016/j.rse.2010.01.010.

Brown, M., de Beurs, K., & Marshall, M. (2012). Global phenological response to climate change in crop areas using satellite remote sensing of vegetation, humidity and temperature over 26 years. *Remote Sensing of Environment*, *126*, 174–183. http://dx.doi.org/10.1016/j.rse.2012.08.009.

Cescatti, A., Marcolla, B., Santhana Vannan, S.K., Pan, J.Y., Román, M.O., Yang, X., et al. (2012). Intercomparison of MODIS albedo retrievals and in situ measurements across the global FLUXNET network. *Remote Sensing of Environment*, *121*, 323–334. http://dx.doi.org/10.1016/j.rse.2012.02.019.

Chen, J. M. (1999). Spatial scaling of a remotely sensed surface parameter by contexture. *Remote Sensing of Environment*, *69*, 30–42. http://dx.doi.org/10.1016/S0034-4257(99)00006-1.

Chen, J., Menges, C., & Leblanc, S. (2005). Global mapping of foliage clumping index using multi-angular satellite data. *Remote Sensing of Environment*, *97*, 447–457. http://dx.doi.org/10.1016/j.rse.2005.05.003.

Cracknell, A.P. (1998). Synergy in remote sensing—What's in a pixel? *International Journal of Remote Sensing*, *19*, 2025–2047.

de Bie, C., Nguyen, T.T.H., Ali, A., Scarrott, R., & Skidmore, a. K. (2012). LaHMa: A landscape heterogeneity mapping method using hyper-temporal datasets. *International Journal of Geographical Information Science*, *26*, 2177–2192. http://dx.doi.org/10.1080/13658816.2012.712126.

de Wit, A.J., Duveiller, G., & Defourny, P. (2012). Estimating regional winter wheat yield with WOFOST through the assimilation of green area index retrieved from MODIS observations. *Agricultural and Forest Meteorology*, *164*, 39–52. http://dx.doi.org/10.1016/j.agrformet.2012.04.011.

Duchemin, B., Maisongrande, P., Boulet, G., & Benhadj, I. (2008). A simple algorithm for yield estimates: Evaluation for semi-arid irrigated winter wheat monitored with green leaf area index. *Environmental Modelling & Software*, *23*, 876–892.

Duveiller, G. (2012). Caveats in calculating crop specific pixel purity for agricultural monitoring using MODIS time series. In C.M.U. Neale, & A. Maltese (Eds.), *Ecosystems, and Hydrology XIV, 24–27 September 2012. Proc. of SPIE: Remote Sensing for Agriculture.* (pp. 85310J). http://dx.doi.org/10.1117/12.974625 (Edinburgh, UK.).

Duveiller, G., Baret, F., & Defourny, P. (2011). Crop specific green area index retrieval from MODIS data at regional scale by controlling pixel-target adequacy. *Remote Sensing of Environment*, *115*, 2686–2701. http://dx.doi.org/10.1016/j.rse.2011.05.026.

Duveiller, G., Baret, F., & Defourny, P. (2012). Remotely sensed green area index for winter wheat crop monitoring: 10-Year assessment at regional scale over a fragmented landscape. *Agricultural and Forest Meteorology*, *166–167*, 156–168. http://dx.doi.org/10.1016/j.agrformet.2012.07.014.

Duveiller, G., & Lopez-Lozano, R. (2013). Exploiting MODIS observation geometry to identify crop specific time series for regional agricultural monitoring. In L. Ouwehand (Ed.), *Proc. 4th ESA Living Planet Symposium, (ESA SP-722) held on 9–13 September 2013 in Edinburgh*. Noordwijk, The Netherlands: ESA Communications ESTEC.

Efford, M.G. (2015). *secr: Spatially explicit capture-recapture models. R package version 2.9.3.* (http://CRAN.R-project.org/package=secr.).

Fleming, S.W. (2010). Signal-to-noise ratios of geophysical and environmental time series. *Environmental and Engineering Geoscience*, *16*, 389–399. http://dx.doi.org/10.2113/gseegeosci.16.4.389.

Foody, G. M. (1996). Incorporating mixed pixels in the training, allocation and testing stages of supervised classifications. *Pattern Recognition Letters*, *17*, 1389–1398. http://dx.doi.org/10.1016/S0167-8655(96)00095-5.

Fritz, S., See, L., McCallum, I., You, L., Bun, A., Moltchanova, E., et al. (2015). Mapping global cropland and field size. *Global Change Biology*, *21*, 1980–1992. http://dx.doi.org/10.1111/gcb.12838.

Garrigues, S., Allard, D., Baret, F., & Weiss, M. (2006). Influence of landscape spatial heterogeneity on the non-linear estimation of leaf area index from moderate spatial resolution remote sensing data. *Remote Sensing of Environment*, *105*, 286–298. http://dx.doi.org/10.1016/j.rse.2006.03.013.

Gómez-Chova, L., Zurita-Milla, R., Alonso, L., Amorós-lópez, J., Guanter, L., & Camps-valls, G. (2011). Gridding artifacts on medium-resolution satellite image time series: MERIS case study. *IEEE Transactions on Geoscience and Remote Sensing*, *49*, 2601–2611.

Jacob, F., & Weiss, M. (2014). Mapping biophysical variables from solar and thermal infrared remote sensing: Focus on Agricultural Landscapes with spatial heterogeneity. *IEEE Geoscience and Remote Sensing Letters*, *11*, 1844–1848. http://dx.doi.org/10.1109/LGRS.2014.2313592.

Justice, C.O., & Becker-Reshef, I. (2007). *Report from the Workshop on Developing a Strategy for Global Agricultural Monitoring in the framework of Group on Earth Observations (GEO).* Technical Report*Geography Dept.* University of Maryland.

Kouadio, L., Duveiller, G., Djaby, B., El Jarroudi, M., Defourny, P., & Tychon, B. (2012). Estimating regional wheat yield from the shape of decreasing curves of green area index temporal profiles retrieved from MODIS data. *International Journal of Applied Earth Observation and Geoinformation*, *18*, 111–118. http://dx.doi.org/10.1016/j.jag.2012.01.009.

Kuusinen, N., Tomppo, E., & Berninger, F. (2013). Linear unmixing of MODIS albedo composites to infer subpixel land cover type albedos. *International Journal of Applied Earth Observation and Geoinformation*, *23*, 324–333. http://dx.doi.org/10.1016/j.jag.2012.10.005.

Lobell, D.B., Ortiz-Monasterio, J.I., Sibley, A.M., & Sohu, V. (2013). Satellite detection of earlier wheat sowing in India and implications for yield trends. *Agricultural Systems*, *115*, 137–143. http://dx.doi.org/10.1016/j.agsy.2012.09.003.

López-Lozano, R., Duveiller, G., Seguini, L., Meroni, M., Garca-Condado, S., Hooker, J. D., Leo, O., & Baruth, B. (2015). Towards regional grain yield forecasting with 1 km-resolution EO biophysical products: strengths and limitations at pan-European level. *Agricultural and Forest Meteorology*, *206*, 12–32. http://dx.doi.org/10.1016/j.agrformet.2015.02.021.

Löw, F., & Duveiller, G. (2014). Defining the spatial resolution requirements for crop identification using optical remote sensing. *Remote Sensing*, *6*, 9034–9063. http://dx.doi.org/10.3390/rs6099034.

Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., & Hornik, K. (2014). *Cluster: Cluster analysis basics and extensions.* R package version 1.15.2.

Mattiuzzi, M. (2014). MODIS acquisition and processing package for R. Version: 0.10–11. https://r-forge.r-project.org/R/?group_id=1252

McGarigal, K., Cushman, S., & Ene, E. (2012). *FRAGSTATS v4: Spatial pattern analysis program for categorical and continuous maps.*

Meroni, M., Verstraete, M.M., Rembold, F., Urbano, F., & Kayitakire, F. (2014). A phenology-based method to derive biomass production anomalies for food security monitoring in the Horn of Africa. *International Journal of Remote Sensing*, *35*, 2472–2492. http://dx.doi.org/10.1080/01431161.2014.883090.

Mu, Q., Zhao, M., & Running, S.W. (2011). Improvements to a MODIS global terrestrial evapotranspiration algorithm. *Remote Sensing of Environment*, *115*, 1781–1800. http://dx.doi.org/10.1016/j.rse.2011.02.019.

Myneni, R.B., Hoffman, S., Knyazikhin, Y., Privette, J.L., Glassy, J., Tian, Y., et al. (2002). Global products of vegetation leaf area and fraction absorbed PAR from year one of MODIS data. *Remote Sensing of Environment*, *83*, 214–231.

Pinty, B., Widlowski, J.J.L., Gobron, N., Verstraete, M., & Diner, D. (2002). Uniqueness of multiangular measurements. I. An indicator of subpixel surface heterogeneity from MISR. *Geoscience and Remote Sensing, IEEE Transactions on*, *40*, 1560–1573. http://dx.doi.org/10.1109/TGRS.2002.801148.

Pittman, K., Hansen, M.C., Becker-Reshef, I., Potapov, P.V., & Justice, C.O. (2010). Estimating global cropland extent with multi-year MODIS data. *Remote Sensing*, *2*, 1844–1863. http://dx.doi.org/10.3390/rs2071844.

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65. http://dx.doi.org/10.1016/0377-0427(87)90125-7.

Saura, S., & Martinez-Millan, J. (2000). Landscape patterns simulation with a modified random clusters method. *Landscape Ecology*, *15*, 661–678. http://dx.doi.org/10.1023/A:1008107902848.

Schaaf, C.B., Gao, F., Strahler, A.H., Lucht, W., Li, X., Tsang, T., et al. (2002). First operational BRDF, albedo nadir reflectance products from MODIS. *Remote Sensing of Environment*, *83*, 135–148. http://dx.doi.org/10.1016/S0034-4257(02)00091-3.

Schowengerdt (2007). *Remote sensing: Models and methods for image processing* (3rd ed.). San Diego: Academic Press.

Tan, B., Woodcock, C.E., Hu, J., Zhang, P., Ozdogan, M., Huang, D., et al. (2006). The impact of gridding artifacts on the local spatial properties of MODIS data: Implications for validation, compositing, and band-to-band registration across resolutions. *Remote Sensing of Environment*, *105*, 98–114.

Vermote, E., Justice, C.O., & Bréon, F.M. (2009). Towards a generalized approach for correction of the BRDF effect in MODIS directional reflectances. *IEEE Transactions on Geoscience and Remote Sensing*, *47*, 898–908.

Vrieling, A., de Leeuw, J., & Said, M. (2013). Length of growing period over Africa: Variability and trends from 30 years of NDVI time series. *Remote Sensing*, *5*, 982–1000. http://dx.doi.org/10.3390/rs5020982.

Wan, Z. (2008). New refinements and validation of the MODIS land-surface temperature/emissivity products. *Remote Sensing of Environment*, *112*, 59–74. http://dx.doi.org/10.1016/j.rse.2006.06.026.

Whitcraft, A.K., Becker-Reshef, I., & Justice, C.O. (2014). Agricultural growing season calendars derived from MODIS surface reflectance. *International Journal of Digital Earth*, 1–25. http://dx.doi.org/10.1080/17538947.2014.894147.

Wolfe, R.E., Nishihama, M., Fleig, A.J., Kuyper, J.A., Roy, D.P., Storey, J.C., et al. (2002). Achieving sub-pixel geolocation accuracy in support of MODIS land science. *Remote Sensing of Environment*, *83*, 31–49.

Wolfe, R.E., Roy, D.P., & Vermote, E. (1998). MODIS Land data storage, gridding, and compositing methodology: Level 2 grid. *Geoscience and Remote Sensing, IEEE Transactions on*, *36*, 1324–1338.

Woodcock, C.E., & Strahler, A. (1987). The factor of scale in remote sensing. *Remote Sensing of Environment*, *21*, 311–332.

Zhang, X., Friedl, M.A., Schaaf, C.B., Strahler, A.H., Hodges, J.C., Gao, F., et al. (2003). Monitoring vegetation phenology using MODIS. *Remote Sensing of Environment*, *84*, 471–475. http://dx.doi.org/10.1016/S0034-4257(02)00135-9.