Field Crops Research 143 (2013) 56–64

Contents lists available at SciVerse ScienceDirect

Field Crops Research

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

The use of satellite data for crop yield gap analysis

David B. Lobell

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

ARTICLE INFO

Article history:
Received 19 February 2012
Received in revised form 18 July 2012
Accepted 8 August 2012

Keywords:
Remote sensing
GIS
Agronomy

ABSTRACT

Field experiments and simulation models are useful tools for understanding crop yield gaps, but scaling up these approaches to understand entire regions over time has remained a considerable challenge. Satellite data have repeatedly been shown to provide information that, by themselves or in combination with other data and models, can accurately measure crop yields in farmers' fields. The resulting yield maps provide a unique opportunity to overcome both spatial and temporal scaling challenges and thus improve understanding of crop yield gaps. This review discusses the use of remote sensing to measure the magnitude and causes of yield gaps. Examples from previous work demonstrate the utility of remote sensing, but many areas of possible application remain unexplored. Two simple yet useful approaches are presented that measure the persistence of yield differences between fields, which in combination with maps of average yields can be used to direct further study of specific factors. Whereas the use of remote sensing may have historically been restricted by the cost and availability of fine resolution data, this impediment is rapidly receding.

1. Introduction

Two goals underlie most discussions of crop yield gaps (Van Ittersum et al., 2013). The first is to measure the size of the yield gap, defined as the difference between yield potential (Yp) and average yields, in order to identify the potential scope for raising average yields via management changes. The second is to identify the key causes of the yield gap, in order to prioritize efforts in extension, research, and policy to raise land and labor productivity.

A fundamental challenge in pursuit of either of these goals is the considerable spatial and temporal heterogeneity of agricultural landscapes. In the measurement of yield gaps, for example, actual yields are often reported for administrative units that span hundreds or thousands of fields. Yield potential, meanwhile, is most readily estimated for individual fields, using either agronomic trials or well-tested crop simulation models (Lobell et al., 2009). How should the measurements at these two different spatial scales be compared when computing a yield gap? Some studies ignore the scale mismatch, implicitly assuming that point-level estimates of Yp are a good proxy for average Yp across the spatial domain of the reported average yield. Other studies attempt to estimate Yp at multiple points within the domain and then take an average, a sensible approach provided that data of sufficient quality exist to estimate Yp at multiple points.

Similarly, studies to understand causes of the yield gap may reasonably start by evaluating yield responses to different management changes in an experimental station, or on farmers' fields. However, the fields analyzed may not be representative of the entire region, or the year(s) in which the study was done may not be representative of the full range of conditions that farmers face.

Agronomists have long appreciated the challenge of generalizing results from a small handful of sites and years to the broader scales relevant to regional measures of performance. Over the past two decades, remote sensing has emerged as a useful tool for dealing with heterogeneity, to complement more traditional approaches such as field trials or simulation models. In particular, remote sensing from airplane- or satellite-mounted sensors can potentially provide observations for every single field in a region for every single growing season. Although remote sensing-based estimates of quantities such as crop yield are often less accurate than field-based measures, the unprecedented spatial and temporal coverage of remote sensing can often outweigh the negatives for many applications.

The goal of this paper is to specifically address the potential value of satellite-based remote sensing for efforts to measure and explain crop yield gaps. The premise of the paper is that as efforts to understand yield gaps intensify, new approaches that can complement the traditional toolbox of agronomists have great potential value, and remote sensing may be one such tool.

The promise of satellite data is enhanced by at least two recent developments. One is the decision in 2008 by the United States Geological Survey to make the entire archive of Landsat data available at no charge (http://landsat.usgs.gov/products_data_at_no_charge.php). This change, coupled with improvements

E-mail address: dlobell@stanford.edu

0378-4290
http://dx.doi.org/10.1016/j.fcr.2012.08.008

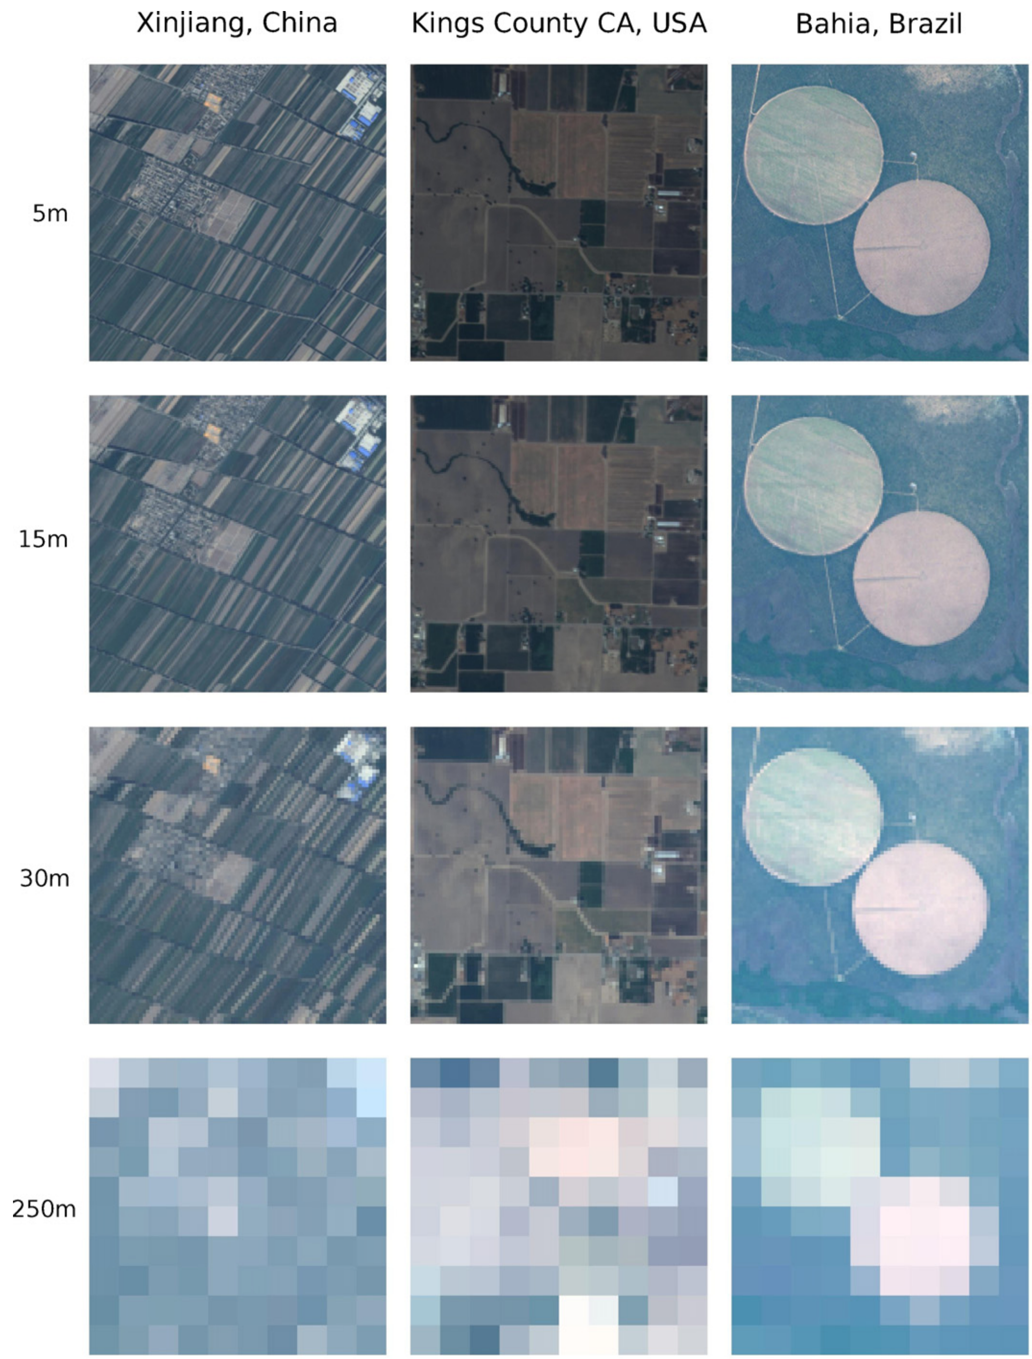

Fig. 1. Examples of the effect of image resolution on the ability of remote sensing to monitor individual fields. The top row displays a 2.5 km × 2.5 km section of a RapidEye 5 m × 5 m resolution image of (left) Xinjiang, China, (middle) Kings County, United States, and (right) Bahia, Brazil. Lower rows show images resampled to represent the coarser resolutions of some other common sensors: ASTER (15 m, second row), Landsat (30 m, third row), and MODIS (250 m, bottom row).

in preprocessing algorithms to geographically register the images, have vastly reduced the expense and time required to obtain images at relatively fine spatial resolution (30 m × 30 m) for the period of 1982 to present. This resolution is sufficient to delineate individual fields that are roughly 1 ha in size or greater, which includes many regions of the world (see Fig. 1). Second, new commercial systems are delivering even higher spatial resolution (5 m × 5 m or finer) at costs that are approaching 1 USD$ per km^2 (or $0.01 per ha). In the next decade, it should be increasingly feasible to obtain multiple years of data for regions where field sizes have been too small to distinguish with traditional sensors like Landsat.

The next section briefly summarizes the capabilities and limitations of remote sensing for measuring crop yields. The following two sections discuss examples and potential uses of remote sensing for the two main goals of crop yield analysis: measurement and explanation. Finally, some conclusions and recommendations for future work are presented.

2. Remote sensing of crop yield

2.1. Approaches

Numerous approaches exist for estimating crop yields with remote sensing. Several reviews on this topic are available (Moulin et al., 1998; Gallego et al., 2010), and so only a brief summary of approaches is given here. Early efforts relied on simple vegetation indices (VIs) computed from remote sensing measurements of light at red and near-infrared (NIR) wavelengths. Crop plants, like all vegetation, are very reflective in the NIR and absorptive at red wavelengths, and therefore some combination of the two are a good measure of vegetation vigor (Tucker, 1979; Sellers, 1987). Measures at these two wavelengths remain the mainstay of nearly all approaches to crop yield estimation, although other parts of the spectrum are commonly utilized in more sophisticated approaches (e.g., Gitelson et al., 2003).

The simplest approach to estimating crop yields is to establish empirical relationships between ground-based yield measures and VIs measured on a single date or integrated over the growing season. Early applications with wheat and maize indicated that variations in VI can explain over 80% of the observed variation in crop yields within individual fields (Tucker et al., 1980; Wiegand and Richardson, 1990; Shanahan et al., 2001). However, as with any purely empirical approach, extrapolation of equations to new locations or years can be problematic, and for this reason many efforts have been made toward more general techniques.

One class of models relies on the light-use efficiency approach pioneered by Monteith (1977), which states that total biomass production is directly proportional to total absorption of photosynthetically active radiation (PAR) over the course of the growing season. The ratio of biomass to PAR, known as radiation use efficiency (RUE), is relatively constant because plants adjust total leaf area, and thus capture of sunlight, in response to other growth constraints such as nutrient or temperature stress (Bloom et al., 1985). Numerous studies have confirmed the utility of this concept for predicting biomass in different crops, although variations in RUE do occur in some settings, particularly when plants experience acute moisture stress (Steinmetz et al., 1990; Sinclair and Muchow, 1999).

Models that rely on the RUE concept to predict yield have at least four required inputs, as indicated in Eq. (1):

$$\text{Yield} = \left(\sum_{t=1}^{n} \text{PAR}_t \cdot \text{fPAR}_t\right) \cdot \text{RUE} \cdot \text{HI} \tag{1}$$

where the summation indicates a sum over the growing season, often using a daily time step, fPAR_t is the fraction of PAR absorbed by the crop canopy at time t, and HI is the harvest index, or the ratio of yield to total biomass. The first input is PAR throughout the growing season, which is typically obtained from local meteorological station data or satellite-based estimates. Second, estimates of fPAR throughout the growing season are required. For remote sensing-based studies, estimates of fPAR are most often derived from established relationships with VIs. However, because remote sensing data are not available on a daily basis, some interpolation is needed to estimate daily fPAR.

Three general approaches are used to address the interpolation issue. In the case of regularly spaced satellite-based measures, such as the eight-day composites of fPAR from the Moderate Resolution Imaging Spectroradiometer (MODIS) (Myneni et al., 2002), the time step in Eq. (1) can simply be coarsened to match the satellite data. A second approach is to interpolate based on some statistical function, such as local linear interpolation or fitting double logistic curves (Nightingale et al., 2009). These two approaches can also be combined, for instance when trying to account for missing or low-quality data points within an eight-day product (Zhao et al., 2005; Nightingale et al., 2009).

Third, interpolation can be achieved by using a crop model to simulate the daily evolution of fPAR between available observations. This can be achieved, for example, by calibrating parameters of a simple crop growth model until the simulated fPAR values most closely match the remote sensing estimates (Lobell et al., 2003). This approach has the benefit of allowing flexibility in the timing of remote sensing data, and being able to provide estimates even when only two or three images are available during the growing season, as is often the case when using higher resolution imagery. Another feature is the ability of the crop model to account for effects of temperatures on crop development rates.

The final two factors in Eq. (1) are estimates of RUE and HI. Both of these are typically assumed constant, and derived from field data or calibration to reported statistics. In comparison to simple empirical relationships, RUE models can capture variations in yield due to changes in PAR, and interpolation of fPAR allows more flexibility in the timing of images relative to crop growth stage (whereas empirical methods are typically specific to a narrow range of timing).

Beyond simple RUE models, many approaches attempt to more fully integrate crop simulation models with remote sensing data. For example, Doraiswamy et al. (2004, 2005) adjusted parameters in a simple climate-based crop model to match MODIS estimates of leaf area index (LAI) over maize and soybean fields in the United States. Several studies have used different sensors, spatial scales, and crop models (Weiss et al., 2001; Prévot et al., 2003; Launay and Guerif, 2005; Dente et al., 2008), but all with the same principle of adjusting crop model parameters on a pixel-by-pixel basis to match remote sensing estimates of some crop attribute, often LAI. The simulated yield by the crop model then provides an estimate of yield for each pixel. (Note the contrast with the use of crop models only to interpolate fPAR (e.g., Lobell et al., 2003), with yield then calculated from Eq. (1).)

The use of crop simulation models allows the possibility of accounting for factors such as acute moisture stress, which can result in variations in RUE and HI that confound estimates using Eq. (1). However, this added complexity comes at a cost of requiring more model inputs, such as soil properties, and greater computation time.

2.2. Accuracies and limitations

Overall, there is clear evidence that crop yield estimation is possible with remote sensing, with good accuracies in some cases. Most evaluations of remote sensing are at scales broader than individual fields, for example by comparing reported yields for counties or crop reporting districts with the average of remotely sensed yields over this domain (Doraiswamy et al., 2005; Becker-Reshef et al., 2010; Lobell et al., 2010). Such comparisons will typically present overly optimistic views of the field-scale accuracies that are most relevant to yield gap assessments. As an example of the deterioration of accuracy at finer scales, Reeves et al. (2005) used 1 km MODIS data to estimate wheat yields in North Dakota and Montana, and found accuracies to within 5% for state level estimates but much lower accuracies at the county level.

Nonetheless, field level accuracies are sometimes quite high, with root mean square errors of less than 10% for predicting farmer-reported yields for individual commercial fields (Clevers, 1997; Lobell et al., 2005, 2007b). In most situations, non-negligible errors exist in the farmer-reported data, suggesting that the true accuracies of remote sensing data can be even higher. However, the conditions necessary for accurate yield estimation are not always met, and many steps in the process of yield estimation introduce errors. Even at coarse spatial scales, accuracies can be too low to provide useful results. It is telling, for example, that both the United States' and European agencies in charge of forecasting domestic and international crop production currently use remote sensing only in a qualitative fashion (Allen et al., 2002; Baruth et al., 2008).

Sources of error in yield estimation include potential misclassification of which crops are growing in which pixels, errors in estimating fPAR with reflectance data, difficulties in interpolating values between available observations, imperfect relationships between measures such as fPAR and final biomass, and difficulties in predicting variations in HI. The first of these issues, misclassification of crop type, is particularly problematic in regions that grow multiple crops with very similar phenologies, or in regions with intercropped fields. In the former case, crops can sometimes be distinguished based on spectral differences, but accuracies are generally low relative to cases that separate crops based on clearly distinct phenologies (Thenkabail, 2001). In the latter case, such as mixed-cropping systems found throughout Africa, both crop type and yield estimation are extremely difficult and to the author's knowledge there are no good examples of reliable yield estimates in these situations.

Other sources of error reflect the difficulty of predicting aspects of crop growth which are not related to absorption of PAR, such as RUE and HI. Recent studies have indicated that remote sensing based indicators of chlorophyll content are related to RUE, and thus incorporation of wavelengths associated with chlorophyll can improve estimates of instantaneous carbon uptake (Peng and Gitelson, 2011). However, these approaches have not yet been demonstrated to improve yield estimates. Harvest index variations are particularly hard to predict from remote sensing, and the main approach remains either assuming a constant HI or using crop models that predict changes in HI. Although irrigated cereal crops tend to have stable HI, rainfed and non-cereal crops can exhibit very low values of HI when exposed to stress (Hay, 1995). For example, past efforts by the author to estimate safflower yields were hampered by large variations in HI.

Further sources of error stem from inherent tradeoffs between acquiring data with sufficiently high spatial resolution to delineate individual fields, and data with sufficiently high temporal resolution to obtain multiple cloud-free observations during the growing season. Historically, Landsat (or similar sensors such as SPOT) has been the main source of data with sufficient spatial resolution in most agricultural areas, but with a 16-day gap between successive images, and frequent cloud cover in most cropping regions (with the exception of dry, irrigated areas), it can be difficult to obtain more than one or two clear images within a growing season. The main source of fine temporal resolution data has been MODIS, but with 250 m spatial resolution it is extremely difficult to identify individual fields in most regions (Fig. 1). The tradeoff between spatial and temporal resolution is most extreme in regions with small field sizes, for which even Landsat cannot resolve individual fields.

Moving forward, the lower cost and increased number of fine spatial resolution sensors should help to partly mitigate the inherent tradeoff between spatial and temporal resolution. Another promising development has been recent progress in data fusion methods that combine observations from different sensors. One approach, termed the spatial and temporal adaptive reflectance fusion model (STARFM) uses a pair of fine and coarse resolution images (e.g., Landsat and MODIS) acquired on the same day, along with coarse images acquired on other dates, to generate high temporal resolution data at the finer spatial scale (Gao et al., 2006). An important step in this approach is to identify, for each pixel, spectrally similar pixels within a neighborhood, which are then assumed to change similarly over time. An enhanced STARFM (ESTARFM) was also introduced recently (Zhu et al., 2010), which relies on a pair of fine resolution images to define similarity. A third approach to fusing fine and coarse resolution images was introduced by Zurita-Milla et al. (2009), which relies on modeling each pixel's reflectance as the linear sum of reflectance from individual land cover types within the pixel. The critical assumptions in this method are that a reliable map of land cover types at fine resolution can be obtained and that all pixels of the same land cover type behave identically within a window around the central pixel.

These approaches have been tested in a few cases with promising results (Gao et al., 2006; Hilker et al., 2009), but have yet to be rigorously evaluated in agricultural settings. All three approaches assume that nearby pixels that appear similar at a single date behave similarly thereafter. This assumption is problematic in agricultural settings, where for example fields are sown at different times and therefore may appear similar at one date but very different in another. Nonetheless, data fusion methods are likely to improve over the next few years, and offer opportunities for improving yield estimation at field scales.

3. Measurement of yield gaps

Yield gap estimation requires two quantities: Yp and actual yields. Given that remote sensing can only assess actual yields, additional information or assumptions about Yp are required to estimate yield gaps. One simple approach would be to pair independent estimates of Yp (for example based on methods used in various papers in this special issue) with remote sensing-based estimates of actual yields for the corresponding sites and locations. No examples of this exist, to the author's knowledge, presumably because average yields are already known in cases where Yp estimates are available.

An alternative is to use the maximum yield within the remote sensing estimates as a proxy for Yp, which assumes that some farmers in a given year achieve Yp. This assumption has been tested in some cases (Lobell et al., 2009), but in general is made without independent data to test its validity (if such data existed, the assumption would not be necessary). Early examples of this approach include a study of various crops in Pakistan using 1.1 km AVHRR data (Bastiaanssen and Ali, 2003), and a study of wheat in Mexico using Landsat (Lobell et al., 2002). In both cases, the authors did not use the single maximum value of yields in the region, but instead define the gap using a high (e.g., 95th) percentile of the yield distribution.

In both cases, it was also assumed that a single value of Yp was applicable to the entire study region. This assumption is likely a good approximation for relatively small study regions, but otherwise could introduce significant errors into the analysis. There is no inherent reason, however, why a single value of Yp must be used for the entire study region. For example, one could compute Yp for each pixel based on the maximum or 95th percentile of yield observed in a small surrounding region (e.g., a 5 km^2 area centered on the pixel).

Such an approach, though, would still rely on the assumption that the highest yielding fields are approaching Yp. A hybrid approach could use independent estimates of Yp for specific points within the study region, obtained for instance using crop simulation models. The remote sensing estimates of Yp, obtained by the moving window approach described above, could then be compared to these point estimates and a regression equation defined to adjust remote sensing estimates to match the point estimates. Thus, one could leverage the ability of remote sensing to detect spatial gradients in Yp without relying exclusively on the assumption that some farmers achieve yields close to Yp.

In general, very little effort has been devoted to using remote sensing to measure the magnitude of yield gaps. There are some reasons to believe satellite data can complement point-level estimates of Yp from other methods, or even in some cases circumvent the need for other methods altogether. However, much more work is needed to establish the true utility of remote sensing in this application.

4. Analysis of yield gap causes

Much more common than work focused explicitly on the size of yield gaps has been satellite based studies of spatial variation in observed yields. A substantial amount of work has been devoted to mapping and interpreting yield variations *within* a single field in order to inform precision agriculture applications (Yang et al., 2001; Zarco-Tejada et al., 2005). Of more relevance here, however, are studies that span a large number of fields within a region.

Work to understand causes of spatial yield variations has an obvious but potentially limited relevance to the question of what causes yield gaps. In particular, if all fields in a region are well below Yp, then explaining variations between the fields will not necessarily help to explain how to raise average yields close to Yp. In such cases, however, it is very likely that at least the main proximate causes of the yield gap are already known, such as the case of insufficient fertilizer inputs and soil fertility throughout much of Africa. In these instances it is also likely that variations between fields will be dominated by the same constraint that is so binding to overall productivity (Tittonell et al., 2008).

In general, two types of approaches have characterized studies of landscape yield heterogeneity. First, maps of yields derived from remote sensing can be compared to ancillary datasets on factors thought to control yields, such as soil properties or management practices. Statistical analyses can then be used to evaluate the relative importance of each factor in driving observed yield variations. An important step in such studies is accounting for spatial correlation of model errors: given the high spatial density of remote sensing measurements, it is relatively easy to underestimate standard errors if spatial correlation is ignored.

A series of studies conducted in the Yaqui Valley by the author and colleagues provide several examples of this approach, with comparisons of satellite-derived yield data to management surveys (Lobell et al., 2005), sowing date and weed maps (Ortiz-Monasterio and Lobell, 2007), and datasets on irrigation deliveries (Lobell and Ortiz-Monasterio, 2008). Studies in other regions include effects of position along irrigation canals for rice in Mali (Zwart and Leclert, 2010), effects of distance to main irrigation canals and roads on wheat yields in Punjab, India (Lobell et al., 2010), and effects of soil sodicity on wheat yields in Mexicali, Mexico (Seifert et al., 2011).

In all of the above cases, statistically significant yield effects of the variables in question were detected. Moreover, because of the large sample sizes afforded by remote sensing, it is possible to identify important interactions, nonlinearities, or thresholds that are not uncovered with traditional linear regression techniques applied to smaller datasets. For example, Seifert et al. (2011) detected a yield decline beginning at values of soil exchangeable sodium percentage (ESP) values of 6, whereas the common definition of sodic soils is for ESP above 15.

A second approach to studying yield gaps with remote sensing relies exclusively on spatial and temporal patterns in the yield maps themselves. Although these approaches cannot provide insight into specific yield controls, and are thus not sufficient by themselves, they provide a relatively rapid assessment of what types of factors deserve further scrutiny. That is, different spatial and temporal patterns in the data will suggest different types of potential yield controls. For example, if farmer characteristics like technical knowledge or access to credit can explain much of the yield variation, then one would expect to see yields consistently high year after year in fields where farmers have the favorable traits. Similarly, if soil quality and degradation present major obstacles to yield improvement, and these soil properties are not themselves driven by management (Tittonell et al., 2008), one might expect to see yield patterns that reflect relatively smooth variation of soil properties, as opposed to the abrupt transitions in management across field boundaries (Lobell and Ortiz-Monasterio, 2006). In contrast, if the yield responses to management or soil properties depend strongly on weather conditions, one would expect a more variable spatial pattern of yields from year to year.

One simple way to gain insight into yield gap causes is to examine the spatial distribution of average yields, with the average calculated over varying lengths of time. In general, averages calculated over longer periods of time will show less spatial variation than averages over shorter periods, since factors that are idiosyncratic will tend to cancel out across years. An example in Fig. 2 for wheat yields in the Yaqui Valley, Mexico illustrates this effect. For a single year, a relatively wide (and asymmetric) distribution of yields is observed with many locations (i.e., 30 m × 30 m pixels) achieving both comparatively high (e.g., 20% above average) and low (e.g., 20% below average) yields. The average yields over eight seasons (shown only for fields that had wheat in at least five of these seasons), displays a much narrower distribution, with virtually no locations achieving 20% above the regional average and much fewer laggards at 20% below average.

This visual impression can be formalized by extracting some key statistics of the yield distribution, such as the difference between the maximum yields (or 95th percentile) and average yields. We can refer to this quantity as YGM_L, where YGM refers to yield gap relative to maximum yields and the subscript L references the length of the record (in # of seasons) used to compute the average yields for each pixel. Values of YGM_L for L varying from one to the length of the record can then be plotted on a "yield gap curve," as outlined Fig. 3. The steepness of this curve then provides insight into how persistent spatial yield differences are throughout the study period, and thus how important persistent factors like soil quality or farmer skill are in explaining the overall yield gap.

Fig. 4 presents yield gap curves computed for three different regions where at least six seasons of wheat yield estimates have been generated based on Landsat (data taken from Lobell et al., 2007a, 2010). For reference, each plot also includes a gray line indicating the shape of the curve expected if one randomly scrambles yields across space for each image. This gray line thus represents the null distribution if yield patterns were entirely inconsistent (i.e., random) from year to year, and the distance between the two lines provides an indication of the overall persistence in spatial yield patterns.

Evident in Fig. 4 are important differences between the three locations, with the Yaqui Valley possessing the steepest curve that closely resembles the null curve. Sangrur, India shows a shallower curve and Bhatinda, India is shallower still. By themselves, these curves suggest that there are very few persistent factors explaining the yield gap in Yaqui Valley, in contrast to the Indian sites where at least some persistence is evident.

These impressions are supported by more detailed analyses in each region. In Yaqui Valley, farmers are relatively advanced in their understanding of crop management and have fairly reliable access to irrigation and chemical inputs. Comparisons with field surveys in two years indicated that different management factors become important in different years, namely timing of irrigation in one year and amount of fertilizer inputs in another (Lobell et al., 2005). Of course, a lack of persistence in yield patterns does not suggest that farmer skill or access is not important, but that *variations* in these factors are not sufficient to explain yield differences, presumably because of uniformly high levels of famer skill and access.

In Sangrur, India, comparisons with ancillary datasets on irrigation indicate that farmers at the tail end of irrigation canals have consistently lower yields (Lobell et al., 2010). Although there are several possible explanations for this trend, the most likely is less reliable access to surface water. A similar explanation applies to Bhatinda, except the magnitude of this effect is greater because that district is much more reliant on surface canals relative to groundwater than Sangrur (Lobell et al., 2010). Moreover, a substantial fraction of land in southwestern Bhatinda is routinely planted late because of cotton cultivation in the summer that delays the wheat season. Thus, persistent variations in both sowing date and position along irrigation canals contribute to the shallower slope of the yield gap curve in Bhatinda compared to the other two cases.

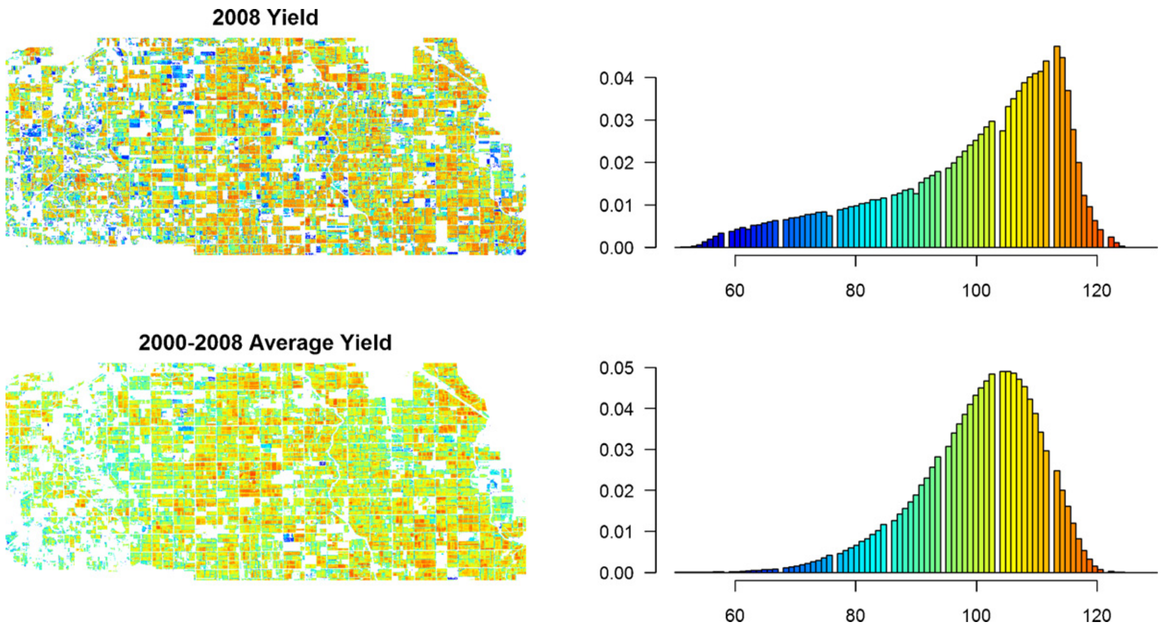

Fig. 2. The effect of multi-year averaging on spatial yield patterns in part of the Yaqui Valley, Mexico. (a) Map and (b) histogram of Landsat-based yield estimates for wheat for the 2007–2008 growing season. Values are expressed as a percentage of the mean yield estimate for the entire region (mean yield = 100). (c) Map and (d) histogram of the average yield estimates for eight growing seasons between 1999 and 2008, for the same locations. Only fields with wheat in at least five growing seasons are included. The distribution of yields tends to narrow as one averages yields over longer time periods, because some of the factors driving spatial differences in any single year are not persistent.

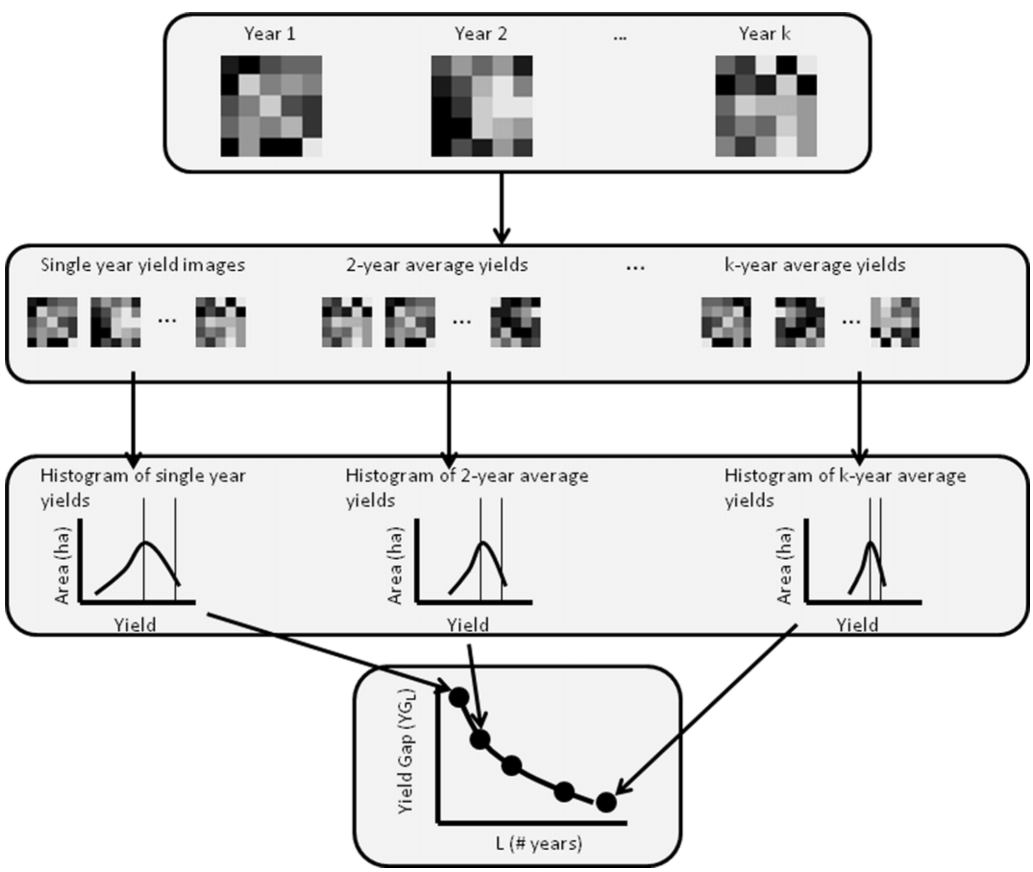

Fig. 3. Summary of procedure to compute yield gap profiles. Images from multiple years are averaged to create maps of average yields for varying periods of time. The maps are then used to compute the difference between maximum and average yielding fields, and this difference is plotted versus the number of years used in the average.

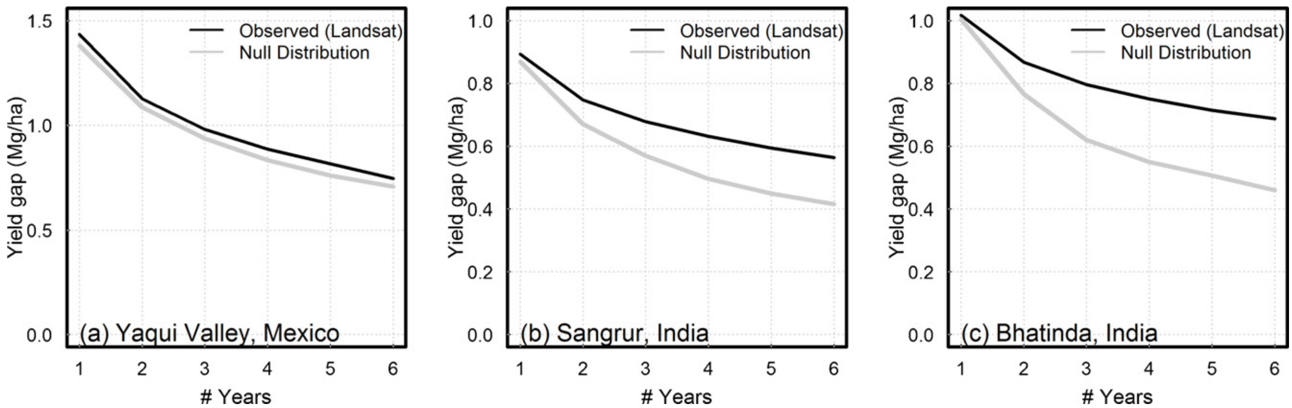

Fig. 4. Yield gap profiles for wheat in different regions. (a) Yaqui Valley District, Mexico, (b) Sangrur District, India, (c) Bhatinda District, India. Each point shows the difference between maximum yield and average yield within the district, where yields are calculated as the average over 1–6 years. Differences diminish in time because of a lack of persistence in spatial patterns. The gray line represents the expected change in yield gap with increasing years if yield patterns were entirely random in space (computed by randomly re-ordering the spatial distribution of yields in each year).

The yield gap curves in Fig. 4 are just one way to summarize the persistence of yield differences. Another simple summary is shown in Fig. 5, which splits fields into 10 groups according to the yield deciles (0–10%, 10–20%, etc.) for the most recent yield image. The figure summarizes the yield distribution for each of these groups across all years except the year used to define the groups. The boxes show the interquartile range of the data (25th and 75th percentile), the horizontal line indicates the median, and the whiskers show the 10th and 90th percentiles of the data. The significant overlap between the distributions indicates again the lack of strong persistence in yield patterns; fields with very different yields in a single year have relatively similar yield distributions in other years. The degree of overlap is greatest in the case of the Yaqui Valley, which similar to Fig. 4 indicates the least persistence in this region relative to the Indian sites. At the same time, all three sites exhibit a positive slope in Fig. 5, which indicates a tendency for high yielding fields to remain in the high yielding category, and for fields in the lowest decile of yields to remain relatively low yielding.

An important issue when analyzing spatial yield distributions is setting the extent of the analysis. In the above examples, the boundaries of the crop districts were used to define the extent. If the extent had been larger, such as an analysis for the entire state, the amount of persistent yield differences would likely increase as soil, climate, and management conditions at opposite ends of the extent would likely grow. For yield gap analysis, it seems preferable to keep the extent small enough such that variations in climate, and thus yield potential, across the domain are small, but large enough such that a large number of fields are included in the sample. However, the sensitivity of results to choosing different spatial extents has not been considered in past studies, and is a topic worthy of more research.

The data required to assemble Figs. 4 and 5 are substantial: six years of Landsat data, with 2–3 images per year for each figure panel. However, the methods to process such data, including radiometric and geometric correction of the remote sensing data and incorporation of temperature and radiation data to estimate yields, can be increasingly automated. In practice, the main obstacle is typically performing image classifications for each season to identify which pixels are sown to the crop of interest – a task that can be made difficult if multiple crops with similar phenologies are grown in the same region.

The two approaches for analyzing yield gap causes with remote sensing – those with and without explicit use of ancillary datasets – can obviously be used as complements to each other. In a region where multiple seasons of cloud-free images are available, a useful starting point is to construct yield gap curves (e.g., Fig. 4) to identify the persistence of yield-limiting factors. If persistence is found to occur, the maps of average yields over 5 or more years can be used to

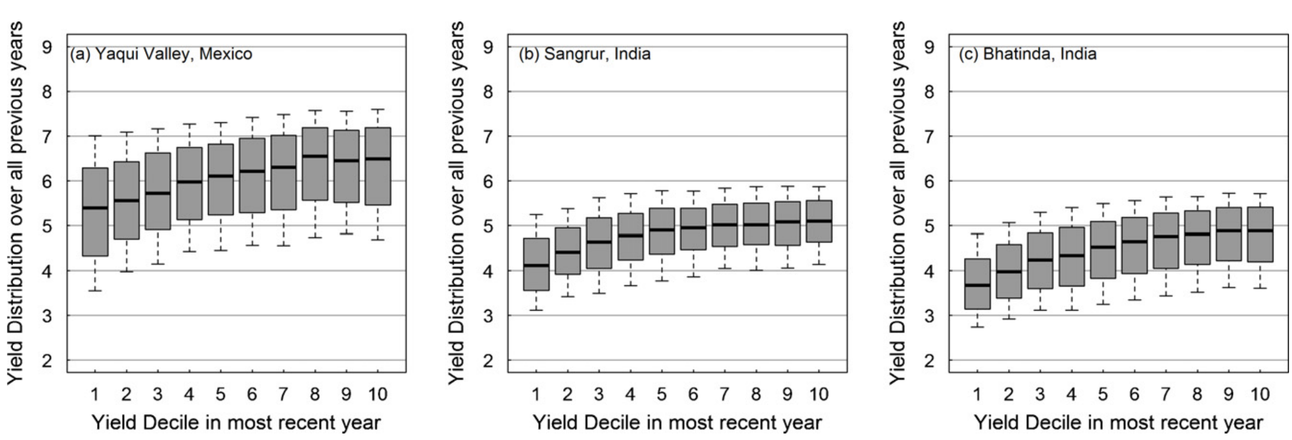

Fig. 5. A summary of wheat yield persistence in (a) Yaqui Valley District, Mexico, (b) Sangrur District, India, (c) Bhatinda District, India. Each plot shows yield distributions for 10 groups of fields, where the groups are define by the yield deciles in a single year (in this case, the last year of the study period). The yield distributions are calculated from yields on these fields from seven years prior to the year used to define the groups. Distributions are represented by the median (horizontal line), 25th and 75th percentiles (box), and 10th and 90th percentiles (whiskers). The significant overlap between the distributions indicates that the relative ranking of yields across fields tends to vary a lot from one year to the next (i.e., a lack of strong persistence). However, the positive trend in all three regions indicates some level of persistence, with high (low) yielding fields in one year more likely than average to be high (low) in other years.

generate hypotheses about the specific causes of yield variation and the types of ancillary datasets that warrant effort to obtain. Maps of persistent yield differences can also be used to target ground efforts at management or soil data collection. Conversely, in regions where fewer years of imagery are currently available, or where ancillary datasets are already in possession, then detailed analyses of an individual year or two would be a useful starting point. Subsequent studies of multiple years can then be used to test whether the factors identified in the single year are persistent.

5. Conclusions

To date, satellite data have played a relatively small role in understanding the magnitude and causes of yield gaps in most regions. However, the few examples that exist indicate that remote sensing can help to overcome some of the inherent spatial and temporal scaling issues associated with field-based approaches. Although the cost or availability of satellite data with sufficient spatial resolution to discriminate agricultural fields was an obstacle in the past, this barrier is rapidly diminishing. Improved algorithms to pre-process remote sensing data and estimate yields, and the increased availability of new, large geospatial datasets on soils, management, and weather should also benefit future efforts in this area.

Thus, remote sensing is poised to become a regular tool in the analyst's toolbox for assessing the magnitude and causes of yield gaps. The scope for applications in the near-future will likely be limited to regions without extensive intercropping within fields, given difficulties of assessing yields with remote sensing in mixed-cropping systems. Important directions for future work include further development and testing of algorithms to estimate yields (particularly for rainfed and non-cereal crops), and comparison and integration of remote sensing with studies of yield gaps based on simulation and experimental approaches. As described throughout this special issue, improved knowledge of yield gaps will play a critical role in meeting future crop demands at affordable prices and with minimal environmental impacts. The use of satellite data can accelerate the pace of discovery, and as such it represents an important area for future work.

Acknowledgements

This work was supported by NASA New Investigator grant no. NNX08AV25G to D.B.L. I thank two anonymous reviewers for very helpful comments on the manuscript.

References

Allen, R., Hanuschak, G., Craig, M., 2002. Limited Use of Remotely Sensed Data for Crop Condition Monitoring and Crop Yield Forecasting in NASS. http://www.usda.gov/nass/nassinfo/remoteuse.htm

Baruth, B., Royer, A., Klisch, A., Genovese, G., 2008. The use of remote sensing within the MARS crop yield monitoring system of the European commission. ISPRS Arch. 36, 935–959.

Bastiaanssen, W.G.M., Ali, S., 2003. A new crop yield forecasting model based on satellite measurements applied across the Indus Basin, Pakistan. Agric. Ecosyst. Environ. 94, 321–340.

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 Sens. Environ. 114, 1312–1323.

Bloom, A.J., Chapin, F.S., Mooney, H.A., 1985. Resource limitation in plants: an economic analogy. Annu. Rev. Ecol. Syst. 16, 363–392.

Clevers, J.G.P.W., 1997. A simplified approach for yield prediction of sugar beet based on optical remote sensing data. Remote Sens. Environ. 61, 221–228.

Dente, L., Satalino, G., Mattia, F., Rinaldi, M., 2008. Assimilation of leaf area index derived from ASAR and MERIS data into CERES-Wheat model to map wheat yield. Remote Sens. Environ. 112, 1395–1407.

Doraiswamy, P.C., Hatfield, J.L., Jackson, T.J., Akhmedov, B., Prueger, J., Stern, A., 2004. Crop condition and yield simulations using Landsat and MODIS. Remote Sens. Environ. 92, 548–559.

Doraiswamy, P.C., Sinclair, T.R., Hollinger, S., Akhmedov, B., Stern, A., Prueger, J., 2005. Application of MODIS derived parameters for regional crop yield assessment. Remote Sens. Environ. 97, 192–202.

Gallego, J., Carfagna, E., Baruth, B., 2010. Accuracy, Objectivity and Efficiency of Remote Sensing for Agricultural Statistics. Agricultural Survey Methods. John Wiley & Sons, Ltd., West Sussex, United Kingdom, pp. 193–211.

Gao, F., Masek, J., Schwaller, M., Hall, F., 2006. On the blending of the Landsat and MODIS surface reflectance: predicting daily Landsat surface reflectance. IEEE Trans. Geosci. Remote Sens. 44, 2207–2218.

Gitelson, A.A., Viña, A., Arkebauer, T.J., Rundquist, D.C., Keydan, G., Leavitt, B., 2003. Remote estimation of leaf area index and green leaf biomass in maize canopies. Geophys. Res. Lett. 30, 1248.

Hay, R.K.M., 1995. Harvest index: a review of its use in plant breeding and crop physiology. Ann. Appl. Biol. 126, 197–210.

Hilker, T., Wulder, M.A., Coops, N.C., Seitz, N., White, J.C., Gao, F., Masek, J.G., Stenhouse, G., 2009. Generation of dense time series synthetic Landsat data through data blending with MODIS using a spatial and temporal adaptive reflectance fusion model. Remote Sens. Environ. 113, 1988–1999.

Launay, M., Guerif, M., 2005. Assimilating remote sensing data into a crop model to improve predictive performance for spatial applications. Agric. Ecosyst. Environ. 111, 321–339.

Lobell, D., Ortiz-Monasterio, J., Addams, C., Asner, G., 2002. Soil, climate, and management impacts on regional wheat productivity in Mexico from remote sensing. Agric. Forest Meteorol. 114, 31–43.

Lobell, D.B., Asner, G.P., Ortiz-Monasterio, J.I., Benning, T.L., 2003. Remote sensing of regional crop production in the Yaqui Valley, Mexico: estimates and uncertainties. Agric. Ecosyst. Environ. 94, 205–220.

Lobell, D.B., Cassman, K.G., Field, C.B., 2009. Crop yield gaps: their importance, magnitudes, and causes. Annu. Rev. Environ. Resour. 34, 179–204.

Lobell, D.B., Ortiz-Monasterio, J.I., 2006. Regional importance of crop yield constraints: linking simulation models and geostatistics to interpret spatial patterns. Ecol. Modell. 196, 173–182.

Lobell, D.B., Ortiz-Monasterio, J.I., 2008. Satellite monitoring of yield responses to irrigation practices across thousands of fields. Agron. J. 100, 1005–1012.

Lobell, D.B., Ortiz-Monasterio, J.I., Asner, G.P., Naylor, R.L., Falcon, W.P., 2005. Combining field surveys, remote sensing, and regression trees to understand yield variations in an irrigated wheat landscape. Agron. J. 97, 241–249.

Lobell, D.B., Ortiz-Monasterio, J.I., Falcon, W.P., 2007a. Yield uncertainty at the field scale evaluated with multi-year satellite data. Agric. Syst. 92, 76–90.

Lobell, D.B., Ortiz-Monasterio, J.I., Gurrola, F.C., Valenzuela, L., 2007b. Identification of saline soils with multiyear remote sensing of crop yields. Soil Sci. Soc. Am. J. 71, 777–783.

Lobell, D.B., Ortiz-Monasterio, J.I., Lee, A.S., 2010. Satellite evidence for yield growth opportunities in Northwest India. Field Crops Res. 118, 13–20.

Monteith, J.L., 1977. Climate and the efficiency of crop production in Britain. Philos. Trans. R. Soc. Lond. B 281, 277–294.

Moulin, S., Bondeau, A., Delecolle, R., 1998. Combining agricultural crop models and satellite observations: from field to regional scales. Int. J. Remote Sens. 19, 1021–1036.

Myneni, R., Hoffman, S., Knyazikhin, Y., Privette, J., Glassy, J., Tian, Y., Wang, Y., Song, X., Zhang, Y., Smith, G., Lotsch, A., Friedl, M., Morisette, J., Votava, P., Nemani, R., Running, S., 2002. Global products of vegetation leaf area and fraction absorbed PAR from year one of MODIS data. Remote Sens. Environ. 83, 214–231.

Nightingale, J.M., Morisette, J.T., Wolfe, R.E., Tan, B., Gao, F., Ederer, G., Collatz, G.J., Turner, D.P., 2009. Temporally smoothed and gap-filled MODIS land products for carbon modelling: application of the fPAR product. Int. J. Remote Sens. 30, 1083–1090.

Ortiz-Monasterio, J.I., Lobell, D.B., 2007. Remote sensing assessment of regional yield losses due to sub-optimal planting dates and fallow period weed management. Field Crops Res. 101, 80–87.

Peng, Y., Gitelson, A.A., 2011. Remote estimation of gross primary productivity in soybean and maize based on total crop chlorophyll content. Remote Sens. Environ. 117, 440–448.

Prévot, L., Chauki, H., Troufleau, D., Weiss, M., Baret, F., Brisson, N., 2003. Assimilating optical and radar data into the STICS crop model for wheat. Agronomie 23, 297–303.

Reeves, M.C., Zhao, M., Running, S.W., 2005. Usefulness and limits of MODIS GPP for estimating wheat yield. Int. J. Remote Sens. 26, 1403–1421.

Seifert, C., Ortiz-Monasterio, J.I., Lobell, D.B., 2011. Satellite-based detection of salinity and sodicity impacts on wheat production in the Mexicali Valley. Soil Sci. Soc. Am. J. 75, 699.

Sellers, P.J., 1987. Canopy reflectance, photosynthesis, and transpiration. II. The role of biophysics in the linearity of their interdependence. Remote Sens. Environ. 21, 143–183.

Shanahan, J.F., Schepers, J.S., Francis, D.D., Varvel, G.E., Wilhelm, W.W., Tringe, J.M., Schlemmer, M.R., Major, D.J., 2001. Use of remote-sensing imagery to estimate corn grain yield. Agron. J. 93, 583–589.

Sinclair, T.R., Muchow, R.C., 1999. Radiation use efficiency. Adv. Agron. 65, 215–265.

Steinmetz, S., Guerif, M., Delecolle, R., Baret, F., 1990. Spectral estimates of the absorbed photosynthetically active radiation and light-use efficiency of a winter wheat crop subjected to nitrogen and water deficiencies. Int. J. Remote Sens. 11, 1797–1808.

Thenkabail, P.S., 2001. Optimal hyperspectral narrowbands for discriminating agricultural crops. Remote Sens. Rev. 20, 257–291.

Tittonell, P., Shepherd, K.D., Vanlauwe, B., Giller, K.E., 2008. Unravelling the effects of soil and crop management on maize productivity in smallholder agricultural

systems of western Kenya—an application of classification and regression tree analysis. Agric. Ecosyst. Environ. 123, 137–150.
Tucker, C.J., 1979. Red and photographic infrared linear combinations for monitoring vegetation. Remote Sens. Environ. 8, 127–150.
Tucker, C.J., Holben, B.N., Elgin, J.H., McMurtrey, J.E., 1980. Relationship of spectral data to grain yield variation. Photogram. Eng. Remote Sens. 45, 600–608.
Van Ittersum, M.K., Cassman, K.G., Grassini, P., Wolf, J., Tittonell, P., Hochman, Z., 2013. Yield gap analysis with local to global relevance—a review. Field Crops Res. 143, 4–17.
Weiss, M., Troufleau, D., Baret, F., Chauki, H., Prevot, L., Olioso, A., Bruguier, N., Brisson, N., 2001. Coupling canopy functioning and radiative transfer models for remote sensing data assimilation. Agric. Forest Meteorol. 108, 113–128.
Wiegand, C.L., Richardson, A.J., 1990. Use of spectral vegetation indexes to infer leaf-area, evapotranspiration and yield. 2. Results. Agron. J. 82, 630–636.
Yang, C., Bradford, J.M., Wiegand, C.L., 2001. Airborne multispectral imagery for mapping variable growing conditions and yields of cotton, grain sorghum, and corn. Trans. ASAE 44, 1983–1994.
Zarco-Tejada, P.J., Ustin, S., Whiting, M., 2005. Temporal and spatial relationships between within-field yield variability in cotton and high-spatial hyperspectral remote sensing imagery. Agron. J. 97, 641–653.
Zhao, M., Heinsch, F.A., Nemani, R.R., Running, S.W., 2005. Improvements of the MODIS terrestrial gross and net primary production global data set. Remote Sens. Environ. 95, 164–176.
Zhu, X., Chen, J., Gao, F., Chen, X., Masek, J.G., 2010. An enhanced spatial and temporal adaptive reflectance fusion model for complex heterogeneous regions. Remote Sens. Environ. 114, 2610–2623.
Zurita-Milla, R., Kaiser, G., Clevers, J.G.P.W., Schneider, W., Schaepman, M.E., 2009. Downscaling time series of MERIS full resolution data to monitor vegetation seasonal dynamics. Remote Sens. Environ. 113, 1874–1885.
Zwart, S.J., Leclert, L.M.C., 2010. A remote sensing-based irrigation performance assessment: a case study of the Office du Niger in Mali. Irrig. Sci. 28, 371–385.