Available online at www.sciencedirect.com

SCIENCE DIRECT®

Remote Sensing of Environment 93 (2004) 412–422

Remote Sensing of Environment

www.elsevier.com/locate/rse

Cropland distributions from temporal unmixing of MODIS data

David B. Lobell*, Gregory P. Asner

Department of Global Ecology, Carnegie Institution of Washington, Stanford, CA 94305-1297, United States
Department of Geological and Environmental Science, Stanford University, Stanford, CA 94305, United States

Received 3 February 2004; received in revised form 4 June 2004; accepted 2 August 2004

Abstract

Knowledge of the distribution of crop types is important for land management and trade decisions, and is needed to constrain remotely sensed estimates of variables, such as crop stress and productivity. The Moderate Resolution Imaging Spectroradiometer (MODIS) offers a unique combination of spectral, temporal, and spatial resolution compared to previous global sensors, making it a good candidate for large-scale crop type mapping. However, because of subpixel heterogeneity, the application of traditional hard classification approaches to MODIS data may result in significant errors in crop area estimation. We developed and tested a linear unmixing approach with MODIS that estimates subpixel fractions of crop area based on the temporal signature of reflectance throughout the growing season. In this method, termed probabilistic temporal unmixing (PTU), endmember sets were constructed using Landsat data to identify pure pixels, and uncertainty resulting from endmember variability was quantified using Monte Carlo simulation. This approach was evaluated using Landsat classification maps in two intensive agricultural regions, the Yaqui Valley (YV) of Mexico and the Southern Great Plains (SGP). Performance of the mixture model varied depending on the scale of comparison, with R^2 ranging from roughly 50% for estimating crop area within individual pixels to greater than 80% for crop cover within areas over 10 km^2. The results of this study demonstrate the importance of subpixel heterogeneity in cropland systems, and the potential of temporal unmixing to provide accurate and rapid assessments of land cover distributions using coarse resolution sensors, such as MODIS.

Keywords: Agriculture; Croplands; Decision tree; Landsat; Mixture modeling; MODIS

1. Introduction

Remote sensing of the extent and distribution of individual crop types has proven useful to a wide range of end-users, including governments, farmers, and scientists (Allen et al., 2002). Maps of cropland distributions are usually generated by supervised classification of multiple Landsat images throughout the growing season. These approaches require amounts of manual interpretation and cloud-free high spatial resolution imagery that are prohibitive for operational implementation over large areas and in multiple years.

The NASA Moderate Resolution Imaging Spectroradiometer (MODIS) instrument provides a unique opportunity for monitoring agricultural systems (Justice et al., 1998; Townshend & Justice, 2002). Several attributes of MODIS, including the daily global coverage, moderate spatial resolution (0.25 to 1 km), rapid availability of various products, and cost-free status may allow for operational mapping of croplands. However, the large size of even the 250-m MODIS data relative to most fields results in MODIS pixels containing mixtures of different fields, crop types, and noncrop surfaces. As a result, approaches that assign a single "hard" classification to each pixel may be prone to significant errors when mapping crop types (Defries et al., 1995; Fisher, 1997; Wu et al., 2002). While the potential errors associated with classification of mixed pixels are widely recognized, the severity depends upon factors such as landscape heterogeneity, sensor resolution, and the

* Corresponding author. Department of Global Ecology, Carnegie Institution of Washington, 260 Panama Street, Stanford, CA 94305-1297, United States. Tel.: +1 650 325 1521x329; fax: +1 650 325 6857.
E-mail address: dlobell@stanford.edu (D.B. Lobell).

0034-4257/$ - see front matter © 2004 Elsevier Inc. All rights reserved.
doi:10.1016/j.rse.2004.08.002

intended scale of analysis (Moody & Woodcock, 1994). For example, crop area estimates for large regions may be improved relative to finer scales because random errors incurred at finer scales will tend to cancel out.

The problem of the mixed pixel has been addressed in numerous studies using linear unmixing (e.g., Adams et al., 1986; Elmore et al., 2000; Smith et al., 1990). These approaches model the reflectance of a pixel as a linear combination of endmember reflectances, weighted by the areal fraction of each endmember within the pixel:

$$\rho = \sum_{i=1}^{m} C_i \rho_i + \varepsilon \quad (1)$$

where ρ is the observed pixel reflectance, C_i and ρ_i are the fractional cover and reflectance, respectively, of the ith endmember, and ε is a residual representing model error. To solve for the fractions of each of m endmembers requires at least m equations, which are most commonly generated by repeating Eq. (1) for different wavelengths, as well as including an equation to constrain the sum of endmember fractions to equal one. However, it is also valid to use reflectance from different images through time, provided that endmember fractions do not vary between image dates (DeFries et al., 1999; Quarmby et al., 1992). In this case, endmembers are defined in terms of temporal signatures of reflectance instead of or in addition to spectral signatures.

In this study, we investigated the impact of subpixel mixing in MODIS data on crop classification, and developed a linear unmixing approach that captures, rather than ignores, subpixel heterogeneity. There were three main goals of this study: (1) quantify, using resampled Landsat classification images, the potential errors resulting from classification of 250 m and 1 km MODIS data; (2) develop a linear unmixing model that estimates the subpixel fraction of individual crop types using time series of MODIS data; and (3) evaluate estimates of crop cover from MODIS unmixing applied to two agricultural regions, using Landsat estimates of crop area as a proxy for true area at a range of spatial scales.

2. Methods

2.1. Site descriptions

This study focused on two agricultural regions where Landsat-derived maps of major crop types were available: the Yaqui Valley (YV) in Northwest Mexico and the Southern Great Plains (SGP) in the United States. The YV comprises roughly 225,000 ha of irrigated cropland situated between the Gulf of California and Sierra Madre Mountains (27.5°N, 109.5°W). The majority of fields are planted with spring wheat in November–December and harvested in the following April–May. Maize is the main secondary crop, averaging 10% of planted area, and is typically sown in September and harvested in March.

The SGP (36.5°N, 97.5°W) is dominated by rain-fed winter wheat, which is planted in the fall and harvested mainly in June (National Agriculture Statistics Service, 1997). Wheat accounts for roughly 80% of cropped area in the region (National Agriculture Statistics Service, 2002), with secondary crops including maize, sorghum, soybean, and alfalfa, all of which are planted in spring and harvested in fall.

2.2. Landsat crop classification

Landsat data acquired on January 14 and March 3, 2002 were used to map wheat in YV for the 2001–2002 growing season. The images were corrected to top-of-atmosphere (TOA) reflectance and transformed to the simple ratio (SR=band 4/band 3), which exhibited a bimodal distribution that easily distinguished cropped from uncropped lands (Lobell et al., 2003). Pixels that contained an active crop in both January and March were classified as wheat, with a total regional wheat area estimate of 176,517 ha. This value is 5.3% lower than the reported area of 186,324 ha (Secretaría de Agricultura, 2003).

In SGP, a Landsat TOA reflectance image from April 4, 2000 was transformed to the SR, and all pixels in the upper part of the bimodal distribution were identified as wheat, the only major land cover that is green at this time of year. The Landsat scene covered only a subset of the SGP region. We thus focused exclusively on Garfield County, OK, which was one of few counties fully contained within the Landsat image and is representative of the SGP region. The Landsat estimate of total wheat area in Garfield was 133,869 ha, which is 2.7% lower than the reported area of 137,652 ha (National Agriculture Statistics Service, 2002).

Although the Landsat-based maps of wheat cover were not extensively validated, close agreement with reported regional totals and the well-known ability of Landsat to map crop types allowed us to assign these wheat maps as "ground truth" for subsequent evaluation of MODIS. Due to limited availability of validated Landsat data coinciding with MODIS collects, our analysis was limited to the 2001–2002 growing season in YV and 2000 in SGP, and in both cases to the major crop (wheat). These data sets provided a means to assess MODIS estimates of crop area at a range of spatial scales.

2.3. Analyses of sensor resolution and scale

One way to evaluate potential classification errors associated with course resolution data is to simulate such data using resampled Landsat imagery (e.g., He et al., 2002; Moody & Woodcock, 1994). We resampled the Landsat wheat maps to 250 m and 1 km resolution, with

the values in the resulting image expressing the subpixel proportion of wheat. These resampled images were then classified according to a simple majority rule, with all pixels containing 50% or more wheat classified as "wheat", and those with less than 50% classified as "not wheat" (see Fig. 1).

To evaluate the impacts of this classification, we considered a range of spatial scales, from a single pixel (e.g., 250×250 m) to the entire region. Here, we use scale to refer to the ground area (i.e., extent) over which crop area is computed. At each spatial scale, the total area of wheat was calculated for both the classified course resolution image and the original Landsat image, representing the true area. For example, at a scale of 10×10 pixels (equal to 625 ha for the 250-m data), the region was divided into 10×10 pixel square subregions, and the total area was computed for each subregion. We then computed the root mean square error (RMSE) and bias of the area estimates, defined as:

$$\mathrm{RMSE}(s) = \sqrt{\sum_{i=1}^{n} (\hat{a}_i - a_i)^2 / n} \tag{2}$$

$$\mathrm{bias}(s) = \sum_{i=1}^{n} (\hat{a}_i - a_i) / n \tag{3}$$

where s is the spatial scale (ground area), n is the number of subregions of size s, and $\hat{a}_i$ and a_i are the estimated and true crop area in the ith subregion, respectively. The RMSE is a combined measure of the bias and variance associated with an estimator. To isolate the variance, we also computed the coefficient of determination (R^2) between the estimated and true area, reflecting the ability of the classification to

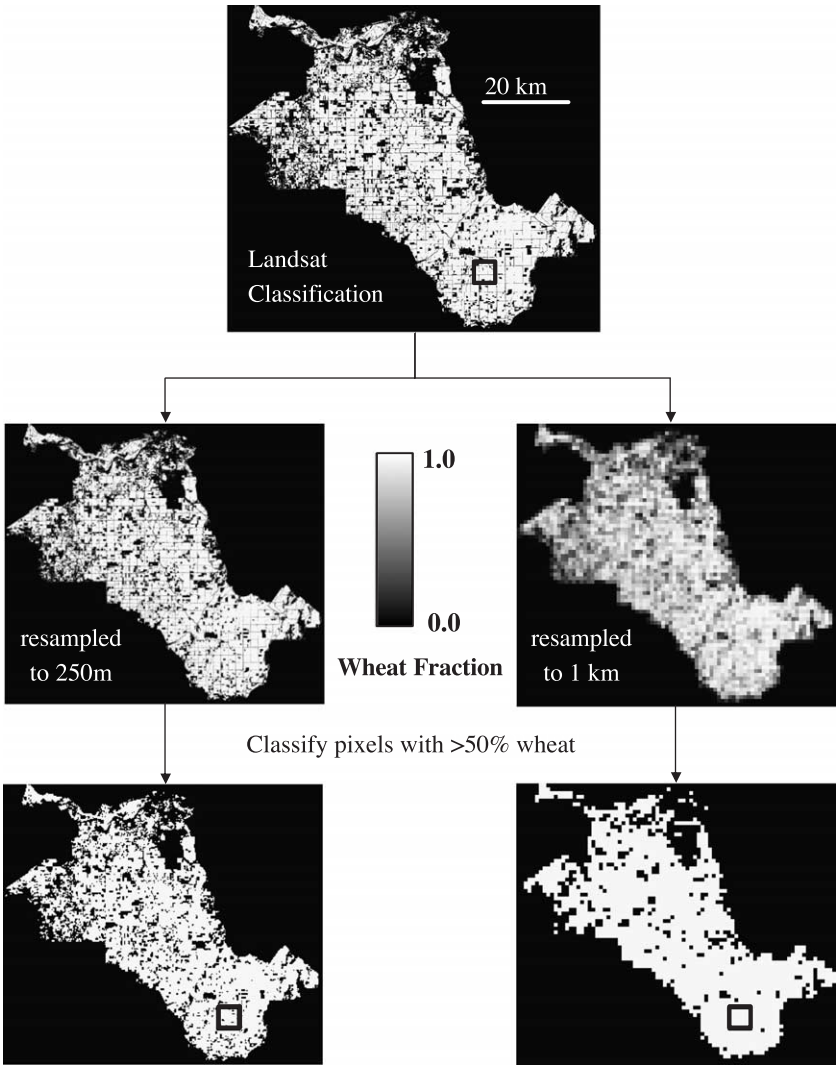

Fig. 1. Overview of procedure used to evaluate the effect of pixel size on classification-based crop area estimates. A Landsat classification-based crop map, in this case, for wheat in YV, was resampled to 250 m and 1 km resolution. These images were then classified based on a simple threshold criterion (wheat>50%). The total area of wheat was then computed for multiple nonintersecting subregions, where the size of the subregion was defined by the scale of interest, illustrated here by the square in the southeast corner.

capture spatial patterns of crop area at the given spatial scale:

$$R^2 = \frac{\mathrm{cov}(a,\hat{a})^2}{\mathrm{var}(a)\mathrm{var}(\hat{a})} \quad (4)$$

where $\mathrm{cov}(a,\hat{a})$ is the covariance between true and estimated crop area, and $\mathrm{var}(a)$ and $\mathrm{var}(\hat{a})$ are the variances of true and estimated area, respectively.

2.4. MODIS data

MODIS 16-day 250-m vegetation index (VI) composite products were obtained for both study regions for the period 2000–2002. These collection 4 products include constrained view angle–maximum value composites (CV–MVC) of the normalized difference VI (NDVI) and enhanced VI (EVI), as well as the corresponding red and near-infrared (NIR) reflectance (MODIS bands 1 and 2) and quality assessment (QA) flags (Huete et al., 2002). For this study, we used only the red and NIR values from the composite, along with the VI Usefulness Index in the QA data set. Reflectance values for dates with a VI Usefulness Index value lower than "good quality" were replaced by linearly interpolated values from the two closest dates with good, high, or perfect quality.

2.5. Probabilistic temporal unmixing

2.5.1. Temporal unmixing

The time series of MODIS reflectance were used to estimate cropland cover in each region by solving the set of linear unmixing equations, written in matrix form as:

$$\begin{pmatrix} v_{11} & \cdots & v_{1m} \\ \vdots & \cdots & \vdots \\ v_{n1} & \cdots & v_{nm} \\ 1 & \cdots & 1 \end{pmatrix} \begin{pmatrix} C_1 \\ \vdots \\ C_m \end{pmatrix} = \begin{pmatrix} v_1 \\ \vdots \\ v_n \\ 1 \end{pmatrix} \quad (5)$$

where m is the number of endmembers, n is the number of observations (number of wavelengths×number of dates), v_{ij} is the jth value for the ith endmember, C_i is the fractional cover of the ith endmember, and v_j is the value of the jth observation. Eq. (5) represents a set of n+1 individual equations, with one for each wavelength at each time step, and an additional equation constraining the sum of fractions to one.

The success of all unmixing approaches relies on the selection of appropriate endmembers. Insufficient contrast between endmembers often leads to an unstable solution, resulting in noisy and inaccurate fraction images. However, too few endmembers will fail to correctly model the pixel reflectance, leading to large errors. Several approaches have been adopted to address these issues. Roberts et al. (1998) applied hundreds of different models to each pixel, with each model containing two to three endmembers selected from a spectral database, and then selected the single model for each pixel with the lowest RMSE. This approach focused on small differences between candidate endmember spectra to adjust the model on a per-pixel basis, and relied on the RMS as an appropriate measure of model error. Asner and Lobell (2000) used linear combinations of reflectance to reduce differences between similar endmembers (e.g., all green vegetation). We found that a single model using just three endmembers (green vegetation, nonphotosynthetic vegetation, and soil) was able to accurately estimate fractions across multiple images due to the reduced endmember variability.

In the context of temporal unmixing of 250 m MODIS data, it is possible to define endmembers in terms of only red reflectance (Red), near-infrared reflectance (NIR), both Red and NIR, or an arbitrary linear combination of the two. An example of the latter is the simple perpendicular vegetation index (PVI; Richardson & Wiegand, 1977):

$$\mathrm{PVI} = \mathrm{NIR} - \mathrm{Red} \quad (6)$$

Nonlinear indices, such as NDVI or EVI, cannot be used because they fail to preserve the linear relation between endmembers and observations. For example, consider a pixel with two endmembers of fractional cover a and b. The observed PVI will equal the appropriate sum of endmember PVI, unlike for NDVI:

$$\begin{aligned} \mathrm{PVI} &= \mathrm{NIR} - \mathrm{Red} = (a\mathrm{NIR}_1 + b\mathrm{NIR}_2) - (a\mathrm{Red}_1 + b\mathrm{Red}_2) \\ &= a(\mathrm{NIR}_1 - \mathrm{Red}_1) + b(\mathrm{NIR}_2 - \mathrm{Red}_2) \\ &= a\mathrm{PVI}_1 + b\mathrm{PVI}_2 \end{aligned} \quad (7)$$

$$\begin{aligned} \mathrm{NDVI} &= \frac{\mathrm{NIR} - \mathrm{Red}}{\mathrm{NIR} + \mathrm{Red}} \\ &= \frac{(a\mathrm{NIR}_1 + b\mathrm{NIR}_2) - (a\mathrm{Red}_1 + b\mathrm{Red}_2)}{(a\mathrm{NIR}_1 + b\mathrm{NIR}_2) + (a\mathrm{Red}_1 + b\mathrm{Red}_2)} \\ &\neq a\mathrm{NDVI}_1 + b\mathrm{NDVI}_2 \end{aligned} \quad (8)$$

To determine the optimal definition of endmembers, we performed the unmixing using Red only, NIR only, Red and NIR, and PVI.

2.5.2. Probabilistic unmixing

In addition to the assumption of linear mixing, Eq. (2) assumes that endmember spectra (here, we use "spectra" to refer to reflectance across multiple wavelengths and/or dates) are known exactly for each pixel. In reality, reflectance is likely to vary across space and time, even for a narrowly defined endmember. For example, a spectrum of a wheat canopy does not appear identical for all locations and years, but varies according to range of environmental and management factors, such as temperature and planting date. Rather than define endmembers with a single spectrum, as in most approaches, we therefore prefer to define endmembers as a set of spectra which represent the full range of potential variability (Asner & Lobell, 2000; Bateson et al., 2000). The uncertainty in endmember fractions arising from endmember variability can then be

quantified using Monte Carlo sampling techniques (that is, a spectrum is randomly selected from each endmember set, fractions are estimated, and this is repeated a large number of times to derive a distribution of fractions for each endmember). Thus, endmember fractions are not estimated as single values, but rather as a probability distribution that can be used to construct confidence intervals appropriate to the desired application.

2.5.3. Endmember selection

Endmember sets were constructed by collecting spectra from pixels judged to consist entirely of a single cover type. Landsat ETM+ data was used for each region to aid in identification of large fields. The temporal domain of the endmembers was defined in each region to span the growing season of the dominant crops, without being so long as to allow multiple crops over time. In YV, the endmembers were the two major crops, wheat and maize, and uncropped land (including bare soil and urban areas; Fig. 2). The number of spectra used to define the three endmembers was 52, 28, and 46, respectively, with the lower number for maize reflecting the difficulty of finding large maize fields. Endmember spectra spanned from day of year (DOY) 241 (Aug. 29), when fall maize planting begins, through DOY 113 (Apr. 23), when wheat harvest is active. Images after late April were not included because cotton, which is commonly planted on fields after maize harvest, begins to develop in April. Defining endmembers across a longer time period would therefore violate the assumption that endmember fractions do not change through time. In practice, it is possible, although not demonstrated here, to separately unmix different times of the year to capture distributions of crops with different growing seasons.

The endmembers for SGP (number of spectra in each endmember set is shown in parentheses) were wheat (22), pasture (24), and summer crops (42), which includes

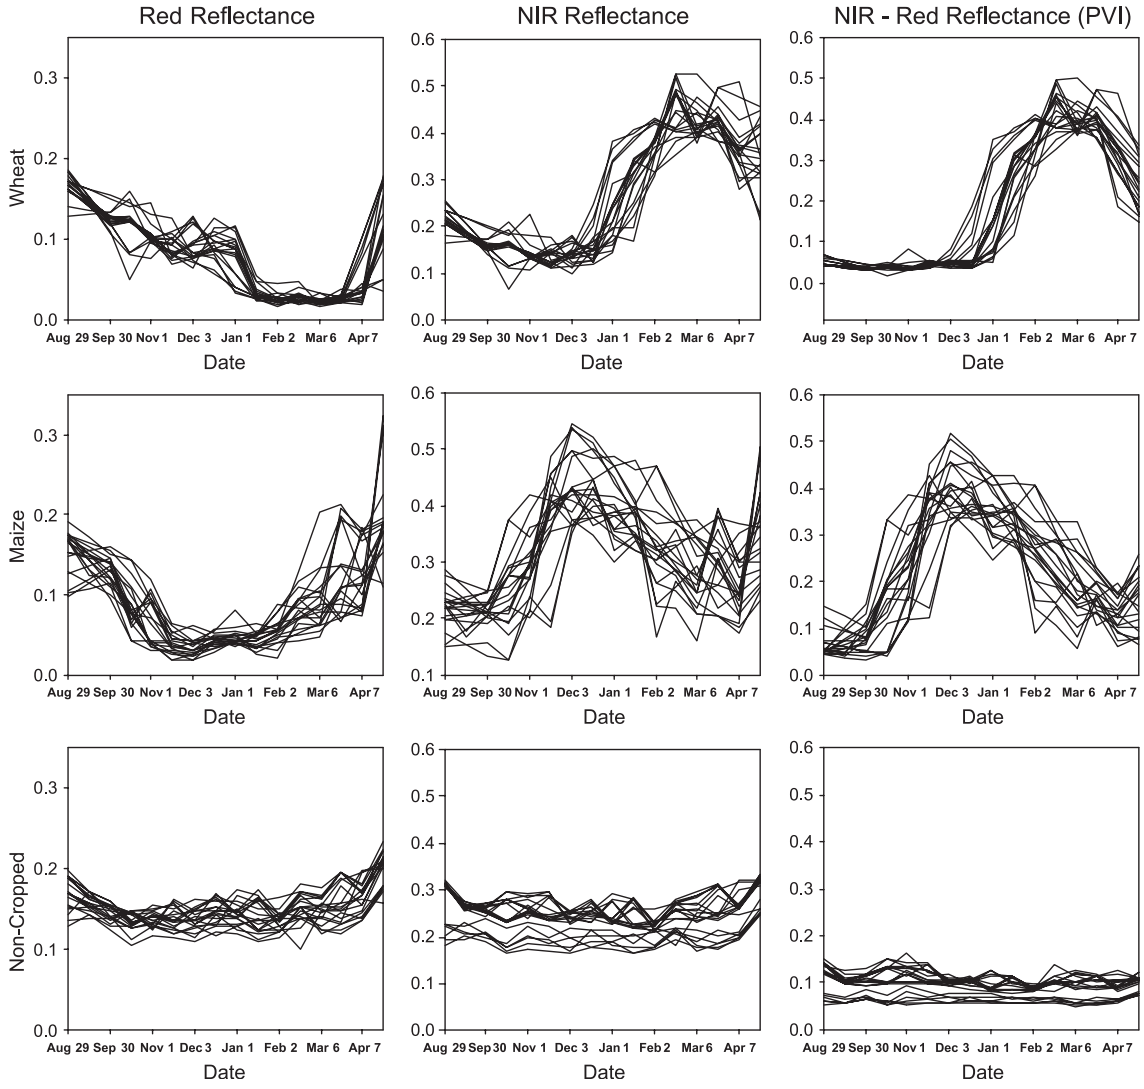

Fig. 2. Time series of red, NIR, and PVI for YV endmembers. Only 20 spectra are shown in each panel for clarity. Wheat reaches peak growth in early March, while maize peaks in December.

sorghum, corn, and soybean (Fig. 3). Endmembers spanned from DOY 49 (Feb. 18), before peak wheat growth, to DOY 305 (Nov. 1), which is after harvest of summer crops. The three summer crops were difficult to distinguish visually, based on phenology due to similar planting and harvest dates (National Agriculture Statistics Service, 1997), and were therefore combined into one endmember set. Potential methods to separate crops with similar phonologies are discussed in the Results and discussion section.

2.5.4. Summary

The entire procedure described above is illustrated in Fig. 4, and referred to hereafter as probabilistic temporal unmixing (PTU) for brevity. First, time series of red and NIR reflectance were constructed based on the MODIS NDVI composite products. Sets of image endmembers were then defined from the reflectance time series; in this case, using Landsat classification data to guide the endmember selection process. We note that endmembers may also be derived from alternative sources, such as combining models of crop growth and radiative transfer to simulate reflectance changes through the growing season. Once endmember sets were defined, each pixel was unmixed repeatedly (i.e., 50 times) with randomly selected endmembers from each set to produce a distribution of fractions. Below, we discuss the results of the PTU model applied to YV in 2001–2002 and SGP in 2000. Subscripts are used hereafter to denote the wavelength combinations used in endmember definition (i.e., PTU_R, PTU_N, PTU_{RN}, and PTU_{PVI} for Red only, NIR only, Red and NIR, and PVI, respectively).

2.6. Model evaluation

Wheat area estimates from PTU were evaluated by comparison with the total area of Landsat pixels classified as wheat within the corresponding MODIS pixel(s). As with

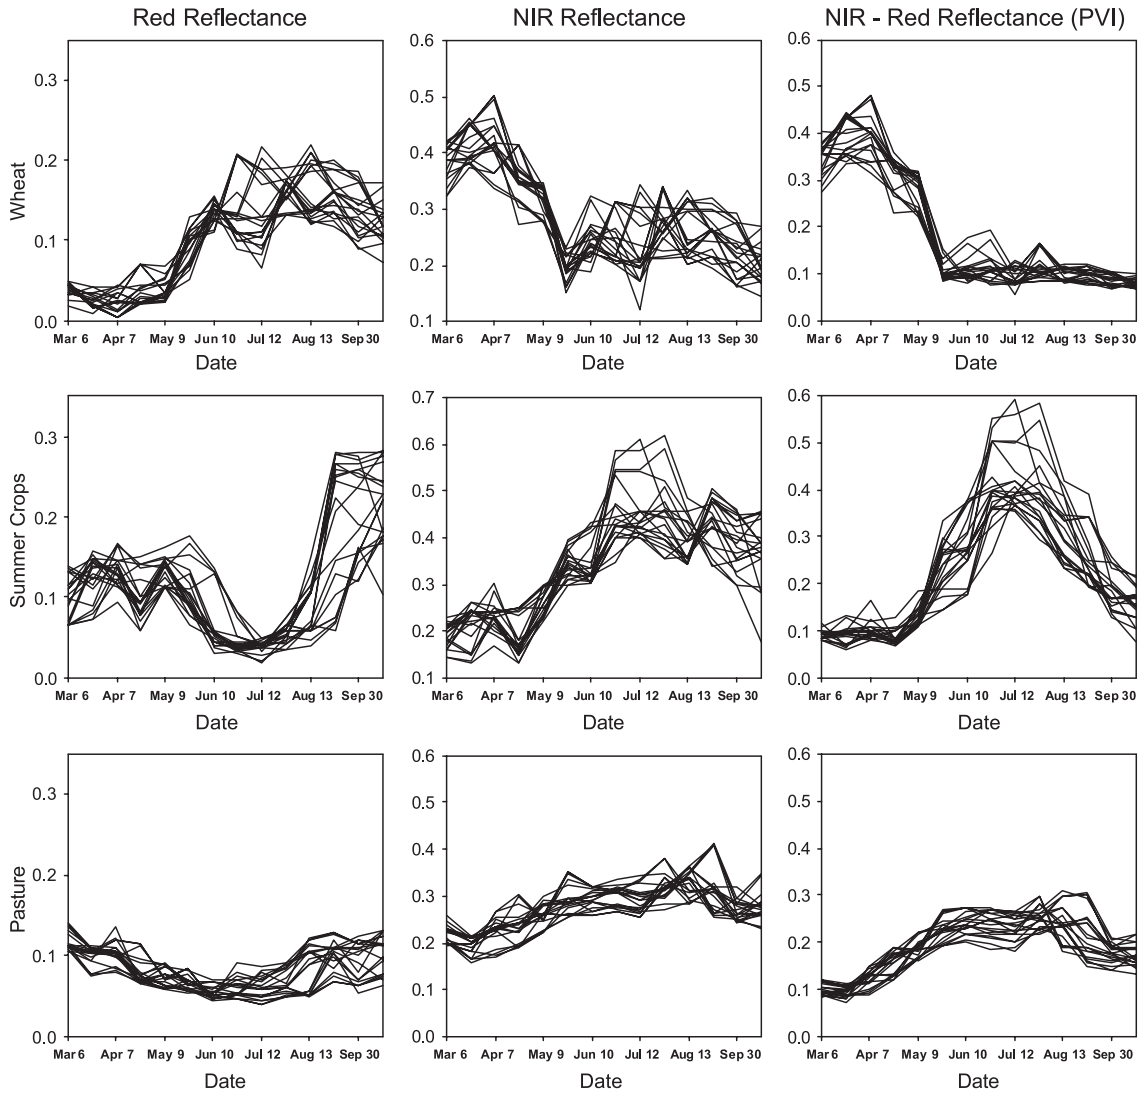

Fig. 3. Time series of red, NIR, and PVI for SGP endmembers. Only 20 spectra are shown in each panel for clarity. Wheat reaches peak growth in early April, while summer crops and pasture peak in mid- and late summer, respectively.

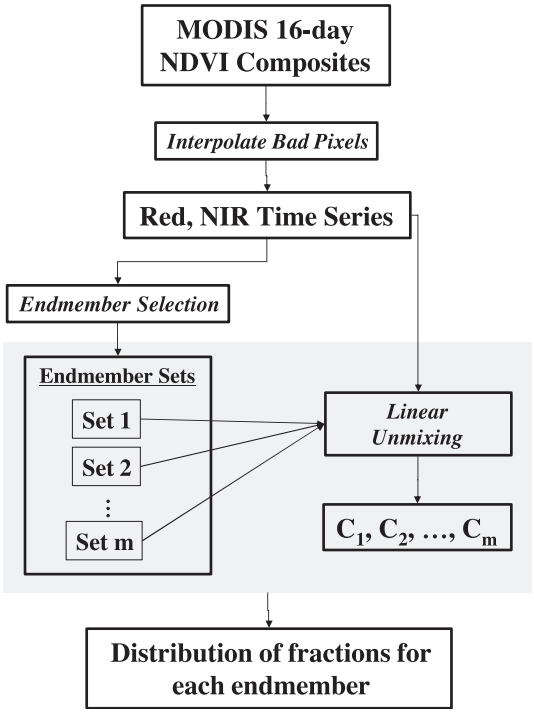

Fig. 4. Outline of probabilistic temporal unmixing (PTU) algorithm. The steps within the gray box, namely, the selection of an endmember from each set and the calculation of endmember fractions, are repeated many (50) times to derive distributions of endmember fractions that reflect the uncertainty associated with endmember variability.

the evaluation of Landsat resampling above, model assessment was performed at scales ranging from individual pixels to over 20 km^2, and was based on RMSE, bias, and R^2 statistics.

3. Results and discussion

3.1. Resampled Landsat classification

Fig. 5 summarizes the errors in estimating wheat area using Landsat classification images resampled to 250 m or 1 km resolution. Not surprisingly, the hard classification resulted in significant errors when estimating percent cover for individual pixels, since the estimate is constrained to two values (0 or 100). As the scale over which computed crop area increased, the errors associated with classification decreased. This reflects the fact that errors for individual pixels partially cancel out when summing over many pixels. However, at 250 m resolution, the RMSE fell below 5% only when computing crop area over a ground area of 300 ha (roughly fifty 250-m pixels).

The RMSE for 1 km resolution was significantly higher than 250 m over small areas, owing to greater mixing of cover types compared to 250 m. Errors for 1 km resolution remained higher even when summing over large areas (>30 km^2), the result of a significant positive bias associated with the classification estimates. This finding concurs with

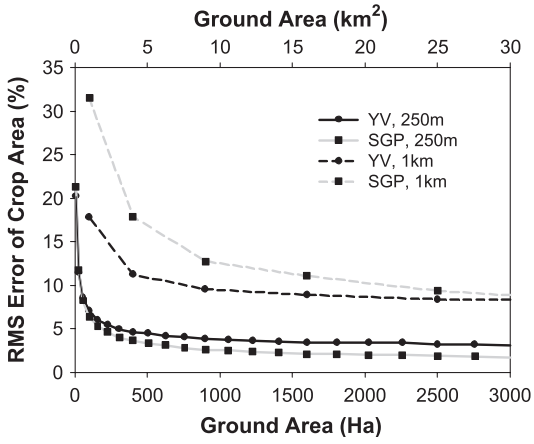

Fig. 5. Root mean square error of crop area at different spatial scales for classification of Landsat-based wheat maps resampled to 250 m and 1 km resolution.

previous studies, which concluded that classification tends to overestimate the most common cover types, such as wheat in these regions (He et al., 2002; Moody & Woodcock, 1995).

The results of the Landsat resampling analysis illustrate several important points. First, classification of mixed pixels can lead to significant errors when estimating crop area. Second, the magnitudes of these errors depend largely on the scale of analysis. For example, a study focused on total regional area may come to very different conclusions than one focused on subregional spatial patterns of land cover. Third, errors also depend on pixel resolution, with larger pixels prone to greater classification errors. In this case, for estimates of crop area over 20 km^2, the RMSE for 1-km pixels were roughly twice as large as those for 250-m pixels.

Finally, the impact of classification will depend on landscape properties, such as mean patch size and proportion of each cover type (Moody & Woodcock, 1995; Wu et al., 2002). Table 1 gives the mean and variability of patch size for wheat in the two regions, along with edge density as computed using the FRAGSTATS software (McGarigal and Marks, 1995). While the two regions in this study demonstrated similar errors from classification, it is interesting to note that SGP exhibited larger errors at fine scales, which likely reflects the greater edge density. Overall, resampling Landsat data provides a simple and convenient way to assess the likely errors associated with classification of coarser resolution imagery in a particular landscape.

Table 1
Patch area and edge density metrics for Landsat wheat classification images

Site	Mean patch size (ha)	Standard deviation of patch size (ha)	Edge density ($m^{-1}\times10{,}000$)
Yaqui Valley	21.6	294.4	19.64
Garfield, OK	17.3	245.7	45.77

Table 2
Average standard deviation of wheat fractions from PTU for different endmember combinations

Site	Red only	NIR only	Red and NIR	PVI
Yaqui Valley	0.13	0.09	0.08	0.07
Garfield, OK	0.12	0.10	0.08	0.08

3.2. Probabilistic temporal unmixing

Application of PTU using different combinations of wavelengths indicated that standard deviations of wheat fractions were generally greater than 10% when using only red or NIR (Table 2). This modeled uncertainty was reduced when using both wavelengths, either as in PTU_{RN} or PTU_{PVI}, with the lowest uncertainty in both regions achieved when using PVI (Table 2). The slightly lower standard deviation for PTU_{PVI} relative to PTU_{RN} reflects the reduced variability in endmembers when defined in terms of PVI.

Fig. 6 shows the mean and standard deviation images of wheat fractions from PTU_{PVI} for both regions. The MODIS-based wheat images were compared to the Landsat "ground truth" maps at a range of spatial scales for YV (Fig. 7) and SGP (Fig. 8). In these figures, the RMSE, bias, and R^2 are shown for PTU outputs using each wavelength combination. Results were generally consistent across the two regions, with RMSE decreasing from roughly 25–30% when evaluated at the scale of individual pixels to roughly 10–15% at scales above 10 km^2. Fractions from PTU_{PVI} and PTU_{RN} were always among the lowest RMSE and highest R^2, with PTU_R and PTU_N performing substantially worse in some cases. This is consistent with the observation above that PTU_{PVI} and PTU_{RN} demonstrated the least modeled uncertainty, in terms of standard deviation of fractions.

For individual pixels, the best PTU results (using PVI) were able to explain roughly 50% of variance in actual wheat area. This value reflects a level of uncertainty significantly higher than modeled based on endmember variability, and suggests that factors beyond endmember variability contribute to estimation error. One such factor may be geolocation errors in the MODIS preprocessing algorithms, which have been improved significantly in collection 4 but remain evident.

At scales greater than 10 km^2, PTU generally estimated true wheat area to within 20%, with much of that error attributable to bias. Over 80% of spatial variability in wheat area was captured by PTU at these coarser scales in YV, with that value rising to over 90% in SGP. Thus, linear unmixing appears very capable of quantifying coarse scale variations in crop area for the landscapes tested in this study.

3.3. Toward operational application

While this study introduced a technique for mapping cropland distributions with MODIS, several problems must be addressed if it is to be used on an operational basis. In particular, we note three main barriers to applying temporal unmixing across multiple sites and years. First, the impact of temporal variability in endmember spectra will need to be evaluated in addition to the effects of spatial variability

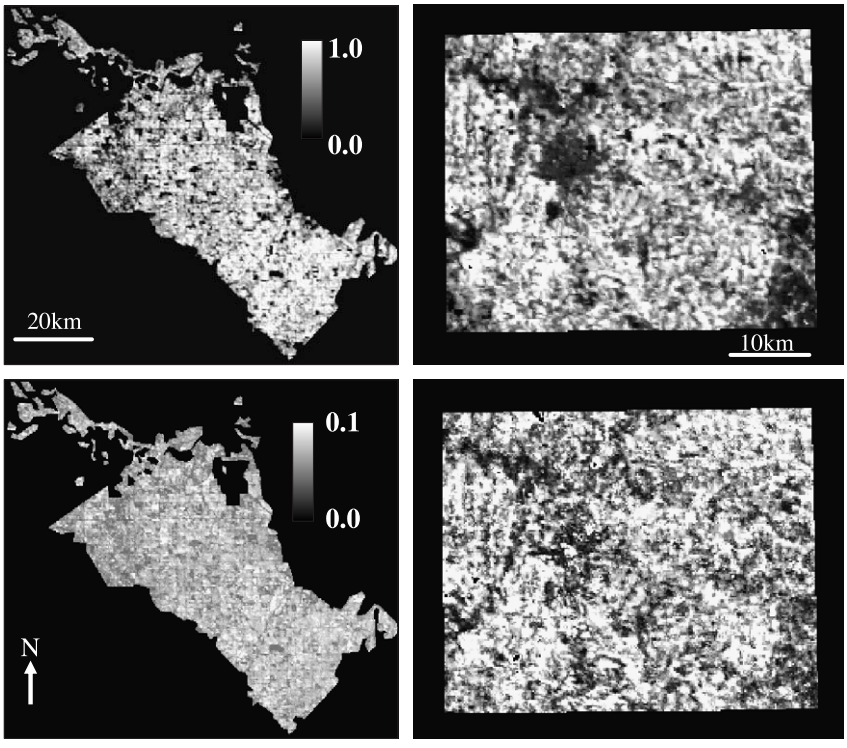

Fig. 6. Mean (top) and standard deviation (bottom) images of wheat fractions from PTU, for YV (left side) and SGP (right side).

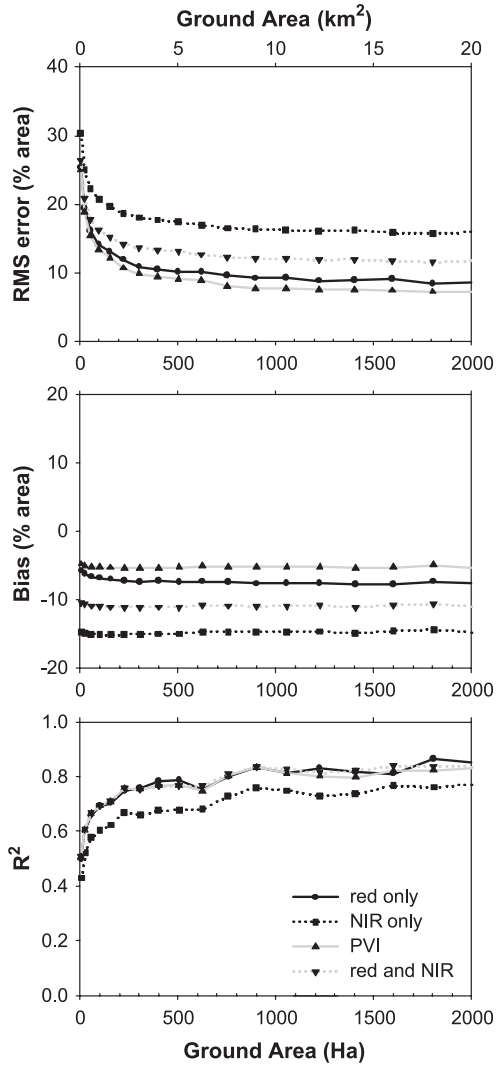

Fig. 7. Root mean square error (top), bias (middle), and R^2 (bottom) for MODIS estimates of wheat area in YV at different spatial scales.

investigated here. For example, if endmembers are defined across multiple years, will interannual differences in crop growth result in temporal signatures that are so variable as to increase fraction uncertainties to an unacceptable level? Preliminary work on this topic suggests that the answer will vary significantly according to climatic variability, management regime (e.g., irrigated vs. rain-fed), and crop type. In cases where temporal variability results in large uncertainties, it may be necessary to derive endmember spectra for each individual year, possibly with automated search procedures (Settle & Drake, 1993).

A second important issue will be that in some cases, different crops may be too similar to separate reliably with linear unmixing. This is exemplified by the case of summer crops in SGP, which were not distinguishable in our study. Similarly, in some years in YV, large areas are planted to safflower, which has a similar phenology to wheat. It may be that, in these cases, approaches such as decision trees that can utilize nonlinear transformations of the data may be more appropriate than linear unmixing. Another option is to employ multiple mixture models with different combinations of endmembers, and to use model selection criteria, such as RMSE, to identify the appropriate subset of endmembers on a per-pixel basis (Roberts et al., 1998). Preliminary work (not shown) suggests that this approach allows separation of wheat and safflower in the YV.

Third, we note that even a small average bias of unmixing results on a per-pixel basis can result in large errors when computing total regional crop cover. Specifically, the total regional bias will be the average bias multiplied by the total number of pixels, while the true crop area may be small relative to this value. A potential solution to this problem is to only compute the area of pixels above a specified threshold of percent cover for the given crop, which will effectively ignore all pixels with zero cover. An important consideration here will be to choose an appropriate threshold, and to determine the variability of this threshold between different crops, regions, and years.

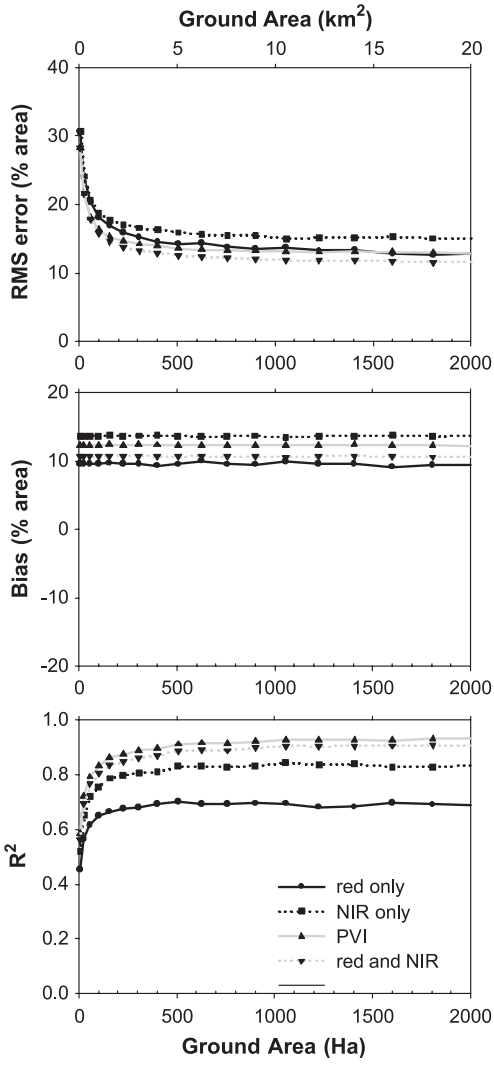

Fig. 8. Root mean square error (top), bias (middle), and R^2 (bottom) for MODIS estimates of wheat area in SGP at different spatial scales.

4. Conclusions

We have presented a probabilistic linear unmixing approach that employs time series of MODIS data to estimate subpixel fractions of land cover types along with their uncertainties. While this approach is applicable in any ecosystem, it is perhaps most relevant in systems, such as croplands, where subpixel heterogeneity is significant. In the agricultural landscapes of Northwest Mexico and the Southern Great Plains, results from the unmixing approach (PTU) were assessed at a range of spatial scales, with R^2 ranging from roughly 50% for individual pixels to greater than 80% for ground extents over 10 km^2.

The success of any technique, such as PTU, will vary based on sensors and landscape characteristics, as well as the extent of training (in this case, the specification of endmember sets). Comparison with other techniques, such as regression tree modeling and nonlinear mixture models, and additional regions are therefore needed to further test the utility of PTU. Importantly, the success of a technique also depends on the scale and required accuracies of the intended application. The significant change in accuracy with scale in this study demonstrates the need to evaluate model estimates at a range of spatial scales. This places a premium on extensive validation sets, such as those derived from higher resolution Landsat data.

Finally, we note that even with subpixel unmixing methods, MODIS data were able to capture only half of the variability expressed in Landsat data at the scale of individual fields (~10 ha). While MODIS offers great promise for characterizing croplands at larger scales, continued acquisition of Landsat resolution data is imperative for the many agricultural applications that require field-level observations.

Acknowledgements

We thank Alice Cialella of the ARM External Data Center for providing the SGP Landsat image, and A. Elmore and three anonymous reviewers for helpful comments on the manuscript. MODIS data are distributed by the Land Processes Distributed Active Archive Center (LP DAAC), located at the U.S. Geological Survey's EROS Data Center http://LPDAAC.usgs.gov. This work was supported by an NSF Graduate Research Fellowship, NASA New Investigator Program grant NAG5-8709, and the Packard Foundation. This is CIW Department of Global Ecology publication 81.

References

Adams, J. B., Smith, M. O., & Johnson, P. E. (1986). Spectral mixture modeling: A new analysis of rock and soil types at the Viking Lander I Site. *Journal of Geophysical Research*, *91*, 8098–8112.

Allen, R., Hanuschak, G., Craig, M. (2002). History of remote sensing for crop acreage in USDA's National Agricultural Statistics Service. Available at http://www.usda.gov/nass/nassinfo/remotehistory.htm

Asner, G., & Lobell, D. (2000). A biogeophysical approach for automated SWIR unmixing of soils and vegetation. *Remote Sensing of Environment*, *74*(1), 99–112.

Bateson, C. A., Asner, G. P., & Wessman, C. A. (2000). Endmember bundles: A new approach to incorporating endmember variability into spectral mixture analysis. *IEEE Transactions On Geoscience and Remote Sensing*, *38*(2), 1083–1094.

Defries, R. S., Field, C. B., Fung, I., Justice, C. O., Los, S., Matson, P. A., Matthews, E., Mooney, H. A., Potter, C. S., Prentice, K., Sellers, P. J., Townshend, J. R. G., Tucker, C. J., Ustin, S. L., & Vituosek, P. M. (1995). Mapping the land-surface for global atmosphere–biosphere models—Toward continuous distributions of vegetations functional-properties. *Journal of Geophysical Research—Atmospheres*, *100*(D10), 20867–20882.

DeFries, R. S., Townshend, J. R. G., & Hansen, M. C. (1999). Continuous fields of vegetation characteristics at the global scale at 1-km resolution. *Journal of Geophysical Research—Atmospheres*, *104*(D14), 16911–16923.

Elmore, A. J., Mustard, J. F., Manning, S. J., & Lobell, D. B. (2000). Quantifying vegetation change in semiarid environments: Precision and accuracy of spectral mixture analysis and the Normalized Difference Vegetation Index. *Remote Sensing of Environment*, *73*(1), 87–102.

Fisher, P. (1997). The pixel: A snare and a delusion. *International Journal of Remote Sensing*, *18*(3), 679–685.

He, H. S., Ventura, S. J., & Mladenoff, D. J. (2002). Effects of spatial aggregation approaches on classified satellite imagery. *International Journal of Geographical Information Science*, *16*(1), 93–109.

Huete, A., Didan, K., Miura, T., Rodriguez, E., Gao, X., & Ferreira, L. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. *Remote Sensing of Environment*, *83*(1–2), 195–213.

Justice, C. O., Vermote, E., Townshend, J. R. G., Defries, R., Roy, D. P., Hall, D. K., Salomonson, V. V., Privette, J. L., Riggs, G., Strahler, A., Lucht, W., Myneni, R. B., Knyazikhin, Y., Running, S. W., Nemani, R. R., Wan, Z. M., Huete, A. R., van Leeuwen, W., Wolfe, R. E., Giglio, L., Muller, J. P., Lewis, P., & Barnsley, M. J. (1998). The Moderate Resolution Imaging Spectroradiometer (MODIS): Land remote sensing for global change research. *IEEE Transactions On Geoscience and Remote Sensing*, *36*(4), 1228–1249.

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. *Agriculture, Ecosystems and Environment*, *94*, 205–220.

McGarigal, K., & Marks, B. J. (1995). FRAGSTATS: spatial pattern analysis program for quantifying landscape structure. USDA Forest Service, Pacific Northwest Research Station, Gen. Tech. Report PNW-GTR-351.

Moody, A., & Woodcock, C. E. (1994). Scale-dependent errors in the estimation of land-cover proportions—Implications for global land-cover datasets. *Photogrammetric Engineering and Remote Sensing*, *60*(5), 585–594.

Moody, A., & Woodcock, C. E. (1995). The influence of scale and the spatial characteristics of landscapes on land-cover mapping using remote sensing. *Landscape Ecology*, *10*(6), 363–379.

National Agriculture Statistics Service, U.S.D.o.A. (2002). Published Estimates Data Base. Available at http://www.nass.usda.gov:81/ipedb/

Quarmby, N. A., Townshend, J. R. G., Settle, J. J., White, K. H., Milnes, M., Hindle, T. L., & Silleos, N. (1992). Linear mixture modeling applied to Avhrr data for crop area estimation. *International Journal of Remote Sensing*, *13*(3), 415–425.

Richardson, A. J., & Wiegand, C. L. (1977). Distinguishing vegetation from soil background information. *Photogrammetric Engineering and Remote Sensing*, *43*(12), 1541–1552.

Roberts, D. A., Gardner, M., Church, R., Ustin, S., Scheer, G., & Green, R. O. (1998). Mapping chaparral in the Santa Monica Mountains using multiple endmember spectral mixture models. *Remote Sensing of Environment*, *65*(3), 267–279.

Secretaría de Agricultura, G., Desarrollo Rural, Pesca y Alimentación (SAGARPA). (2003). Secretaría de Agricultura, Ganadería, Desarrollo Rural, Pesca y Alimentación (SAGARPA). Available at http://www.siea.sagarpa.gov.mx

Settle, J. J., & Drake, N. A. (1993). Linear mixing and the estimation of ground cover proportions. *International Journal of Remote Sensing*, *14*, 1159–1177.

Smith, M. O., Ustin, S. L., Adams, J. B., & Gillespie, A. R. (1990). Vegetation in deserts: I. Regional measure of abundance from multispectral images. *Remote Sensing of Environment*, *31*, 1–26.

Townshend, J., & Justice, C. (2002). Towards operational monitoring of terrestrial systems by moderate-resolution remote sensing. *Remote Sensing of Environment*, *83*(1–2), 351–359.

Wu, J. G., Shen, W. J., Sun, W. Z., & Tueller, P. T. (2002). Empirical patterns of the effects of changing scale on landscape metrics. *Landscape Ecology*, *17*(8), 761–782.