On-the-Fly Massively Multitemporal Change Detection Using Statistical Quality Control Charts and Landsat Data

Evan B. Brooks, Randolph H. Wynne, *Member, IEEE*, Valerie A. Thomas, Christine E. Blinn, and John W. Coulston

***Abstract*—One challenge to implementing spectral change detection algorithms using multitemporal Landsat data is that key dates and periods are often missing from the record due to weather disturbances and lapses in continuous coverage. This paper presents a method that utilizes residuals from harmonic regression over years of Landsat data, in conjunction with statistical quality control charts, to signal subtle disturbances in vegetative cover. These charts are able to detect changes from both deforestation and subtler forest degradation and thinning. First, harmonic regression residuals are computed after fitting models to interannual training data. These residual time series are then subjected to Shewhart X-bar control charts and exponentially weighted moving average charts. The Shewhart X-bar charts are also utilized in the algorithm to generate a data-driven cloud filter, effectively removing clouds and cloud shadows on a location-specific basis. Disturbed pixels are indicated when the charts signal a deviation from data-driven control limits. The methods are applied to a collection of loblolly pine (*Pinus taeda*) stands in Alabama, USA. The results are compared with stands for which known thinning has occurred at known times. The method yielded an overall accuracy of 85%, with the particular result that it provided afforestation/deforestation maps on a per-image basis, producing new maps with each successive incorporated image. These maps matched very well with observed changes in aerial photography over the test period. Accordingly, the method is highly recommended for on-the-fly change detection, for changes in both land use and land management within a given land use.**

***Index Terms*—Degradation, statistical process control, thinning, trajectory.**

Manuscript received January 16, 2013; revised May 21, 2013; accepted June 19, 2013. Date of publication August 1, 2013; date of current version February 27, 2014. This work was supported in part by the U.S. Department of Agriculture Forest Service Cooperative Agreement with Virginia Polytechnic Institute and State University (Virginia Tech) under Grant 10-CA-11330145-158, by the Landsat Science Team under U.S. Geological Survey Contract G12PC00073, by the Pine Integrated Network: Education, Mitigation, and Adaptation Project (PINEMAP, Coordinated Agricultural Project funded in 2011 by the USDA National Institute of Food and Agriculture), by the McIntire-Stennis Cooperative Forestry Research Program through the USDA CSREES under Project VA-136614, and by the Department of Forest Resources and Environmental Conservation at Virginia Tech.

E. B. Brooks, R. H. Wynne, V. A. Thomas, and C. E. Blinn are with the Department of Forest Resources and Environmental Conservation, Virginia Polytechnic Institute and State University, Blacksburg, VA 24060 USA (e-mail: evbrooks@vt.edu).

J. W. Coulston is with the USDA Forest Service Southern Research Station, Forest Inventory and Analysis Unit, Knoxville, TN 37919 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TGRS.2013.2272545

I. Introduction

A. Background

Detection of disturbances and changes in forest cover is a major application of remote sensing of the earth. The simplest method of detecting changes from one time to another is to compare two images from the same area and measure the changes, often signaling changes outside some threshold. Depending on the features in the scene and the timing of the images, this method may be used to identify major changes such as clearcuts, but it is difficult to identify minor, yet still important, changes such as forest degradation or thinning [1]–[4]. While small changes may appear negligible on a pixel basis, on a cumulative basis, they can be the most important feature of change across a scene. By and large, such subtle changes have gone unreported [4], while results of fires and wholesale logging are reported. This accounts for some of the uncertainty in forest mass estimates, for example [1]. The difficulty in detecting subtle changes is mostly due to the spatial resolution of the sensor typically being coarser than the resolution of such minor changes. Subtle changes in forest cover can also affect the spectral characteristics of the area in question, although these changes may be very small and difficult to detect on an image-to-image basis.

In addition to standard bitemporal change detection methods, there are a variety of other methods that make use of multiple Landsat images, often over many years' worth of data, to improve accuracy, resolution, or to provide an ongoing record. These methods are best characterized as "massively multitemporal" in their nature, and they rely in large part on the free availability of Landsat data as of October 2008. Trajectory-based image analysis would be one subset of this class of methods. One corporate example of this paradigm is the persistent change detection used by MDA Federal, Inc. [5]. In this method, changes are only registered if a certain number of consecutive images before the candidate change are significantly different from a consecutive number after the change.

Some more examples of massively multitemporal change detection methods utilizing large Landsat time series applied to forestry follow. The vegetation change tracker [6], [7] uses the concept of normalized values for pixels based on the mean and standard deviation of known forest pixels from the same scene, using multiple images to improve the forest/nonforest classification. This concept has been used to detect and track

changes in forest cover over several regions in the U.S. [6], such as Mississippi [8]. In another vein, LandTrendr [9], [10] is an automated trajectory-based image analysis algorithm. This algorithm allows users to automatically segment and identify different trends in land cover change for pixels over time. While not specifically designed to signal disturbances as they occur, the vertices in the segments allow for an easy classification of trends in time series, representing shifts from one stage of growth to another. The CLASlite software [2], [3] applies its algorithm to a variety of satellite sources in a user-friendly format to produce forest change maps in tropical regions. This free software is distributed by the Carnegie Institution for Science to governments around the South American continent for monitoring of tropical deforestation and degradation, including selective logging.

The interest in developing massively multitemporal change detection methods continues to increase. As a prime example, during the time of drafting this paper, another method was published by Zhu *et al.* [11]. This publication merits special interest here because of the similarity in initial approaches, to be discussed in later parts of this paper. Works such as [5]–[11] and this paper are evidence of a change in paradigm for change detection, transitioning from comparisons of temporally distant images to looking for deviations from model predictions.

While these other methods exist and are effective, there is still room for a simple algorithm that can produce change maps for a scene, updated from previous maps, with each new image of that scene. This sort of "on the fly" methodology would allow for a continuous monitoring paradigm to be implemented, rather than one of yearly summaries. Our goal in this paper is to introduce such an algorithm, based on a well-established statistical methodology used in other disciplines. We note here that Zhu *et al.* [11] has also produced an algorithm that can incorporate incoming data, using different methods for determining when a pixel has undergone significant change. One clear distinction between the methods is that, in this paper, images are considered to be a quasi-systematic sample, the derivatives of which are then used as inputs in statistical process control tools.

In previous work [11], [12], it was shown that Fourier regression may be used to generate smooth curves fitted to interannual Landsat data. In particular, by collating multiple years' worth of Landsat data, one can take advantage of the fact that the Landsat satellites do not have precise yearly return times to "fill in" even more days of the year. This allows one to simulate weather-free conditions on any given day of the year, with good accuracy, even if cloud cover prevents direct use of some images.

It was further shown that the residuals left from subtracting the curve from the data produce a record of shifts from the expectation in the time series development. Two things thus become apparent from this record. If the pixel remains stable over the course of the years, then an extremely good fitted curve will result, in terms of accuracy and detail. If, on the other hand, the pixel undergoes changes or disturbances, either in the context of catastrophic disturbance or in subtle changes due to forest growth or climate, then the residuals left over from subtracting the fitted "averaging" curve will produce a profile of the changes. Using this profile, it becomes possible to use incoming scene data to detect disturbances to the pixels, in a form of real-time environmental monitoring.

In the areas of industry and manufacturing, there has naturally been great interest in the idea of real-time monitoring of processes. Accordingly, there is a field of statistics, quality control, which addresses this interest through the creation of statistical tools for the purpose of actively monitoring processes. The key among these tools are *quality control charts*, which take systematically measured data (although not necessarily regular data) and *signal* the operator in the event that the monitored process goes out of control. For reference, a brief explanation of the quality control charts used in this paper follows.

B. Shewhart Charts

Shewhart charts, originally developed by Walter A. Shewhart in 1924 [13], are the foundation of the idea of statistical control charts. They have historically been used as the standard to test newer control charts against [14], and while they have drawbacks, their simplicity and versatility makes them reasonable to use even today. Consider a time sequence $t_1, t_2, \ldots$ with $t_n < t_{n+1}$ and associated measurements from a process $x_1, x_2, \ldots$, where the measurements are assumed to be independently normally distributed with mean μ and standard deviation σ. If the mean of the process remains at μ, then one should expect the standardized value $z = (x - \mu/\sigma)$ to be within 3σ about 99.7% of the time. If the value is beyond this, it may be evidence that the process mean has shifted. In general, a Shewhart X-bar (or $\overline{X}$) chart at time t_n is given by

$$z_n = \frac{x_n - \mu}{\sigma}. \tag{1}$$

This chart *signals* (note the use as a verb) if the value moves outside the range $\pm L\sigma$, where $L\sigma$ is considered the generic *control limit* of the chart. An example of a Shewhart X-bar chart and its lower control limit is shown in Fig. 1. The pixel in the figure is in a forested area that is stable, suggested by the lack of any clear trend in the residual values in the plot. Due to the statistical properties of the sample, the chart is bound to signal eventually, even if the process remains in control, but the control limits may be set so that this happens only rarely.

Shewhart charts are good for detecting large process mean shifts; however, due to the fact that they only use the present point, they may miss smaller consistent shifts [13], [15], [16]. They are ideally suited for signaling anomalies in the history, such as a passing cloud. Fig. 1 offers an example of this property. The red values in Fig. 1 correspond to dates for which this pixel was shaded by clouds. They do not represent a sustained disturbance because the subsequent values demonstrate a rapid return to the original trend. X-bar charts could be also useful for quickly detecting that a forest has been clearcut, but they are not as useful for detecting a subtle thinning in forest cover.

C. EWMA Charts

Again, consider a time sequence $t_1, t_2, \ldots$ with $t_n < t_{n+1}$, starting at $n = 0$, and associated measurements from a process

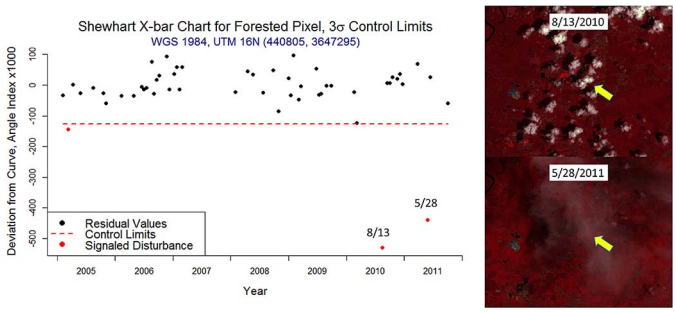

Fig. 1. Shewhart X-bar chart for residual values after removing seasonality. The angle index is a vegetation index, with higher values corresponding to denser vegetation. Signaled dates, outside the control limits, are in red. They correspond to dates for which this pixel (yellow arrow) was shaded by clouds, also evidenced by the subsequent rapid return to the trend.

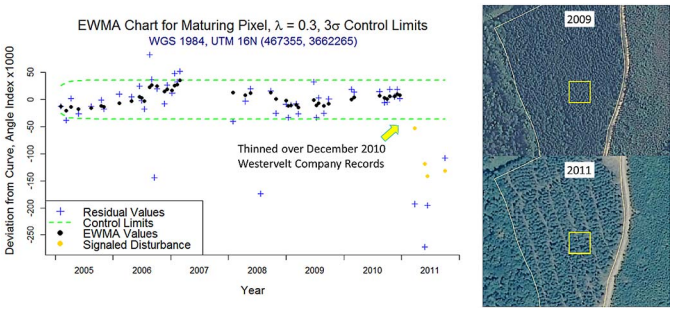

Fig. 2. EWMA chart for residual values after removing seasonality. The pixel in question (yellow boxes at right) underwent a thin in December 2010, according to company records. True color National Agriculture Imagery Program (NAIP) imagery acquired from 2009 and 2011 with 1-m spatial resolution.

$x_1, x_2, \ldots$, where the measurements are assumed to be independently normally distributed with mean μ and standard deviation σ. Then, for a *tuning parameter* $0 < \lambda \leq 1$, the exponentially weighted moving average (EWMA) chart [15], [16] for the process at time t_n is given by

$$z_n^* = (1-\lambda) z_{n-1}^* + \lambda x_n. \tag{2}$$

Thus, the chart value for a given time is a function of the entire history of the chart. The extent to which this history is utilized is characterized by the value of λ, which determines how retrospective the chart is. Values of λ close to 1 will result in a chart that assigns little weight to previous values (the extreme case, i.e., $\lambda = 1$, yields a Shewhart X-bar chart after standardization by μ and σ), whereas a value of λ close to zero relies primarily on historical data. Smaller values for λ may be useful when the variation in process data is very great relative to the shift being detected (low signal-to-noise ratio). In such a case, incoming values should be tempered against the previous trend to avoid frequent false signals. An example of this chart is shown in Fig. 2.

The chart signals when it exceeds the asymptotic control limits (*CL*) given by

$$CL = \mu \pm L_{\mathrm{EWMA}} s \sqrt{\left(\frac{\lambda}{2-\lambda}\right)\left[1-(1-\lambda)^{2n}\right]}. \tag{3}$$

Here, $L_{\mathrm{EWMA}} s$ is the desired number of standard deviations for marking "out of control," and s is the estimated historical standard deviation of the data. Note how the control limit initializes at a lower value and approaches an asymptotic limit of $\mu \pm L_{\mathrm{EWMA}} s \sqrt{(\lambda/2-\lambda)}$ as the chart continues to operate. As a consequence, it is fairly common for EWMA charts to experience "warm-up noise" in the form of false signals during the early training period. For practical purposes, this asymptote is very quickly achieved, typically in the first 15 to 20 samples taken. Thus, the control limits for EWMA charts, while appearing to grow continuously, may be treated as constant limits for all practical purposes.

Equation (2) allows all past measurements to be used in the calculation, with greater weight on the recent measurements. Of particular importance is the fact that the entire chart's history is encapsulated in a single term, the EWMA chart's previous value z_{n-1}. Once that has been calculated, one may discard all the previous data and still compute the next value upon receipt of the next observation. This ease of future processing is the reason that EWMA charts are well suited for on-the-fly monitoring. Additionally, the running average aspect allows the EWMA chart to be fairly robust to the normality assumption, making it particularly appealing. Depending on how λ is set, the corresponding chart is fast in detecting small changes [15], [16]. Thus, we anticipate that appropriately tuned EWMA charts will allow us to detect subtle changes caused by thinning of a forest or possibly changes in leaf area as a result of drought.

Based on the aforementioned reasons, EWMA charts represent an extremely useful tool for environmental monitoring via Landsat time series. There are different types of charts, which are useful for detecting different sorts of disturbance. In the case of catastrophic changes, a simple Shewhart chart would suffice. However, for processes that show gradual change over time, other more time-weighted charts, such as EWMA charts, would be more appropriate. There are other types of control charts designed to signal for small shifts, such as the cumulative sum (CUSUM) chart. These may also have application in an environmental monitoring context, but including them here is beyond the scope of this paper.

Control charts are regularly used in areas of manufacturing [17], computing [18], pharmacology [19], and medicine [20], to name a few examples. The general field of quality control charts is included under the aegis of *statistical process control* [16].

II. Data

Our study area included portions of Mississippi and Alabama, USA, for which the location and timing of private logging activities were known for the 2009–2011 timeframe. To have both training and testing data, we acquired a collection of cloud-free Landsat images (10% nominal cloud cover or less, 51 images total) from January 2005 through December 2011 from the U.S. Geological Survey Global Visualization Viewer (USGS GLOVIS) website [21], using only Landsat 5 imagery. The images were all corrected to L1T prior to downloading. We avoided using Landsat 7 imagery, so that we would not need to contend with the scan line corrector (SLC) problem in the study timeframe. Nevertheless, the process shown here could be used on Landsat 7 images, provided that care is taken to screen away no-data values at the outset, similar to a cloud mask.

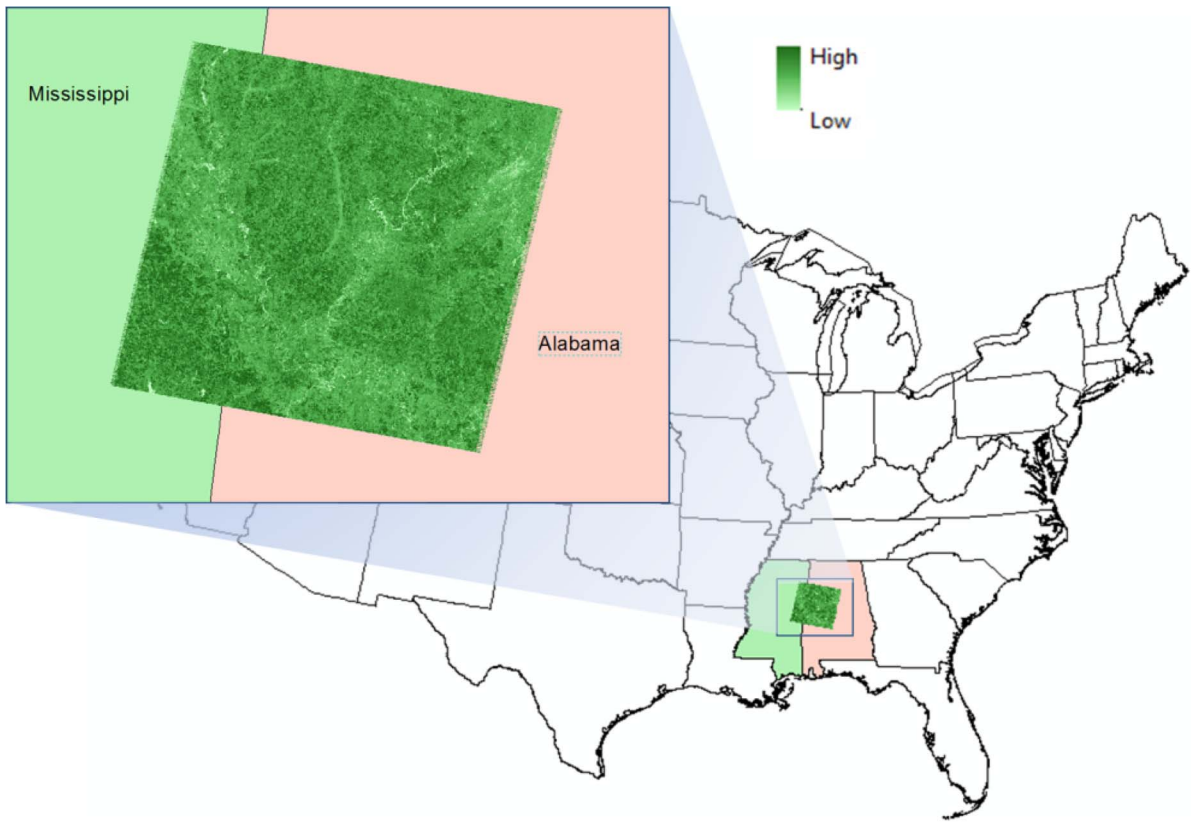

Fig. 3. Study area, Landsat path/row 21/37. The detail is tasseled cap angle index from 10/3/2011, where higher values (here as darker) indicate denser vegetation.

For validation data, we used two data sets. First, we used aerial photo mosaics from 2009 and 2011 as a high-resolution check. The images were obtained from the U.S. Department of Agriculture Natural Resources Conservation service (USDA NRC) Geospatial Data Gateway [22]. We were able to use the mosaics to provide an effective "before and after" check for disturbances within Landsat-resolution regions. We used the entire spatial extent of path/row 21/37, as shown in Fig. 3.

Second, we used the records of timber management of loblolly pine (*Pinus taeda*) from the Westervelt Company (The Westervelt Company, Inc., Tuscaloosa, AL, USA), focusing specifically on harvests from 2009 through 2011, in order to match the aerial mosaics and the Landsat data. This harvesting data consisted of polygons for which the date of harvesting and the general type of harvesting were noted. We used only thinning polygons for this study, as any general method capable of detecting thinning should by default also detect clearcuts.

All Landsat scenes were corrected to surface reflectance by using the Landsat Ecosystem Disturbance Adaptive Processing System (LEDAPS) [23], followed by dark object subtraction using band minima. The dark object subtraction had a significant effect in reducing time series noise, since post-LEDAPS evaluation of the reflectance data revealed that some scenes were uniformly brighter or darker than the remainder of the time series. We converted the resulting base Landsat bands into tasseled cap values [24], and from these, we calculated the tasseled cap angle index (*AI*) [25] by the formula

$$AI = \tan^{-1} \frac{TC_{\text{Greenness}}}{TC_{\text{Brightness}}}. \tag{4}$$

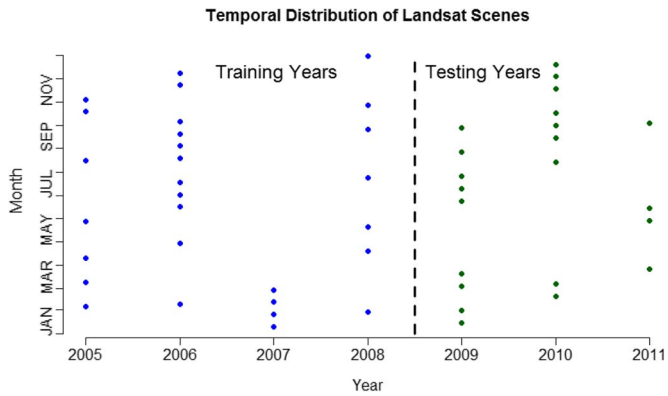

Fig. 4. Temporal distribution of training and testing data used in this study.

Note that, while all Landsat scenes were nominally cloud free, there were still anomalies present in the images, such as "popcorn clouds" scattered across the scenes at times. We chose the angle index because it has been shown to be sensitive to more subtle vegetation changes [25].

In order to have a solid baseline for disturbance detection, we used the first four years' worth of data (2005–2009, 29 images total) for the study area to train the harmonic regression. Fig. 4 shows the temporal distribution of the Landsat images for the study area. Such a graph clearly presents the richness or paucity of the Landsat stack being used. Note that we selected a generous amount of training data to provide a full-year range of values in terms of day of the year. Based on the distribution of image dates, it most likely would have been feasible to use any combination of 2005, 2006, or 2008 as training data.

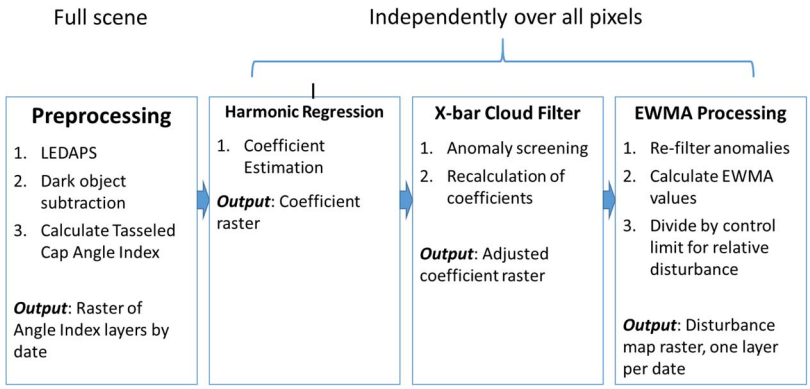

Fig. 5. Flowchart for EWMA detection algorithm.

However, our aerial imagery was from 2009 and 2011, and hence, we chose the cutoff for training and testing based on that, in order to get the closest possible matching between EWMA signals and the validation imagery. It is worth noting the lack of sparsely clouded images in 2007. These temporal features of the data also factored into our decision to use the four-year period as training.

III. Methods

In order to apply statistical process control to Landsat data, the data must first be processed and rendered into roughly independent normally distributed variables. This is accomplished by use of harmonic regression, a modified version of that used in [12] and similar to the initial steps of that used in [11]. Deciduous trees have an easily recognizable phenological curve, and loblolly pine trees lose approximately half of their needles in fall and winter [26], rendering their phenological curves amenable to modeling by harmonic regression by virtue of the clear seasonal pattern. By subtracting the fitted temporal curve from the existing data, in this case, the angle index, we remove seasonality and the bulk of temporal autocorrelation from the data, obtaining a set of residual values that may be treated as being normally distributed and statistically independent. This process works over the entire image but provides results unique to each pixel. A full description of the algorithm follows, which is divided into elements of harmonic regression, adjustment, and EWMA chart processing, with a summary of the method given in Fig. 5.

A. Harmonic Regression Algorithm

From a collection of d images over days of the year (from 1 to 365 or 366, depending on the year, possibly spanning multiple years), denote the dates $\underline{T}$ and the pixel-specific values for those dates $\underline{V_p}$ as the $d \times 1$ column vectors

$$\underline{T}_{d\times 1} = [t_1]_{i\in\{1,2,\ldots,d\}} \text{ and } \underline{V_p}_{d\times 1} = [v_{pi}]_{i\in\{1,2,\ldots,d\}} \tag{5}$$

noting that the p subscript emphasizes the vector's dependence on the pixel in question. For simplicity, scale $\underline{T}$ by converting days of the year to values on $[0, 2\pi]$ by multiplying by $(2\pi/365)$ (or $(2\pi/366)$ when appropriate), yielding

$$\underline{T}_{d\times 1} = [\tau_i]_{i\in\{1,2,\ldots,d\}} = \underline{T}\frac{2\pi}{365}. \tag{6}$$

Let us assume that a correct linear model specification [27] for the time series by day of the year is given by a harmonic series with m harmonics, with m sufficiently smaller than d, and independent identically distributed normal errors, such that

$$\underline{V_p}_{d\times 1} = M_{d\times(1+2m)}\underline{\beta_p}_{(1+2m)\times 1} + \underline{\varepsilon_p}_{d\times 1} \tag{7}$$

where, in this case, the input matrix is given, as a function of m and $\underline{T}_{d\times 1}$, by (8), shown on the bottom of the page. The harmonic coefficients are given by

$$\underline{\beta_p}_{(1+2m)\times 1} = [\,a_{p0} \quad a_{p1} \quad b_{p1} \quad \cdots \quad a_{pm} \quad b_{pm}\,]' \tag{9}$$

with transposition denoted by $'$, and the errors are given by

$$\underline{\varepsilon_p}_{d\times 1} = [\varepsilon_{pi}]_{i\in\{1,2,\ldots,d\}}, \quad \varepsilon_{pi} \sim \text{i.i.d. } N\left(0, \sigma_p^2\right), \sigma_p \in \mathbb{R}^+. \tag{10}$$

Then, estimate the pixel-specific harmonic coefficients

$$\underline{\hat{\beta}_p}_{(1+2m)\times 1} = [\,\hat{a}_{p0} \quad \hat{a}_{p1} \quad \hat{b}_{p1} \quad \cdots \quad \hat{a}_{pm} \quad \hat{b}_{pm}\,]' \tag{11}$$

by the usual least squares method [27], i.e.,

$$\underline{\hat{\beta}_p}_{(1+2m)\times 1} = \left(M'_{(1+2m)\times d}M_{d\times(1+2m)}\right)^{-1} M'_{(1+2m)\times d}\underline{V_p}_{d\times 1}. \tag{12}$$

$$M_{d\times(1+2m)} = \begin{bmatrix} 1 & \sin(1\tau_1) & \cos(1\tau_1) & \sin(2\tau_1) & \cos(2\tau_1) & \cdots & \sin(m\tau_1) & \cos(m\tau_1) \\ 1 & \sin(1\tau_1) & \cos(1\tau_1) & \sin(2\tau_1) & \ddots & \ddots & \cdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \ddots & \cdots & \vdots \\ 1 & \sin(1\tau_d) & \cos(1\tau_d) & \sin(2\tau_d) & \cos(2\tau_d) & \cdots & \sin(m\tau_d) & \cos(m\tau_d) \end{bmatrix} \tag{8}$$

B. X-Bar Cloud Filtering

In practice, a pixel may display anomalous values corresponding to small-scale cloud cover, shading, or other short-lived events that should not be modeled. As an additional precaution, the pixel-specific time series is scrubbed for anomalous values by checking the residuals from the previous model against a low-threshold Shewhart X-bar chart [15], [16], as follows.

If we denote the fitted values from the previous model as

$$\underline{\hat{V}_p}_{d\times 1} = M_{d\times(1+2m)}\underline{\hat{\beta}_p}_{(1+2m)\times 1} \tag{13}$$

then we calculate the residuals the usual way, i.e.,

$$\underline{R_p}_{d\times 1} = [r_{pi}]_{i\in\{1,2,\ldots,d\}} = \underline{V_p}_{d\times 1} - \underline{\hat{V}_p}_{d\times 1}. \tag{14}$$

We then compute an estimated value for the error variance

$$\hat{\sigma}_p^2 = \left(\frac{1}{d-1}\right)\underline{R'_p}_{1\times d}\underline{R_p}_{d\times 1}. \tag{15}$$

We take this value to determine which residuals are beyond a user-defined *control limit*, denoted here as L, and identify these dates as anomalous. Note that this is equivalent to processing the residuals in a Shewhart X-bar chart with control limits $\pm L\sigma$ [15], [16]. We thus obtain a vector of the remaining dates, i.e.,

$$\underline{T^*}_{d^*\times 1} = [\tau_{pi}]_{i\ni(|r_{pi}|>L\hat{\sigma}_p)} = \left[\tau^*_{pj}\right]_{j=\{1,2,\ldots,d^*\}} \tag{16}$$

where $d^* \leq d$, and the new index j reinforces the notion that the elements of $\underline{T^*}_{d^*\times 1}$ do not necessarily correspond to the original elements of $\underline{T}_{d\times 1}$.

From $\underline{T^*}_{d^*\times 1}$, we recompute an estimate of the harmonic coefficients, denoted $\underline{\hat{\beta}_p}^*_{(1+2m)\times 1}$, using (7)–(12) and reindexing the vectors and matrices in a manner similar to (16). This process is illustrated on a sample pixel in Fig. 6. Note how the adjusted curve (green) is not deflected by the unusual points indicated in red.

This approach to filtering out clouds is of interest in its own right, because it relies only on the input Landsat data, a concept also explored in [11]. By extracting as much useful information as possible from each Landsat image, it may be possible to meaningfully utilize scenes with considerable cloud cover. In this study, we used relatively clear data, as our focus was on the detection of subtle changes to land cover, but the notion of using X-bar charts on harmonic residuals as a data-driven cloud filter is worth revisiting.

C. EWMA Chart Algorithm

To illustrate how we process the time series in an EWMA chart, we make some slight modifications to our notation. In the context of the training and testing periods, we treat the training period as the first d dates in a larger timeframe of $D = d + d_1$ dates, and we accordingly extend the scaled date vector for all dates to $\underline{T}_{D\times 1}$, its input matrix $M_{D\times(1+2m)}$, and the pixel-specific value vector to $\underline{V_p}_{D\times 1}$. We first compute the fitted vector for all dates, using the previously computed adjusted coefficients $\underline{\hat{\beta}_p}^*_{(1+2m)\times 1}$ to obtain

$$\underline{\hat{V}_p}^*_{D\times 1} = M_{D\times(1+2m)}\underline{\hat{\beta}_p}^*_{(1+2m)\times 1}. \tag{17}$$

We then compute the residual values

$$\underline{R_p}^*_{D\times 1} = \underline{V_p}_{D\times 1} - \underline{\hat{V}_p}^*_{D\times 1}. \tag{18}$$

As these residuals result from least squares estimation, we treat them as if they had a mean of 0. In practice, this is not the case because of the adjustment for anomalous values. Thus, it is necessary to once again account for these anomalous values, as incorporating them into an EWMA chart could extend their influence well beyond the date of anomaly, resulting in extended false signals. Thus, all of the residuals are reprocessed in a low-threshold X-bar chart with (possibly different) control limits $\pm L^*\sigma$ by reapplication of (15) and (16), resulting in a reduction in the residual time series to one of length D^*, denoted

$$\underline{R^{**}}_{D^*\times 1} = [r_{pi}]_{i\ni\left(|r_{pi}|>L^*\hat{\sigma}^*_p\right)} = \left[r^{**}_{pj}\right]_{j=\{1,2,\ldots,D^*\}}. \tag{19}$$

This reduced set of residuals is the object which we process directly in an EWMA chart. For a tuning parameter $\lambda \in (0, 1]$ ($\lambda = 0$ being undesirable), we generate a transformation matrix Λ [15], [16], as follows:

$$\Lambda_{D^*\times D^*} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ (1-\lambda) & \lambda & 0 & \cdots & 0 \\ (1-\lambda)^2 & (1-\lambda)\lambda & \lambda & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ (1-\lambda)^{D^*} & (1-\lambda)^{D^*-1}\lambda & (1-\lambda)^{D^*-2}\lambda^2 & \cdots & \lambda \end{bmatrix}. \tag{20}$$

Then, the EWMA vector is

$$\underline{Z^*}_{D^*\times 1} = \Lambda_{D^*\times D^*}\underline{R^{**}}_{D^*\times 1}. \tag{21}$$

It is this vector which we plot as a time series against its corresponding dates in the EWMA chart. Here, a great benefit of EWMA charts in general is realized, because for subsequent images on dates $D^* + 1, D^* + 2, \ldots$, we simply compute the pixels' fitted values based on the adjusted harmonic coefficients, take residuals, screen them against the X-bar control limits of $\pm L^*\sigma$ in case of short-lived anomalies, and utilize the standard EWMA definition in (2), as follows:

$$z^*_{jp} = (1-\lambda)z^*_{(j-1)p} + \lambda r^{**}_{jp}, \quad j = \{D^*+1, D^*+2, \ldots\}. \tag{22}$$

This allows new images to be incorporated easily into the existing history, making the detection method on-the-fly.

Recall that the control limits for an EWMA chart with tuning parameter λ are given by

$$CL = \mu \pm L_{\mathrm{EWMA}}s\sqrt{\left(\frac{\lambda}{2-\lambda}\right)[1-(1-\lambda)^{2j}]}. \tag{23}$$

In our case, with the residuals having an assumed stable mean of 0, this becomes

$$\underline{CL}_{D^*\times 1} = \left[0 \pm L_{\mathrm{EWMA}}\hat{\sigma}_p^* \sqrt{\left(\frac{\lambda}{2-\lambda}\right)\left[1-(1-\lambda)^{2j}\right]}\right]_{j=\{1,2,\ldots,D^*\}}. \quad (24)$$

Note that the control limits depend on an estimated value of the standard deviation. We compute this estimate using the training (or historical) data in the absence of prior knowledge about the pixel [15], [16]. For pixels that underwent a disturbance in the training period, both the harmonic coefficients and the estimated standard deviation reflect this in the poor model fit. Since such pixels are not likely to remain in a partly disturbed and partly stable state immediately after the disturbance, the practical effect is that pixels disturbed during the training period tend to signal almost immediately in the testing period.

Recalling that the control limits rapidly move toward the asymptote of $L_{\mathrm{EWMA}}\hat{\sigma}_p^*\sqrt{(\lambda/2-\lambda)}$, and recalling that a variety of land cover classes exist in any given scene, we divide the EWMA chart by the control limits to get the disturbance record, or flag history, for that pixel, i.e.,

$$\underline{F}_{D^*\times 1} = \underline{Z^*}_{D^*\times 1} \div \underline{CL}_{D^*\times 1} \quad (25)$$

where $\div$ represents elementwise division. We insert 0's artificially into the parts of the history that were screened out by (19), although any other code value could be inserted in the interest of easily tracking anomalies. By iterating this process over all pixels p in the scene, we obtain a raster with an equal number of layers as the input raster, corresponding to the initial temporal distribution and giving relative disturbances from the conditions predicted by the training period.

D. Specific Application

For this study, given our stack of Landsat images converted to tasseled cap angle index, we took the dates associated with each image and used the images corresponding to the first four years (2005–2008, inclusive) of data as inputs into the EWMA detection algorithm, first computing harmonic coefficient estimates. In our case, we used $m = 2$ harmonics, resulting in five coefficient estimates for each pixel: a constant term with two sine and cosine terms of increasing frequency. We chose this value based on observations in [12] that two harmonics were usually sufficient to capture the vast majority of the periodic variation in the time series. It is worth noting that one effect of harmonic regression is the removal of seasonal temporal autocorrelation in stable pixels, as the regression captures such periodic behavior. It is also worth noting that, by taking multiple years as training data, the coefficient estimates for pixels undergoing consistent vegetative growth "centered" the curve in the middle years of the training period.

Once the residuals were computed, we then calculated the historical standard deviation of each time series from the pixel-specific residuals and flagged any residual values farther than $L = 2$ standard units away from 0, in effect passing the residual time series through a low-threshold X-bar chart. We used 0 as the mean since the residuals under a stable pixel are assumed to be 0. We treated the flagged values as anomalous and temporarily discarded them, recalculating the least squares estimates of the harmonic coefficients for the remaining time series values. This simple second iteration had the effect of screening away most cloud and shadow interference without the use of cloud masks. The result of the process at this point was a five-layer raster of adjusted harmonic coefficient values, providing baseline information for phenological processes in the study area. This proved much easier to store and manipulate than a raster containing fitted values for all of the dates in the study timeframe.

From this baseline data, on a per-pixel basis, we recomputed fitted values for the entire history of the time series, both in the training and testing periods, 2005–2011, in all. We calculated residual values again by subtracting the fits from the observed angle index values, and once more we screened, from the training period only, anomalous values with a low-threshold (2σ) X-bar chart. This had the effect of reducing the estimate of the historical standard deviation, thus improving the sensitivity of the EWMA chart used next. We then used this standard deviation estimate for calculating the EWMA control limits, setting the mean part of the control limit to 0 for simplicity in light of the empirical observation that the calculated mean estimates were very close to 0 in the overwhelming majority of pixels. In order to filter out short-term anomalies in the testing period and avoid biasing the subsequent EWMA chart values in the event of such an anomaly, we subjected the testing period residuals to a very high threshold $(L = 2)$ X-bar chart. We chose this value based on empirical observation of the charts, balancing the need to screen short-term anomalies with the need to signal for persistent disturbances.

To test which weight parameter would be best for our purposes, we ran the algorithm on a subset of the scene, letting parameters for the charts range from $\lambda = 0.1$ to $\lambda = 1$ by increments of 0.1. These parameters drive the retrospective nature of the chart, with more retrospective charts being less likely to read false positives due to anomalous data. Upon initial testing, even light disturbances signaled across the range of weights. Accordingly, we chose a weight that reduced the chance of false signals due to singularly anomalous data values. That is, we chose a weight which gave charts that were stable in the presence of large singular deviations in the residual time series yet were responsive enough to signal disturbances within one to three dates, depending on the severity of the disturbance in question. Based on our observation of the charts, we determined that an EWMA weight value of $\lambda = 0.3$ worked well enough in our study area, being sensitive to disturbances while not providing excessive false signals. It is uncertain whether this value would generally perform well across other land cover types, but we do note that other weights in the range $0.1 \leq \lambda \leq 0.5$ gave comparable results in our study area.

In order to derive an estimate of the relative severity of the disturbances flagged, we took signaling values (those outside the control limits of the chart) and divided these values by that particular chart's control limits, rounding down to the

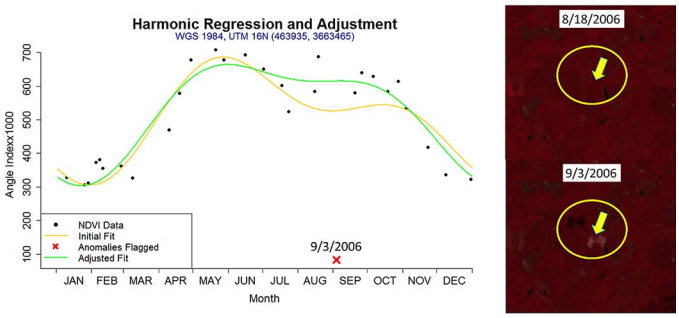

Fig. 6. Illustration of X-bar adjustment for more robust harmonic coefficients. The anomalous value (red x) was excluded when calculating the adjusted fit (green). The images at the right, with the pixel in question marked by the yellow arrow, show that the anomaly was caused by a small cloud.

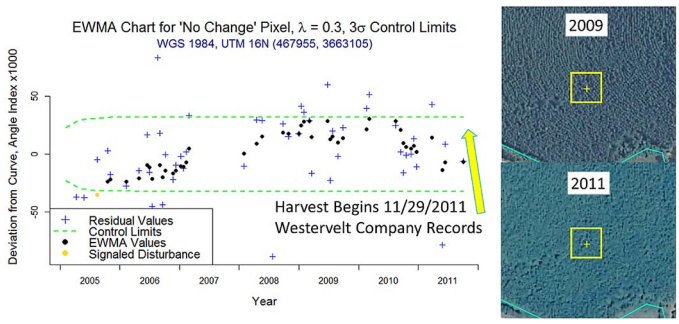

Fig. 7. EWMA chart for a pixel that had a harvest after the timeframe. The images to the right depict the pixel in question (yellow box) according to the aerial images.

nearest integer in absolute value. Since the control limits were stable by the time the testing period was reached, the division provided a way to compare relative disturbances. Thus, for each pixel, we obtained a time series of integer values, indicating both when disturbances were signaled, the relative severity of the disturbance, and the nature of the disturbance (growth or reduction in vegetative cover). We compiled these time series per pixel and generated an output raster with one layer for each date in the test timeframe.

IV. Results

The algorithm's outputs took the form of a stacked raster, each layer corresponding to the signaled disturbances for a particular date in the original Landsat stack. In effect, our results could be checked along three dimensions: space, time, and severity. In the following sections, we attempt to give a sense of the results along these lines, using the aerial images and Westervelt polygons to validate our observations.

The challenge in analyzing the results comes in showing that the disturbances are accurately signaled. In order to do this, we used the Westervelt polygons in conjunction with the aerial imagery. Three questions were of interest in working with the harvest polygon data. First, did the control chart algorithm accurately signal (or not signal) according to the actual changes on the landscape? This is a question of whether the EWMA charts properly identified disturbances in space, although answering this question depends on the disturbance severity and time of occurrence. The second question is one of accurately identifying the severity of the disturbance, distinguishing between subtle and gross disturbances. Was there a relationship between the severity of the disturbance signaled and the severity indicated by the polygon data? The third question was about the timing of the disturbances being signaled. Did the EWMA charts identify the correct times of disturbance?

It is worth noting some considerations in using the polygon information. First, the polygons were all classified as types of thinning. Despite this, cursory observation of the aerial images showed that the polygons' treatments were not uniform in either time or space. That is, some polygons were thinned unevenly (a transport road cut through one section, for example), and commonly, the polygons were large enough that it would have taken weeks to give them a full thinning treatment. Second,

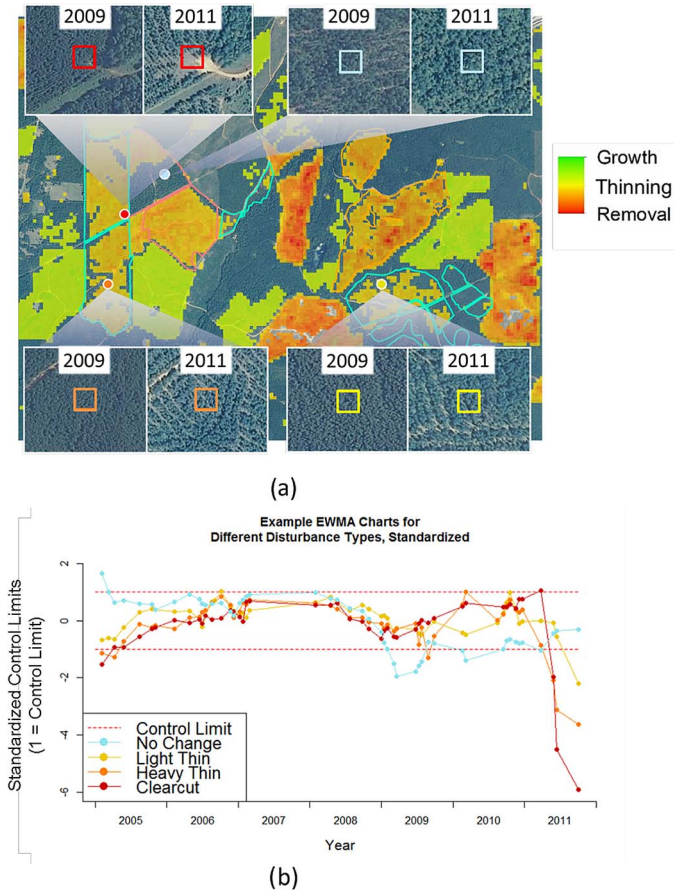

Fig. 8. (a) Example pixels of each type of disturbance, from a variety of Westervelt polygons. The squares represent Landsat pixels in scale and location, and the following descriptions apply only to the space within the squares. Clockwise from upper right: No Change (blue boxes); Light Thin (yellow); Heavy Thin (orange); Clearcut (red). The clearcut in this case was at the head of a logging road within a thinned stand. Green regions are where the algorithm signaled growth; orange and red regions are where the algorithm signaled thinning and removal. (b) EWMA charts for the example pixels in (a). Note that the No Change pixel was actually thinned in 2009, in agreement with that polygon's information.

some of the treatments observed in the aerial imagery did not match the descriptions from the polygon data. For example, one polygon was in actuality completely clearcut based on the aerial photos. Another challenge came from the aerial photos themselves. Being mosaics, the images were taken at varying points throughout their respective years, with the majority being

TABLE I
ACCURACY ASSESSMENT CRITERIA

Aerial Classification	Description
No Change	No visible difference, or possibly slight increase in vegetation
Light Thin	1-25% relative decrease in canopy or tree cover
Heavy Thin	26-75% relative decrease in canopy or tree cover
Clearcut	76-100% relative decrease in canopy or tree cover

taken during the summer months. Thus, it was possible, for polygons harvested on the edge of the testing timeframe, that the photographs missed the contrast in the treatment, either before or after. We illustrate this in Fig. 7, where the documented thinning occurs both after the final Landsat image is taken and after the 2011 aerial photograph.

In light of the aforementioned challenges, we chose to conduct the accuracy assessment for the space and severity dimensions by selecting one point at random from within each Westervelt polygon (where we had records of harvesting), associating each point with the pixel containing it. In all, we selected 141 pixels in this manner. For each of these pixels, we selected the last element in the disturbance time series. We chose the last element in order to have the best possible chance at matching the disturbance record with the 2011 aerial imagery, because our validation involved observing the difference between 2009 and 2011 in the aerial mosaics. In such a situation, attempting to compare the disturbance record during earlier dates in the time series with the 2011 images would have resulted in mismatches due to postdisturbance recovery. An example of this situation is depicted in Fig. 8, in the case of the blue example pixel. The stand at this pixel was thinned early on in 2009 (but after the aerial image for 2009). By 2011, the forest had recovered sufficiently for the aerial photographs to show little change in vegetative cover. If we had classified this pixel as thinned, it would have disagreed with the later aerial assessment.

Thus, the accuracy assessment was carried out as follows. For each test pixel, we recorded the final disturbance signal associated with that pixel. Then, independently observing the region at that pixel, we made a visual comparison between the 2009 image and the 2011 image. For consistency in estimation, we used the aerial classification shown in Table I and illustrated in Fig. 8. Note that the criteria in the table are for relative changes, not absolute changes. Note further the effect of calculating pixel disturbances independently of their surroundings, as the red example pixel in Fig. 8 reads as a clearcut by relative cover change within the pixel, despite its neighborhood being obviously thinned.

A. Accuracy Assessment (Space)

To assess the spatial accuracy, we simply observed the agreement between the EWMA charts for the relevant pixels and the aerial images on a change/no-change basis. We treated values for which the EWMA charts showed net growth as if they were no change, so that "change" in this context was equivalent to vegetative removal. The results of this dichotomous accuracy assessment are given in Table II.

TABLE II
DICHOTOMOUS ACCURACY ASSESSMENT RESULTS

Aerial \ EWMA	No Disturbance	Disturbance
No Disturbance	27	15
Disturbance	6	93

The overall accuracy was 85% (120/141), and the Cohen's kappa for the dichotomous assessment was 0.621. In general, we observed good agreement between the EWMA charts and the aerial interpretations, both in the case of disturbance and in the case of no disturbance. The number of no-disturbance observations may seem surprising, given that each pixel in the assessment was contained in a polygon documented as being harvested between 2009 and 2011, but we have already shown typical examples of the no-disturbance occurrences in Figs. 7 and 8 (blue pixel). Similarly, we have shown typical examples of disturbances in Fig. 2 and Fig. 8 (red, orange, and yellow pixels).

The commission error (false signal rate) was 36% (15/42), although this was impacted by the timing of the aerial photographs and the ability of the algorithm to signal changes at the edge of the timeframe. One example of this is shown in Fig. 9, in which the aerial photograph was taken after the thinning took place. In this case, the aerial interpretation was that of "no change," since the interpretations were based on *relative* forest cover. However, the thinning is documented and is signaled on the EWMA chart, leading to an error of commission.

The omission error rate was 6% (6/99). We show a typical example of this type of error in Fig. 10. In this pixel, the underlying forest was steadily maturing over time, as evidenced by the clear increasing trend in the EWMA chart values. By the time the pixel was thinned, the EWMA values were high enough that the thinning had the effect of returning the pixel to its original condition. Noting that the raw residual (blue cross) for the final value is quite low, relatively, we are confident that the EWMA chart would have signaled given one or maybe two additional images to confirm the new residual mean.

These aforementioned examples suggest that the bulk of these errors, particularly the commission errors, are "startup" and "cutoff" errors, artifacts of the timeframe's finite nature. This is because the EWMA chart necessarily trades a little in response time, for the ability to detect a smaller change, as a result of its weighted average. From the examples, it is clear that the response time of the charts is still usually very fast, typically signaling in the first image after the disturbance. In a continuously running algorithm with no last date of observation, these errors probably would not have occurred, or at the least, would have been mitigated in the next iteration or two. We will discuss the temporal responsiveness of the EWMA charts further when answering our question about the time dimension.

B. Accuracy Assessment (Severity)

The basic question about the severity dimension was whether the EWMA charts produced signals in agreement with the aerial

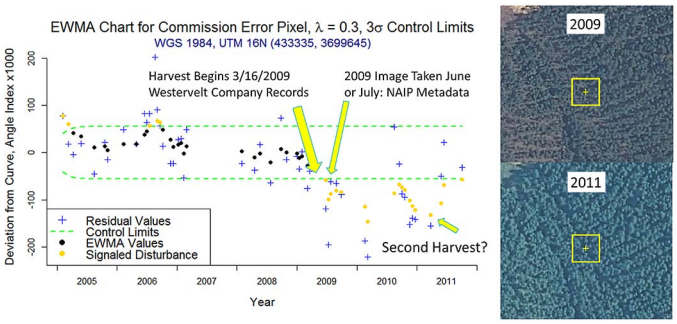

Fig. 9. EWMA chart for a pixel with a commission error (false alarm). The images to the right depict the pixel in question (yellow box) according to the aerial images.

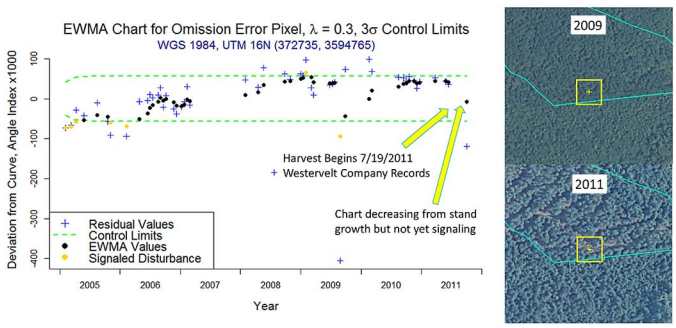

Fig. 10. EWMA chart for a pixel with an omission error (failure to signal). The images to the right depict the pixel in question (yellow box) according to the aerial images.

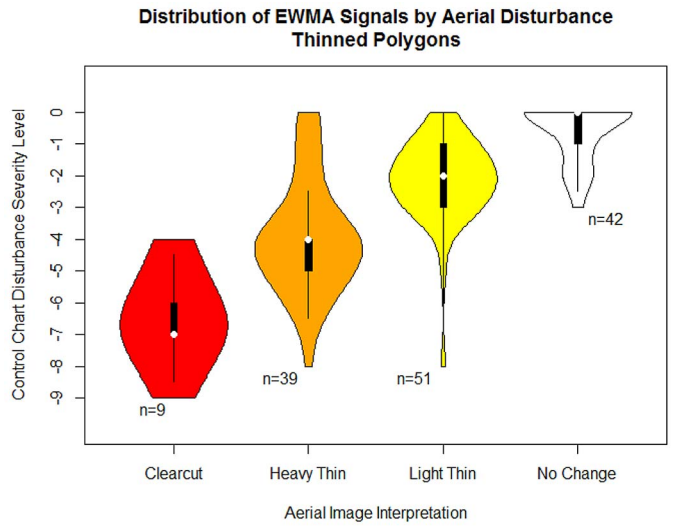

Fig. 11. Comparison of algorithm disturbance level with that observed in aerial imagery. In violin plots [28], the black box and whiskers correspond to the usual first and third quartiles, with the white dot representing the median. The thickness of the violin corresponds to the prevalence of data at the vertical value, similar to a histogram. Note the clear positive association between the EWMA outputs and the aerial categories.

image comparison. In this case, we wanted to look for associations between the two methods. Because we were interested here in forest degradation only, we again treated the EWMA chart values signaling for growth in the pixels as "no change." The results are displayed in Fig. 11. We used violin plots [28] to illustrate the distributions within each class.

It is clear in Fig. 11 that there is a simple relationship between the disturbance signal and the manner of disturbance. The Spearman correlation ([29], used here due to the ordinal nature of the data) was 0.753, highly significantly different from 0. Succinctly put, the greater the signal's deviation from zero, the more severe the disturbance tends to be. Based on this, in our case, we might use a rule of thumb that disturbance signals between -1 and -3 can be considered light thins, signals between -3 and -6 may be considered heavy thins, and signals beyond -6 tend to be clearcuts or wholesale removals. As the classifications in the assessment were in some sense arbitrary and based on regions that were thinned from mature stands, they may not reflect cutoff values across the remaining land use/land cover types in the scene and history. It is worth repeating that the values given by the algorithm are *relative disturbances* against that pixel's baseline history. Thus, perhaps the best way to treat the disturbance signals is as fodder for a "disturbance heat map," illustrated for the final date in the timeframe in Fig. 12. The long diagonal path of disturbance is that of a tornado that moved through the area on April 27, 2011 [30].

Indeed, by stacking the disturbance maps temporally, one may generate a "disturbance heat movie," allowing for an additional dimension to tracking the changes on the land. Note as well that, while the Westervelt polygons used here were only thins or removals of vegetation (shown as yellows, oranges, and reds), the image in Fig. 12 also indicates regions where the vegetation has increased beyond the training baseline (shown in greens). We offer a more specific example of this in Fig. 13, in which we observe development in a young pine stand. While we did not have truth data by which to statistically check for stand growth, the example strongly suggests that one may track afforestation as easily as deforestation with EWMA charts.

We reinforce the notion of signaling for forest growth with the example shown in Fig. 14. In the process, we also illustrate the manner in which harmonic regression captures intraannual variations to prevent false signals due to annually recurring disturbances such as plowing or planting of agricultural fields. In Fig. 14, we see a building with a young tree plantation to the north and an agricultural field to the south. Within a possible subpixel misregistration a few meters south and east, the EWMA algorithm precisely signals both the growth of the stand and the landscaping around the building, while offering no signal for the agricultural fields. It is worth recalling that the EWMA charts were performed on each pixel independently, with no masks to filter results. The algorithm does not signal the fields because the variations are intraannual and periodic, not interannual. While we display the signals only for one date here, the same pattern was evident throughout almost all dates of the testing period, with slight signal variation around the edges of the field and stand. Recall that this is only an example, because our validation data for severity were only for harvested pine stands. However, the example offers compelling evidence of the algorithm's general utility.

As one final example of the sensitivity of the method, we present an area in which two major changes take place from 2009 to 2011 (see Fig. 15). In this example, the western part of the area is a maturing young pine plantation, whereas the eastern side undergoes removal and conversion to a field.

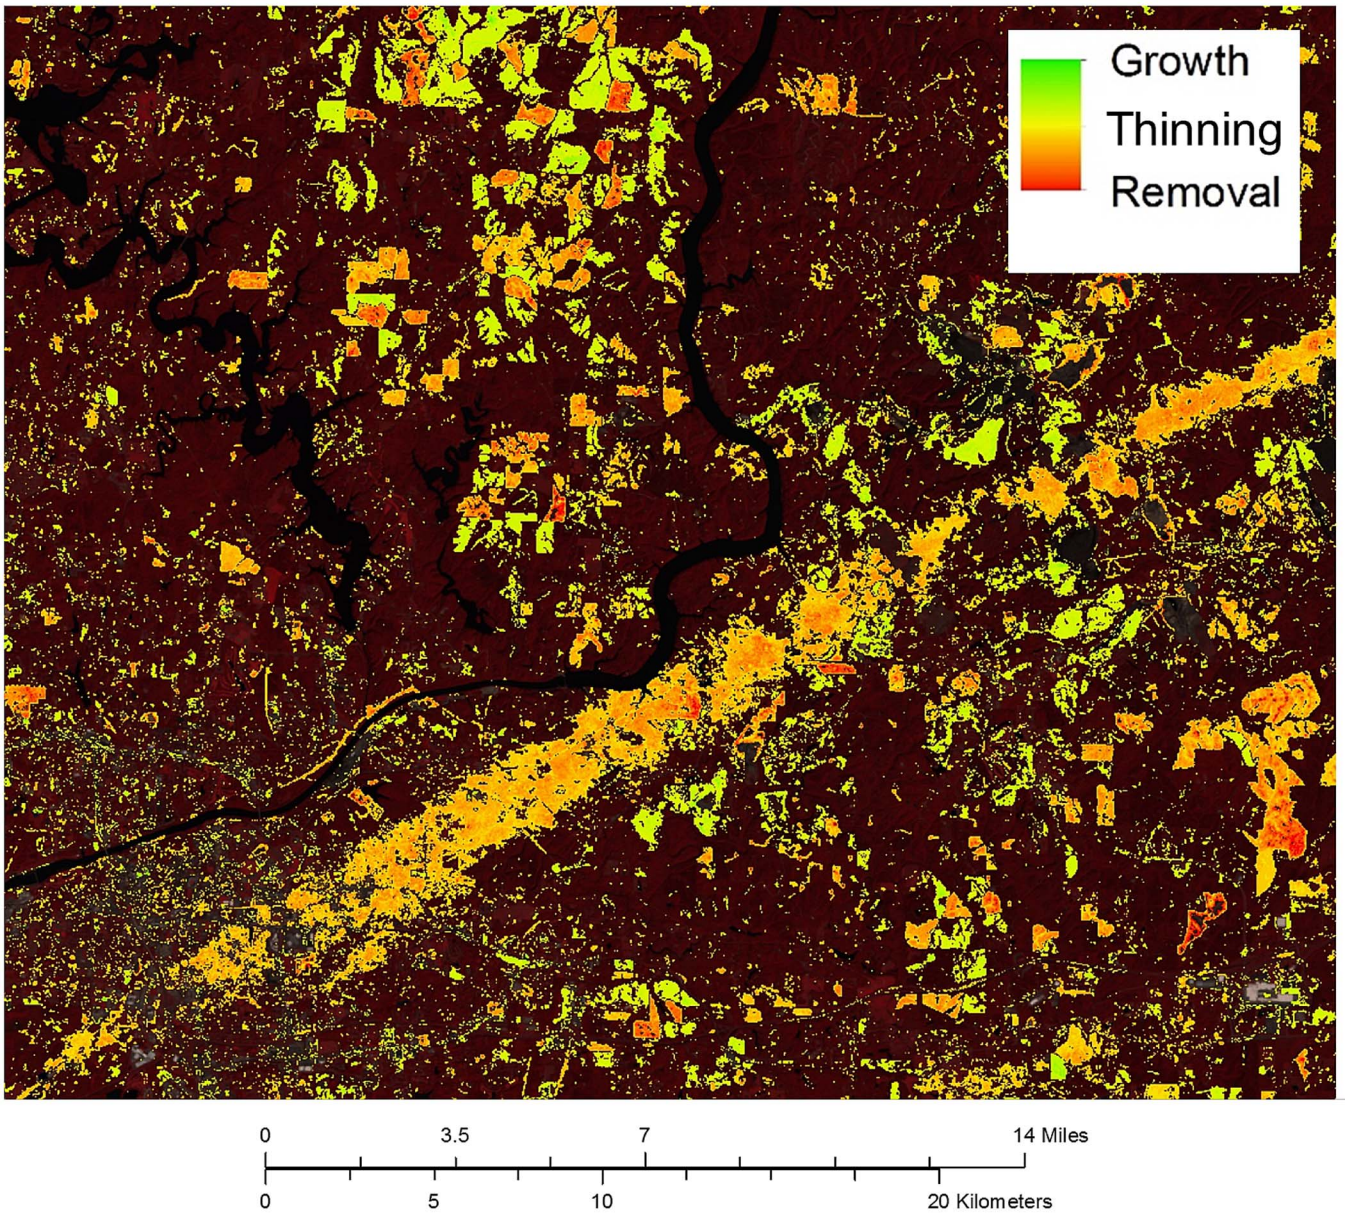

Fig. 12. Disturbance magnitudes for 10/3/2011. The town of Tuscaloosa, AL, USA, is in the lower left. The diagonal linear feature is a tornado path [30].

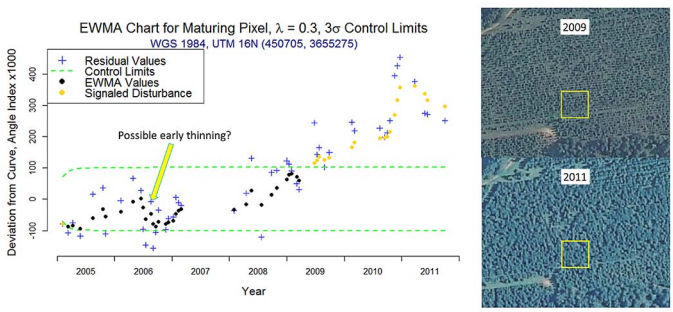

Fig. 13. EWMA chart for a maturing pine stand pixel. The images to the right depict the pixel in question (yellow box) according to the aerial images.

While no documented thinning occurs in this region (the region does not include any Westervelt polygons), it is clear that the algorithm is signaling according to the severity of the disturbances. In particular, note the way that colors in the eastern side correspond to the degree of the removal, with the harshest removals taking place where a dirt road was created (right-center of the region, red signals). It is also interesting to observe how the algorithm ignores areas of no change, such as a clearing in the maturing pine stand (lower left of the region, empty pixel). From this example and the preceding ones, it is clear that EWMA detection is appropriately signaling the severity of the disturbances.

C. Accuracy Assessment (Time)

We were able to use the aerial imagery from 2009 and 2011 to assess the effectiveness of EWMA charts in accurately signaling both the location and severity of disturbances as subtle as light thinning, also showing in the process that the charts can signal for forest growth as well. However, these results were in some sense limited by focusing only on the final image in the timeframe. We did this in order to most effectively utilize the

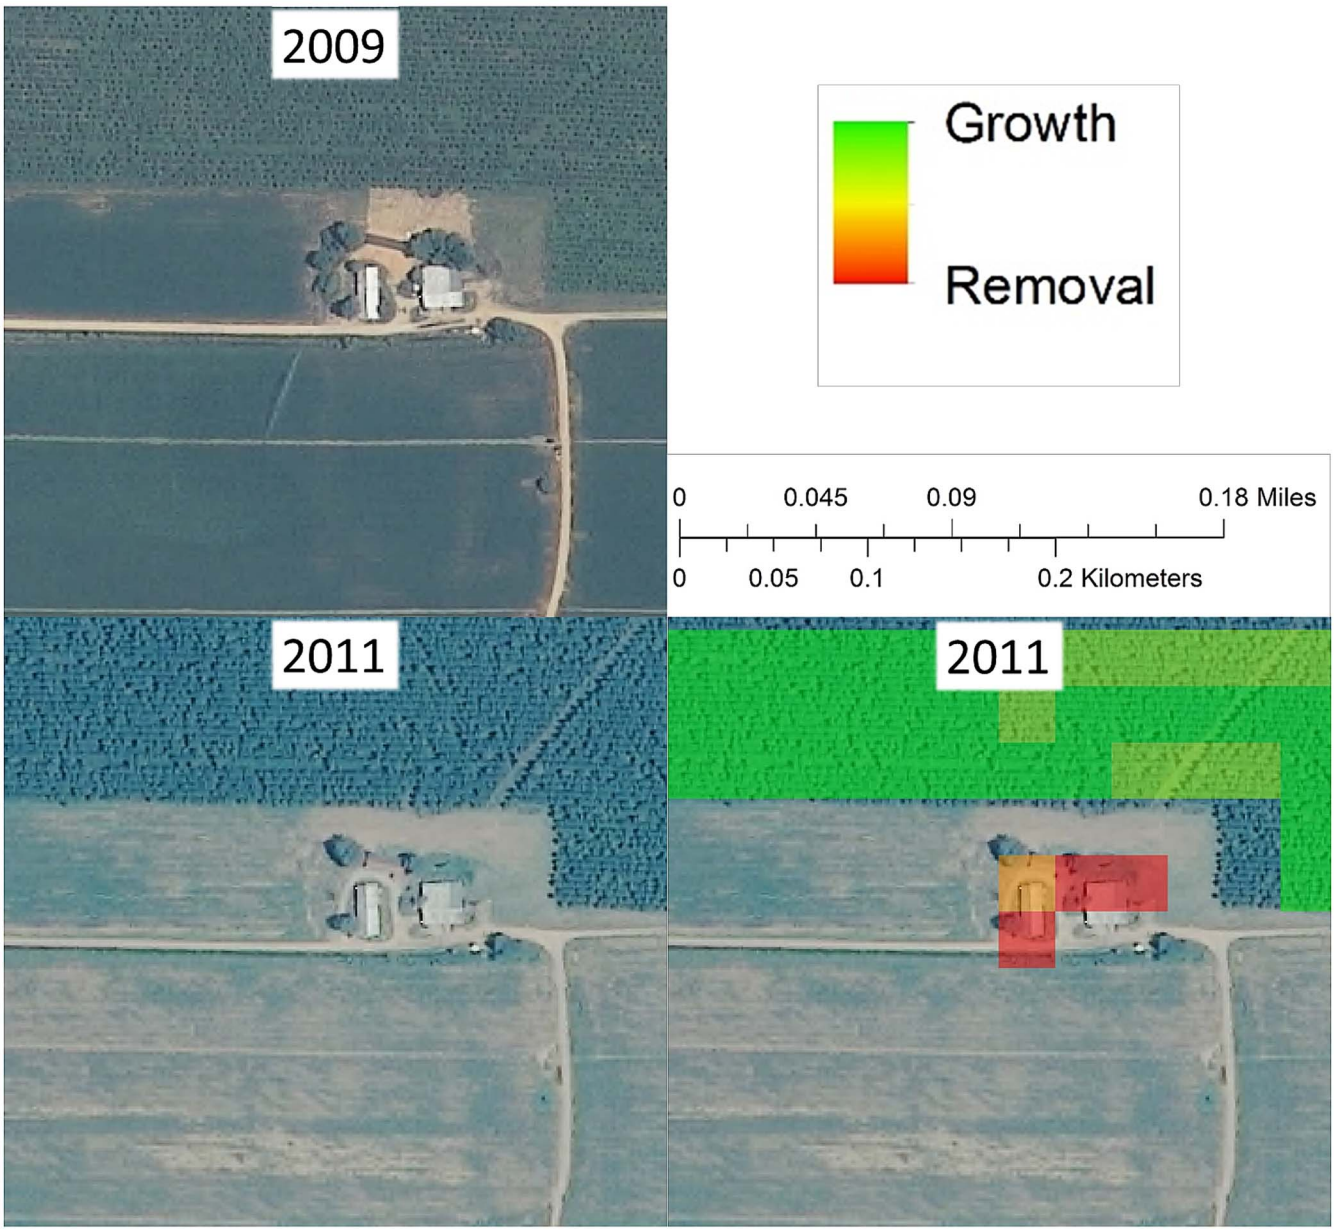

Fig. 14. Region with both agricultural and silvicultural activities. The color code of the EWMA signals is exaggerated slightly for easier interpretation. Note the precision of the signals and the lack of signals for the agricultural fields.

2011 aerial photographs, but the assessment thus far does not speak for the responsiveness of the EWMA charts, outside of the graphical examples.

Fig. 16 depicts a subset of our study area, in which we move our focus from the *magnitude* of the disturbances to the *timing*. For this figure, we show the EWMA disturbances by the year in which they began signaling for at least four images. This consistency constraint is a simple way to avoid displaying anomalous false signals, such as those caused by the fringes of clouds and shadows, that may have slipped through the cloud filter in the algorithm. This approach of seeking persistent change has been used by other methods [2], [3], [5]–[7], [11].

It is apparent in Fig. 16 that the EWMA algorithm does quite well in identifying the disturbances in the Westervelt polygons in the year they happen. The exceptions in this figure derive from stands that were harvested after the final image in the timeframe. In the case of the polygon in the northwest, the start date for the harvest was in late 2010, but the actual harvest extended into January of 2011, according to the records.

To get a sense of the finer responsiveness of EWMA charts, we consider a Westervelt polygon recorded as being thinned in September and October of 2010, as shown in Fig. 17. What is impressive about this polygon is that we can observe not only when the polygon was thinned, we can see the order in which the harvesters performed the thin. Through the use of an existing road (circled in the figure), it is clear that the harvesters opened up the center ahead of schedule before moving first east and then west of this point. This represents near real-time responsiveness, and we stress that the harvest was a thinning, subtle enough to be difficult to detect in its own right. Note that the stand to the northwest of the image was not in the Westervelt records, but it was harvested in early 2010.

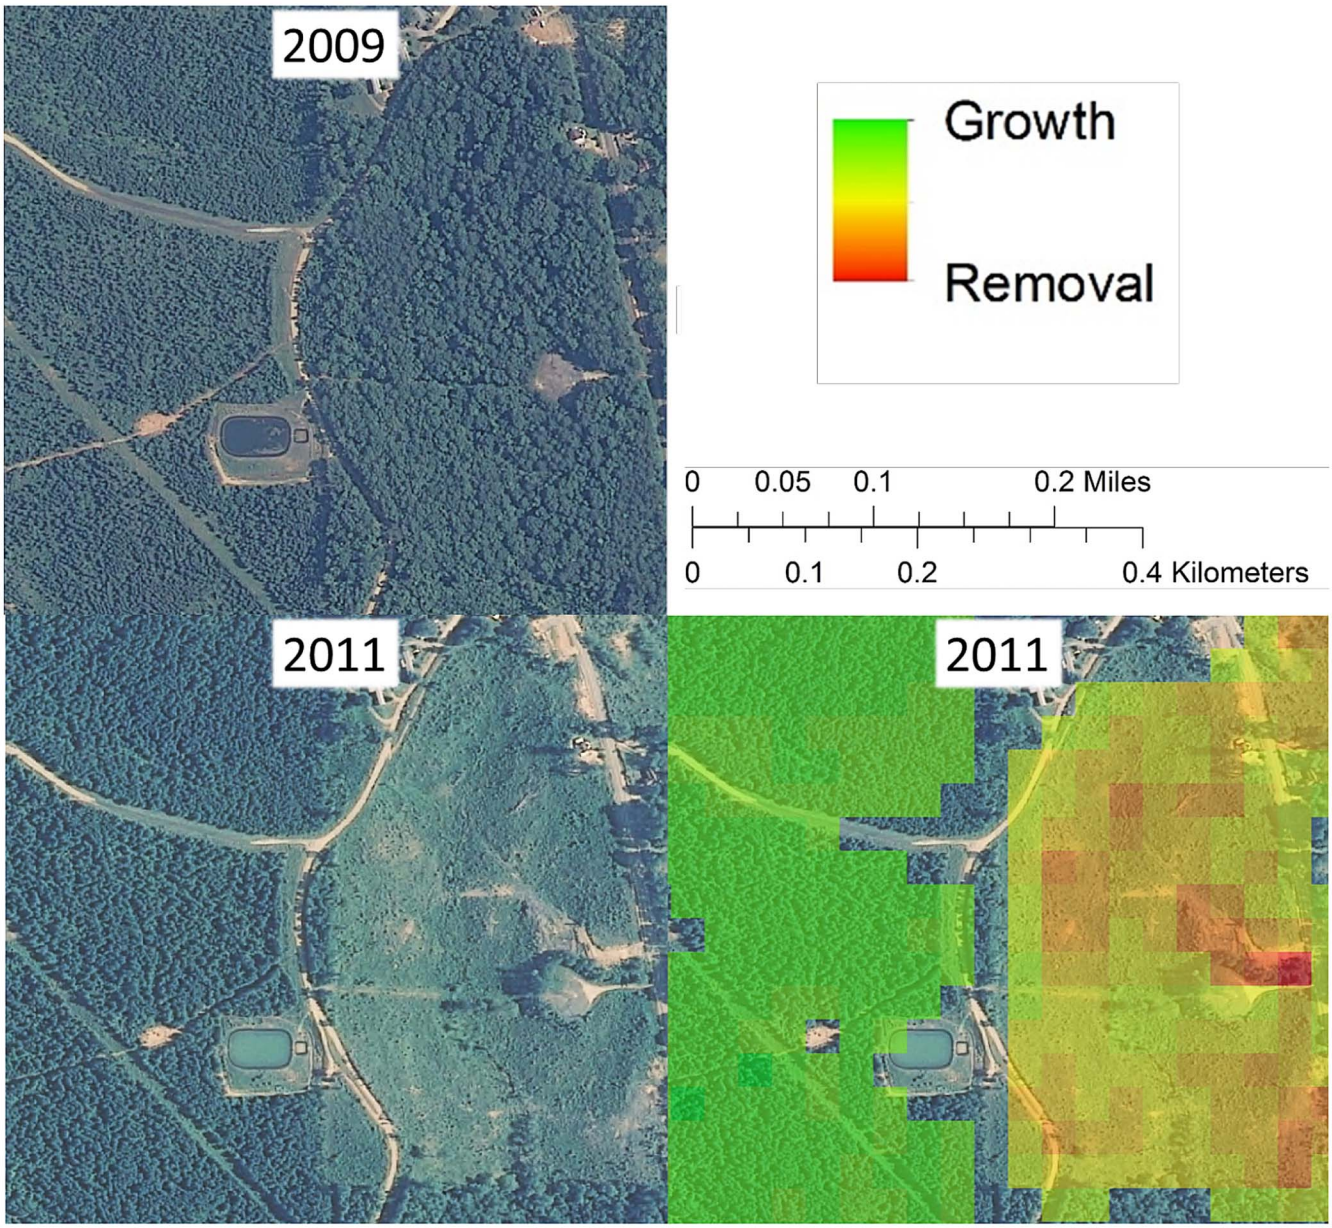

Fig. 15. Region undergoing both stand maturation and stand removal, illustrating the manner in which the EWMA charts indicate severity of disturbances. The pixel colors are exaggerated slightly for easier interpretation.

Because the thinnings typically took place over a range of dates within each polygon, it was difficult to concisely summarize the results of a temporal accuracy assessment. We will simply state that, for the accuracy assessment pixels for which the stands underwent basic harvesting according to the Westervelt records, over half of the EWMA charts for those pixels signaled within the first two images after the recorded start date. This, coupled with the aforementioned examples, leads us to conclude that the EWMA chart method is very responsive in the time dimension, allowing near-immediate signaling for a wide range of disturbances. Given that the charts also accurately signal the magnitude of the disturbance (or growth) and easily incorporate new images as they arrive, we conclude from the aforementioned assessments that EWMA charts are particularly well suited for on-the-fly environmental monitoring.

V. Discussion

It is well worth reviewing a few key points or challenges to the change detection method presented here. Perhaps the first of these may be the nature of our sampling and the validation data we used. In particular, each pixel we sampled was from a polygon designated to have been thinned in the testing period. Theoretically, this allowed us to test the algorithm's sensitivity to known changes, but it allowed no inference regarding the algorithm's specificity. That is, we had no control group of undisturbed pixels with which to look for false alarms.

Interestingly, of the 141 pixels thus sampled, 42 of them showed no visual sign of thinning or other vegetative removal. As we discussed in the results section, there are likely a variety of reasons for this high proportion of unthinned pixels in a collection that was supposed to be all thinned. These reasons

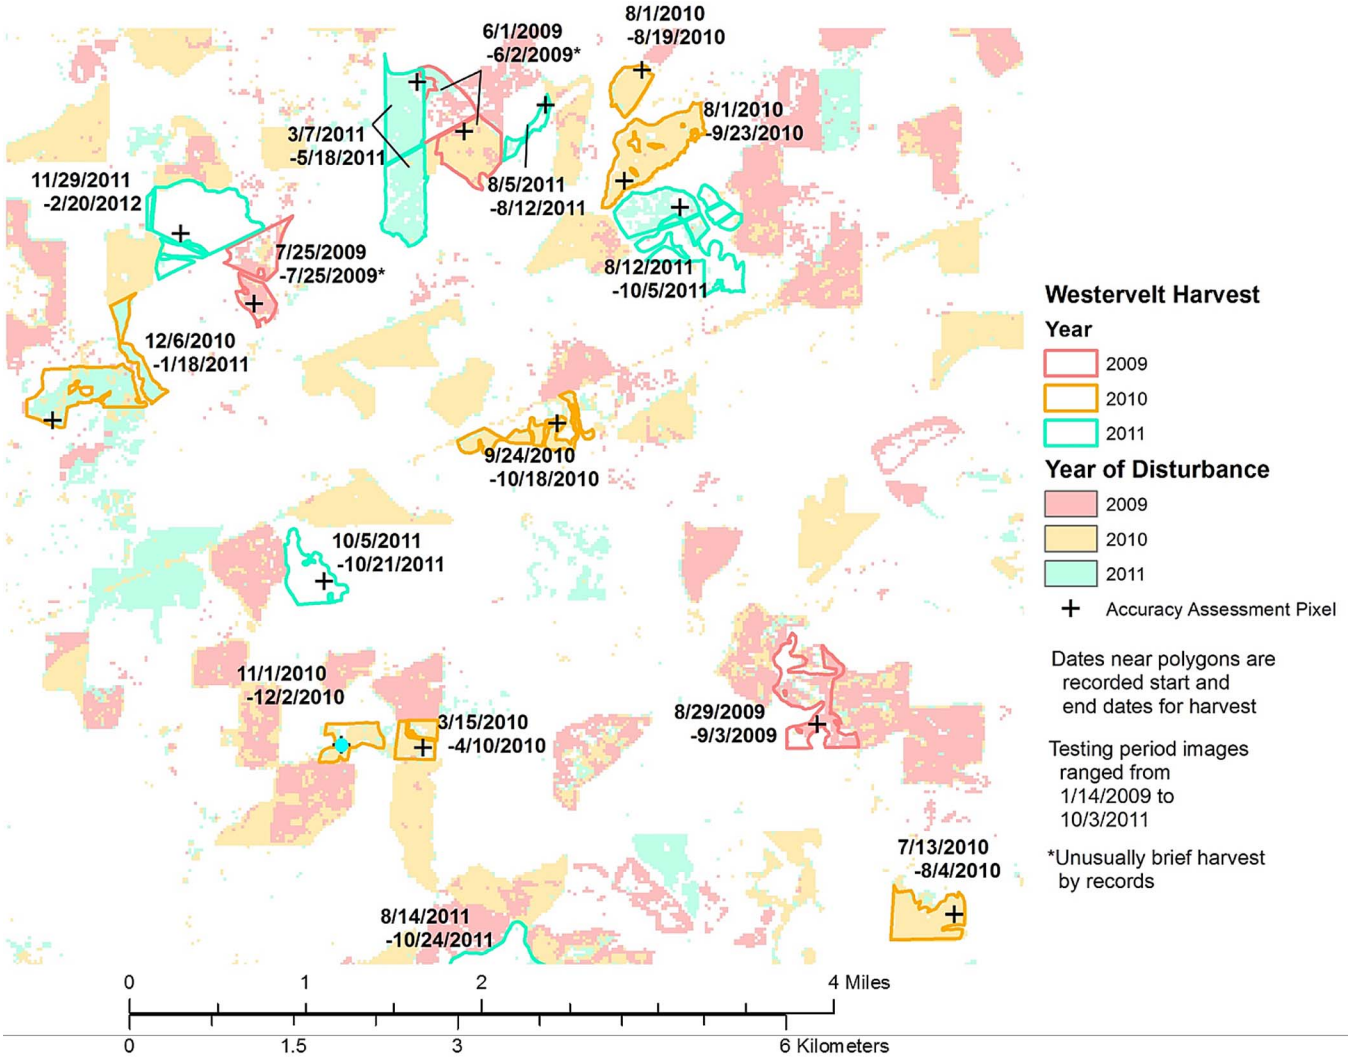

Fig. 16. Disturbance map based on the year of measured disturbance. The polygon outlines are pine stands harvested by the Westervelt Company.

appeared mostly to do with thinnings taking place at the fringes of the testing period, beyond the span of the images. As a result, we have a nice heuristic view of the algorithm's ability to detect false positives. While we cannot make any rigorous statements regarding the specificity of the algorithm here, the results appear promising.

Another issue that arises is that of independent sampling. Clearly there is temporal dependence from image to image. This is the reason that we used harmonic regression to remove all seasonal or periodic influences from the time series, resulting in a set of residuals that is temporally uncorrelated over the training period. In the event of a disturbance (possibly including stand growth), these residuals lose their uncorrelated nature, tracking the change as it occurs. This is true if a pixel undergoes disturbance during the training period as well, with the result that the harmonic regression centers the residuals as best it can on the historical mean value. In either case, the time series appears to violate the original assumptions behind the method.

This, however, is desirable. Within a control chart framework, when everything is in process, errors behave as assumed, in this case independently and identically normally distributed. In effect, the control chart is testing the validity of this assumption at each successive measurement. When the assumptions fail to convincingly explain the observed measurements, the chart signals that the process is out of control. Thus, the method outlined here really extracts meaningful information (vegetative disturbance and severity) from those places where the harmonic regression did not adequately model the pixel.

Another question that may arise is that of how to discriminate between changes in forest land cover as opposed to changes in (or to and from) other land cover classes such as agricultural and urban land covers. Because the EWMA detection algorithm operates on each pixel independently of the others, and because the magnitude of signaled disturbances is relative to the conditions of the pixel during the training period, we did not make any attempt to discriminate between the land cover classes in this paper. As shown in Fig. 14, provided there is no change in the way the fields are managed, the algorithm generally does not signal fields as changing. This represents a major strength of the algorithm and suggests that it may offer a way to capture land use and land use change, as opposed to land cover change. If one was making estimates of change to forest cover over

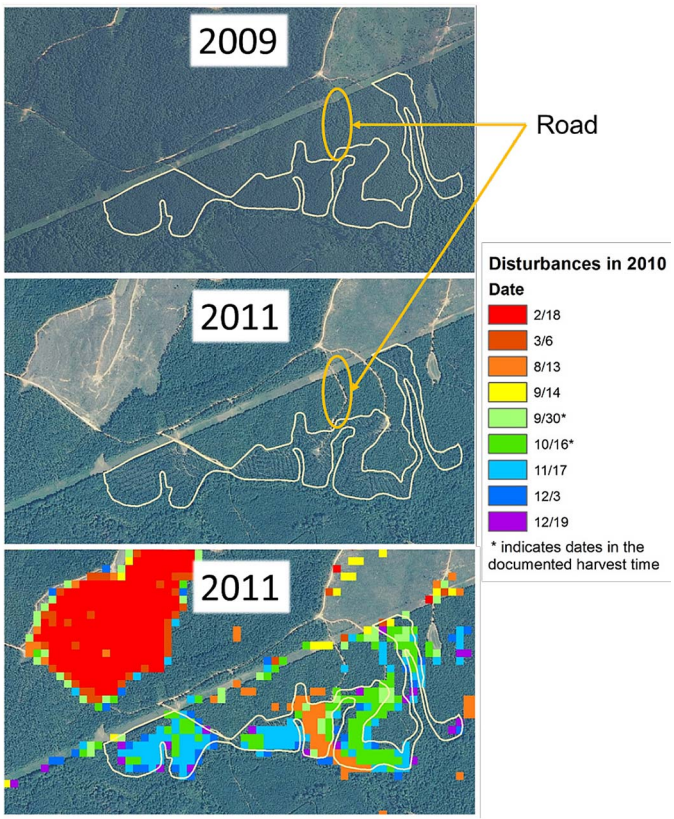

Fig. 17. Disturbances based on date in 2010 for a pine stand undergoing thinning in September and October. There is a clear pattern of harvesting in the east side of the stand before the west side.

a specific area, then one would need a forest-nonforest mask during the training period in order to measure the appropriate areas.

One other question that may arise from this particular study is that of how much training data the EWMA detection algorithm needed. It was shown in [12] that harmonic regression requires at least one image from "key dates" of the year, typically corresponding in a forestry context to phenologically important times such as green-up in the spring and senescence in the fall. In that sense, our use of four years of training data was very generous, being largely a function of having the data on-hand. We are confident that the EWMA detection algorithm could run well off of a single year of training data, and we intend to demonstrate this in future studies.

There are several other areas for future work on and improvement for this method. For example, one of the underlying assumptions behind the EWMA and Shewhart charts is that the sampling is done systematically, with equal temporal intervals between measurements. Owing to the nature of the Landsat stacks after removing pixels or images that are unsuitable for analysis, this assumption is generally not met. In the preceding research, we simply allowed the EWMA charts to act as if the measurements were still systematic, with the useful results presented here. However, it should be possible to modify the chart's construction to allow relative weighting to incoming data, relative to their temporal context. This modification would allow the control chart to work with less lag time, which is important, particularly when data are sparse. For example, after a period of extended silence, a new data point may be given particularly heavy weight; conversely, if measurements become particularly dense (for example, if one uses data from multiple platforms), then it may do to reduce the impact of each successive point. These are issues of appropriateness of temporal scale and were not addressed in this paper, but they are certainly areas worth exploring.

The algorithm as presented here did not use ancillary cloud masks, but it could easily incorporate them. Doing so would improve the quality of the results further and simplify the preprocessing by screening out troublesome data quickly. However, the algorithm incorporates a pixel-specific cloud filter in the form of the low-threshold X-bar charts. Provided that the training data are sufficiently clean to initialize the algorithm with a curve that models the phenology more than the noise, the built-in cloud mask appears to be sufficient.

Another question that arises is what to do with pixels which have signaled. At some point, the pixel needs to be reevaluated and reset to a no-change status, presumably with a new baseline. For this, sufficient retraining data need to be collected, ideally long enough after the disturbance that the pixel has stabilized somewhat. This is a constraint on any continuous monitoring approach to change detection, such as [11], in addition to the EWMA detection method. In other words, the method may be refined by allowing a "penalty period" after a change has been signaled, during which time incoming images may be used to recalibrate the baseline harmonic coefficients. This would enable the algorithm to be run continuously and to be adaptive, so that old changes do not linger until, or unless, the original baseline is restored.

One last area for future work presented here (there are undoubtedly many other areas not mentioned) would be expanding the application of the algorithm to other areas of the world. Our study area was restricted to pine stands in the southeastern United States. The algorithm itself is general in its operation, but it depends on having sufficiently dense image stacks on which to train the harmonic regression. In areas where the stacks are sparser, it is uncertain how long the training period would need to be extended to get a workable baseline. With regard to other land surface features, the algorithm is not confined to vegetation-specific indexes, and hence, one might expect it to perform well across different features, provided that a relevant band or index is used. All of these challenges and questions provide fertile ground for further work.

VI. Conclusion

Since the Landsat archive became available at no cost to the user, there has been a surge of interest in leveraging the temporal richness of the archive to aid in change detection [2], [3], [5]–[7], [9]–[11]. These massively multitemporal approaches represent a fundamental shift in the paradigm of change detection. Rather than treating images as sparse and isolated snapshots, researchers are able to treat them as statistical observations of underlying land surface processes. Accordingly, it is natural to utilize well-developed methods like quality control charts in this area.

When working in the context of change detection using large stacks of images, there are a number of algorithm attributes one may consider desirable. First, accuracy in the signaled changes is requisite. Additionally, the ability to detect more subtle changes, in our case forest thinning in addition to clearcuts, is very desirable. One might also wish to be able to screen for changes from recently acquired Landsat images, providing an up-to-date change map. It would be even more efficient to be able to do this using older maps as the basis for this change detection, yielding a dynamic and iterative process that refreshes itself with minimal additional processing. The EWMA detection algorithm, a method of using EWMA charts on the residuals of harmonic regression fits, possesses all of these features, at least, in the study area presented here. We have shown it to be accurate in this paper in terms of space, severity, and time. It not only detects thinning, it can also discriminate between differing degrees of disturbance, in both afforestation and deforestation, Additionally, because each pixel is treated independently and changes are measured relatively, the EWMA charts presented here appear to detect vegetative changes for other land use and land cover classes, as shown in Fig. 14. While our accuracy assessment and validation did not include these other classes, the records yielded by the algorithm invite future study in this area. The charts typically signal these disturbances in the images immediately following the disturbances. By detecting subtler changes in forest cover, it allows for improved estimates in forest parameter changes across regional or broader levels. Given its sensitivity in the temporal and severity dimensions, it is possible that the algorithm may benefit monitoring and response to migrating insect species, although this remains to be tested. It may also permit more responsive tracking of vegetation stress due to drought, perhaps aiding in determining the greatest threats due to wildfire in a region, although again this has not been tested yet. The method operates in a data-driven manner, requiring only a global specification of the number of harmonics to be used in baseline estimation and the tuning parameters for the charts across the entire scene. By virtue of the design of EWMA charts, the process can easily incorporate new images without reprocessing all of the history. This is possibly the most powerful feature of the method. It allows a smaller archive of data to be stored for tracking recent changes in the landscape, as all the historical information is preserved in the most recent disturbance map. All of these properties, taken together, indicate the great potential of this massively multitemporal on-the-fly method.

ACKNOWLEDGMENT

The authors would like to thank the researchers behind the NASA Earth Exchange (NEX), Ames Research Center, for their contribution of computing resources and programming feedback, specifically J. McCreight for his help in parallelizing R efficiently; K. Sorensen for his work in preprocessing the Landsat data through LEDAPS; B. De Haan of The Westervelt Company for providing us with the polygons used to check the algorithm; and USGS for making the Landsat archive freely available; without that forward-looking decision, any sort of massively multitemporal method would be much costlier.

REFERENCES

[1] B. E. Law and M. E. Harmon, "Forest sector carbon management, measurement and verification, and discussion of policy related to climate change," *Carbon Manage.*, vol. 2, no. 1, pp. 73–84, Feb. 2011.

[2] G. P. Asner, D. E. Knapp, E. N. Broadbent, P. J. C. Oliveira, M. Keller, and J. N. Silva, "Selective logging in the Brazilian Amazon," *Science*, vol. 310, no. 5747, pp. 480–482, Oct. 2005.

[3] G. P. Asner, D. E. Knapp, A. Balaji, and G. Páez-Acosta, "Automated mapping of tropical deforestation and forest degradation: CLASlite," *J. Appl. Remote Sens.*, vol. 3, no. 1, p. 24, Jan. 2009.

[4] J. Penman, "Definitions and methodological options to inventory emissions from direct human-induced degradation of forests and devegetation of other vegetation types," Inst. Global Environ. Strategies (IGES), IPCC, Geneva, Switzerland, 2003.

[5] MDA Federal, Inc., persistent change detection, May 3, 2013. [Online]. Available: http://mdaus.com/Geospatial/Global-Change-Monitoring.aspx

[6] C. Huang, S. Goward, K. Schleewis, N. Thomas, J. Masek, and Z. Zhu, "Dynamics of national forests assessed using the Landsat record: Case studies in eastern United States," *Remote Sens. Environ.*, vol. 113, no. 7, pp. 1430–1442, Jul. 2009.

[7] C. Huang, S. Goward, J. Masek, N. Thomas, Z. Zhu, and J. Vogelmann, "An automated approach for reconstructing recent forest disturbance history using dense Landsat time series stacks," *Remote Sens. Environ.*, vol. 114, no. 1, pp. 183–198, Jan. 2010.

[8] M. Li, C. Huang, Z. Zhu, W. Wen, D. Xu, and A. Liu, "Use of remote sensing coupled with a vegetation change tracker model to assess rates of forest change and fragmentation in Mississippi, USA," *Int. J. Remote Sens.*, vol. 30, no. 24, pp. 6559–6574, Dec. 2009.

[9] R. Kennedy, Z. Yang, and W. Cohen, "Detecting trends in forest disturbance and recovery using yearly Landsat time series: 1. LandTrendr—Temporal segmentation algorithms," *Remote Sens. Environ.*, vol. 114, no. 12, pp. 2897–2910, Dec. 2010.

[10] W. Cohen, Z. Yang, and R. Kennedy, "Detecting trends in forest disturbance and recovery using yearly Landsat time series: 2. TimeSync—Tools for calibration and validation," *Remote Sens. Environ.*, vol. 114, no. 12, pp. 2911–2924, Dec. 2010.

[11] Z. Zhu, C. E. Woodcock, and P. Olofsson, "Continuous monitoring of forest disturbance using all available Landsat imagery," *Remote Sens. Environ.*, vol. 122, pp. 75–91, Jul. 2012, Landsat Legacy Special Issue.

[12] E. Brooks, V. Thomas, R. Wynne, and J. Coulston, "Fitting the multitemporal curve: A Fourier series approach to the missing data problem in remote sensing analysis," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 9, pp. 3340–3353, Sep. 2012.

[13] W. A. Shewhart, *Economic Control of Quality of Manufactured Product*. Milwaukee, WI, USA: Amer. Soc. Qual. Control Qual. Press, 1980.

[14] M. R. Reynolds, Jr. and G. Y. Cho, "Multivariate control charts for monitoring the mean vector and covariance matrix with variable sampling intervals," *Sequential Anal.*, vol. 30, no. 1, pp. 1–40, Jan. 2011.

[15] B. C. Gupta, F. W. Walker, and P. Bloomfield, *Statistical Quality Control for the Six Sigma Green Belt*. Milwaukee, WI, USA: Amer. Soc. Qual. Control Qual. Press, 2007.

[16] D. C. Montgomery, *Introduction to Statistical Quality Control*, 6th ed. Hoboken, NJ, USA: Wiley, 2008.

[17] J. Yu and J. Liu, "LRProb control chart based on logistic regression for monitoring mean shifts of auto-correlated manufacturing processes," *Int. J. Product. Res.*, vol. 49, no. 8, pp. 2301–2326, Apr. 2011.

[18] X. Ning, Y. Shang, and F. Tsung, "Statistical process control techniques for service processes: A review," in *Proc. 6th Int. Conf. Service Syst. Service Manage.*, 2009, pp. 927–931.

[19] S. Shah, P. Shridhar, and D. Gohil, "Control chart: A statistical process control tool in pharmacy," *Asian J. Pharmaceutics*, vol. 4, no. 3, pp. 184–192, 2010.

[20] S. H. Steiner and M. Jones, "Risk-adjusted survival time monitoring with an updating exponentially weighted moving average (EWMA) control chart," *Statist. Med.*, vol. 29, no. 4, pp. 444–454, Feb. 2010.

[21] USGS Global Visualization Viewer (GLOVIS), May 3, 2013. [Online]. Available: http://glovis.usgs.gov/

[22] USDA NRC Geospatial Data Gateway, May 3, 2013. [Online]. Available: http://datagateway.nrcs.usda.gov/

[23] J. G. Masek, E. F. Vermote, N. E. Saleous, R. Wolfe, F. G. Hall, K. F. Huemmrich, F. Gao, J. Kutler, and T.-K. Lim, "A Landsat surface reflectance dataset for North America, 1990–2000," *IEEE Geosci. Remote Sens. Lett.*, vol. 3, no. 1, pp. 68–72, Jan. 2006.

[24] E. P. Crist, "A TM tasseled cap equivalent transformation for reflectance factor data," *Remote Sens. Environ.*, vol. 17, no. 3, pp. 301–306, Jun. 1985.

[25] S. L. Powell, W. B. Cohen, S. P. Healey, R. E. Kennedy, G. C. Moisen, K. B. Pierce, and J. L. Ohmann, "Quantification of live aboveground forest biomass dynamics with Landsat time-series and field inventory data: A comparison of empirical modeling approaches," *Remote Sens. Environ.*, vol. 114, no. 5, pp. 1053–1068, May 2010.
[26] C. E. Blinn, T. J. Albaugh, T. R. Fox, R. H. Wynne, J. Stape, R. A. Rubilar, and H. L. Allen, "A method for estimating deciduous competition in pine stands using Landsat," *Southern J. Appl. Forest.*, vol. 36, no. 2, pp. 71–78, May 2012.
[27] A. C Rencher and G. B. Schaalje, *Linear Models in Statistics*, 2nd ed. Hoboken, NJ, USA: Wiley, 2008.
[28] J. L. Hintze and R. D. Nelson, "Violin plots: A box plot-density trace synergism," *Amer. Statist.*, vol. 52, no. 2, pp. 181–184, May 1998.
[29] M. Hollander and D. A. Wolfe, *Nonparametric Statistical Methods*, 2nd ed. Hoboken, NJ, USA: Wiley Interscience, 1999.
[30] National oceanic and atmospheric administration, Nat. Weather Service Weather Forecast Off., May 3, 2013. [Online]. Available: http://www.srh.noaa.gov/bmx/?n=event_04272011tuscbirm

Evan B. Brooks received the B.A. degree in mathematics and physics and the M.A. degree in mathematics from the University of North Texas, Denton, TX, USA, in 2004 and 2007, respectively, and the M.S. degree in statistics and the Ph.D. degree in forestry from Virginia Polytechnic Institute and State University, Blacksburg, VA, USA, in 2010 and May 2013, respectively.

He is currently a Postdoctoral Associate with Virginia Polytechnic Institute and State University. His future career plans include a combination of research and teaching in environmental and remote sensing applications of statistics and informatics. His particular research interests focus on time-series-based monitoring and communication of climate change issues.

Randolph H. Wynne (M'99) received the Ph.D. degree from the University of Wisconsin, Madison, in 1995.

He is a Professor with the Department of Forest Resources and Environmental Conservation, Virginia Polytechnic Institute and State University, Blacksburg, VA, USA, where he also serves as Co-Director of the Interdisciplinary Graduate Education Program in Remote Sensing and Co-Director of the Center for Environmental Applications of Remote Sensing. His current remote sensing research focuses on algorithm development, image time series analysis of forest disturbance and recovery, active remote sensing of forest biophysical parameters, ecosystem service assessment and modeling, and terrestrial ecology.

Dr. Wynne is a member of the Landsat Science Team.

Valerie A. Thomas received the B.Sc.(Eng) degree in environmental engineering from the University of Guelph, Guelph, ON, Canada, in 1996, the Advanced Diploma in remote sensing from the College of Geographic Sciences, Lawrencetown, NS, Canada, in 1998, and the M.Sc. and Ph.D. degrees from Queen's University, Kingston, ON, in 2001 and 2006, respectively.

In 2006, she was a Postdoctoral Fellow for Fluxnet-Canada with the Department of Geography, Queen's University. Since 2007, she has been an Assistant Professor with the Department of Forest Resources and Environmental Conservation, Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, VA, USA. Her research interests include the use of remote sensing to examine forest canopy physiology, structure, and function across space and time.

Christine E. Blinn received the B.S. degree in natural resource management from Rutgers University, New Brunswick, NJ, USA, in 1997 and the M.S. and Ph.D. degrees in forestry from Virginia Polytechnic Institute and State University (Virginia Tech), Blacksburg, VA, USA, in 2000 and 2005, respectively.

She is a Research Scientist with the Department of Forest Resources and Environmental Conservation, Virginia Tech. Her research interests include the application of geospatial technologies and statistics to forest resource management.

John W. Coulston received the B.S. degree in forestry, the M.S. degree in forest biometrics, and the Ph.D. degree in forest ecology and remote sensing from the North Carolina State University, Raleigh, NC, USA, in 1996, 1999, and 2004, respectively.

He is currently the Team Leader for methods and techniques research with the USDA Forest Service Southern Research Station, Forest Inventory and Analysis Unit, Knoxville, TN, USA. His research focuses on developing and testing quantitative techniques for broad-scale inventory and monitoring of forest resources.