SANDIA REPORT

SAND2011-6876
Unlimited Release
Printed October 2011

Spatial-Temporal Event Detection in Climate Parameter Imagery

Sean A. McKenna and Karen Gutierrez

Prepared by Sandia National Laboratories Albuquerque,
New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2011-6876
Unlimited Release
Printed October 2011

Spatial-Temporal Event Detection in Climate Parameter Imagery

Sean A. McKenna and Karen Gutierrez
Sandia National Laboratories
Albuquerque, New Mexico, USA

Abstract

Previously developed techniques that comprise statistical parametric mapping, with applications focused on human brain imaging, are examined and tested here for new applications in anomaly detection within remotely-sensed imagery. Two approaches to analysis are developed: online, regression-based anomaly detection and conditional differences. These approaches are applied to two example spatial-temporal data sets: data simulated with a Gaussian field deformation approach and weekly NDVI images derived from global satellite coverage. Results indicate that anomalies can be identified in spatial temporal data with the regression-based approach. Additionally, *la Nina* and *el Nino* climatic conditions are used as different stimuli applied to the earth and this comparison shows that *el Nino* conditions lead to significant decreases in NDVI in both the Amazon Basin and in Southern India.

This Page Intentionally Left Blank

Table of Contents

Table of Figures

Table of Tables

Introduction

A general problem in a number of fields is rapid and accurate identification of anomalous measurements from time series data. The problem of anomaly detection is complicated when the time series is composed of more than one variable (multi-variate) and becomes even more complex when the multiple variables make up a random field with spatial correlation. The temporal and spatial correlation of the signal rules out the application of standard statistical tests for change detection. Change detection in these spatial-temporal data sets has received considerable attention over the past 15-20 years within the medical imaging research community (Brett et al., 2003; Friston, et al., 1994 and 1995; Worsley et al., 1994 and 1996) and a significant development of this research has been Statistical Parametric Mapping (SPM).

The practice of statistical parametric mapping has been developed in the field of medical imaging, particularly in brain imaging and in the practice of functional magnetic resonance imaging (fMRI) of the brain while the subject is performing various tasks (functions). Friston et al. (1995, p. 190) provide a concise definition of SPM: "*one proceeds by analyzing each voxel using any (univariate) statistical parametric test. The resulting statistics are assembled into an image, that is then interpreted as a spatially extended statistical process*". In other words, each pixel (voxel) in an image can be analyzed using a univariate statistical test (e.g., t-test) and the resulting values of the test statistic at each pixel are then displayed as a map. The underlying spatial correlation of the image is reflected in the map of the test statistics. This resulting map is analyzed using theory underlying stationary Gaussian fields and techniques developed for excursion sets of these fields. Properties of truncated Gaussian fields (e.g., Adler and Hasofer, 1976; Adler, 1981; Adler and Taylor, 2007; Adler et al., 2009) serve as the basis of the SPM techniques.

In this work, the techniques developed for medical imaging are applied to analysis of spatial-temporal fields measured from satellite imagery of the earth and focused on images of properties that are related to climatic variables. In particular, this study is focused on analysis of spatial-temporal images of vegetation density. Several significant differences are drawn between the work of this study, focused on identification of change detection in climatic indicators and the body of work in medical imaging:

1) Symptoms of climatic change are expressed by a single patient. While this sounds obvious, there is a single earth available for study and elements of clinical trials and experimental design that are important aspects of medical imaging (e.g., replicates, treatments, pre-existing conditions, various covariates, etc.) are not always available in the study of the earth. Additionally, the lag time between cause and effect in the human brain is on the order of milliseconds and is used effectively to determine which regions of the brain respond to certain stimuli. An equivalent lag time between application of a stimulus to the earth and a detectable response in the system may be decades or centuries. As described below, several techniques are utilized to provide some semblance of replication and treatments for the earth in order to better determine relationships between applied stimuli and responses.
2) Functional imaging of the brain is centered on the collection and analysis of 3D images. Each image is discretized into a number of 3D voxels. The climatic

indicator data utilized here are observed from satellites, or derived from combinations of satellite imagery and other inputs to create 2D images discretized into a number of pixels.

To date, the SPM approach has not been applied outside of medical imaging, but it appears to be a technique that could be successfully applied in a number of areas. The goals of this work are to both understand the basis of SPM and then apply SPM to an environmental data set obtained from satellite-based remote sensing of the earth.

Anomaly Detection with Statistical Parametric Mapping

Anomaly detection is defined here as the identification of a region in time and/or space that is anomalous in its shape, size (duration) and/or values within the region (intensity). Two approaches to anomaly detection in spatial-temporal data sets are explored here: 1) Anomaly detection in an online mode where prior data are used to predict future values of the measured variable and anomalies occur in areas and/or times where the predictions are inconsistent with the corresponding measurements; 2) Anomaly detection in historical data sets where differences in some treatment or external forcing condition is suspected to cause a difference in the measured variable. The anomalies in this case are significant differences in measured variables taken observed with and without activation of the external condition.

The subject of this study is anomaly detection in imagery where each image can be considered a two dimensional "time-slice" of an observed variable. Multiple time slices of the same region are available and examined as part of this work. Observational data provide a single variable at each location in space and time and the observed data are assumed to be spatially and temporally correlated, and this correlation has to be taken into account in the anomaly detection. Two measures of anomaly detection are employed: omnibus and localized. Omnibus detection uses a set of calculations to determine if the current image, taken as a whole, is anomalous. Localized detection determines one or more locations within the image where the anomaly occurs. Both measures are examined here.

Anomaly detection is not done directly on the observed images but on a difference of images for each time slice. The differences are calculated in one of two ways corresponding to the online or historical comparison goals of the anomaly detection: 1) differences between an expected, or predicted, image and an observed image. For example, a statistical or mathematical model is used to predict the expected image at a future time step and the difference between this prediction and the image observed at that future time is calculated; and 2) differences of two average values where each average is calculated over a set of images collected under a specific condition. For example, in studies of the human brain, images are often collected under "resting" and "stimulated" conditions and the average image from each condition is then used to create a difference map.

In both cases, the end result is a difference between two images. These difference maps represent values of a t-statistic at each location. The t-distribution is symmetric, has a mean of zero and converges to the standard normal distribution for an infinite number of degrees of freedom. In practicality, for degrees of freedom > 100, the t-distribution is well approximated by the standard normal distribution. Analyses here are done on both the t-statistic map or on z-score maps where the t-statistic values have been transformed to standard normal values in keeping with the original development of the SPM approach. The transformation from t-statistic to z-score is accomplished using the appropriate degrees of freedom. In the terminology of random fields, the z-score, or standard normal maps, represent multiGaussian fields.

SPM was developed to directly address a problem in statistical testing. Direct application of most statistical tests requires independence of the observations, but for many problems, including those studied here, correlation between adjacent observations is generally the norm. Therefore, the results of the statistical tests for adjacent, or even nearby, pixels cannot be effectively evaluated using standard techniques. SPM considers a single map comprised of the results of all local statistical tests and provides a number of measures for comparison of the values in the map to a critical threshold level.

MultiGaussian Fields

The basis of the SPM approach is the analysis of the number, size and degree of excursions from a multiGaussian (mG) random field. For a concise, statistical description of mG fields, see Adler et al., 2009, page 27. In the testing and verification section of this report, mG fields are created and/or modeled as the convolution of a Gaussian kernel with an uncorrelated (white noise) field. The Gaussian kernel is defined as

$$\mathrm{G(x,y)} = \frac{1}{2\pi|\Sigma|^{1/2}} exp\left(-\frac{1}{2} d\Sigma^{-1}\mathrm{d^T}\right)$$

where d is the distance vector containing distances d_x and d_y from any location (x,y) to the origin of the Gaussian function x_0, y_0 (here (0, 0) for the standard normal distribution). In this work, the covariance matrix, $\Sigma=\sigma^2 I$, (where I is the identity matrix) is diagonal for the specific case of the kernel being aligned with the grid axes.

An often-used measure of the size of the Gaussian kernel is the "full width at half maximum" (FWHM):

$$FWHM = \sigma\sqrt{8\ln(2)}$$

If the mG field is not created, but is obtained from some type of imagery or other analyses, then there is no known underlying kernel and it is necessary to estimate the FWHM directly from the image. Estimation can be done using the variance matrix of the partial derivatives of the image values with respect to the discretization of the image. In 2D, the covariance matrix is:

$$\Lambda = \begin{bmatrix} Var\left(\frac{\partial T}{\partial x}\right) & Cov\left(\frac{\partial T}{\partial x}, \frac{\partial T}{\partial y}\right) \\ Cov\left(\frac{\partial T}{\partial x}, \frac{\partial T}{\partial y}\right) & Var\left(\frac{\partial T}{\partial y}\right) \end{bmatrix}$$

This covariance matrix can be seen as a measure of the roughness/smoothness of the image.

Estimation of Λ□□□□□□□□□□□□□□□□□□□□□□r□□□□□□□□□□□□□□□□□□□□□□h□□□and here the simple relationship defined by Worsley et al. (1992) between the FWHM values in each of the principal directions and Λ is utilized. However, if the image being analyzed has been collected, not created, the FWHM values are most likely unknown and this approach is not practical.

Another approach is numerical calculation of the derivatives within Λ. For each spatial dimension, differences between adjacent pixels are calculated as:

$$Z_{xi}(x,y) = \{T_i(x+\delta x, y) - T_i(x,y)\}/\delta_x$$
$$Z_{yi}(x,y) = \{T_i(x, y+\delta y) - T_i(x,y)\}/\delta_y$$

where δ_x and δ_y are the dimensions of the image pixels in the x and y directions. The variances and covariances of the differences are then used to approximate the variances and covariances of the derivatives:

$$V_{xx} = \sum_{i,x,y,z} Z_{xi}(x,y,z)^2/N(n-1)$$
$$V_{yy} = \sum_{i,x,y,z} Z_{yi}(x,y,z)^2/N(n-1)$$
$$V_{xy} = \sum_{i,x,y,z} \{Z_{xi}(x,y,z) + Z_{xi}(x, y+\delta_y, z)\}\{Z_{xi}(x,y,z) + Z_{xi}(x+\delta_x, y, z)\} \quad /4N(n-1)$$

These variance and covariance estimates are used to estimate Λ:

$$\Lambda = \begin{bmatrix} V_{xx} & V_{xy} \\ V_{xy} & V_{yy} \end{bmatrix}$$

Finally, the FWHM in the X and Y directions are calculated as:

$$FWHM_x = \sqrt{\frac{4\ln(2)}{V_{xx}}}$$

$$FWHM_y = \sqrt{\frac{4\ln(2)}{V_{yy}}}$$

Calculating the SPM

The Statistical Parametric Map is the difference image between individual pairs of images or average images, which is typically transformed from a map of t-statistics to a map of Gaussian z-score values. The different methods used in this study for calculating the SPM are described in this section.

Conditional Differences

Following the approach of Worsley, et al., 1992, to look for regions of brain activation in subjects under different conditions, the earth is considered as the single subject of the study. The earth is mapped under different conditions (e.g., background temperature anomaly and either *el Nino* or *la Nina* conditions) with the goal of detecting areas of activation under the stimulus by subtracting the image under the background condition. The matrices *A(x,y)* and *B(x,y)* are the average values at pixels with coordinates *(x,y)* under the two conditions A and B. The number of images that went into these average calculations may be relatively small (i.e., 4 to 10), but should be similar for the two conditions. The two average images are normalized to have the same mean value and then subtracted to provide normalized differences $\Delta(x,y)$:

$$\Delta(x,y) = \frac{B(x,y)}{\bar{B}} - \frac{A(x,y)}{\bar{A}}$$

where $\bar{A}$ and $\bar{B}$ are the spatial average of the observed variable values across all pixels in the study region under conditions A and B, respectively.

The t-test is a traditional measure of the difference between two means (e.g.,Walpole and Myers, 1989). Quite simply, the t-statistic is the difference between two values, at least one of which is a population or sample mean, normalized by the standard error of the mean:

$$t = \frac{\bar{X} - \mu}{s_e} = \frac{\bar{X} - \mu}{s\sqrt{1/n}}$$

where $\bar{X}$ is a sample mean, μ is a population mean, s_e is the standard error of the mean which is the standard deviation of the observations, s, that make up the data vector *X* multiplied by the square root of 1 over the number of samples within *X*. The cumulative probabilities for any value of *t* are available from the Student's t distribution and require knowledge of the degrees of freedom, ν, within the test. For the analyses done here, ν is generally n-1.

In the case of comparing two sample means to each other at each location ,i.e., A(x,y) and B(x,y), instead of comparing a sample mean to a theoretical population mean, the value of s_e must be calculated from both sample sets as:

$$s_e = s_p\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}$$

where n_1 and n_2 are the number of images that were used in calculating the average maps A and B and s_p is the average pooled standard deviation:

$$s_p(x,y) = \sqrt{\frac{(n_1 - 1)s_1^2(x,y) + (n_2 - 1)s_2^2(x,y)}{n_1 + n_2 - 2}}$$

Here we are assuming that n_1 and n_2 are constant for all locations and therefore not a function of (x,y). The t-statistic image (map), based on the pooled standard deviation, is:

$$t(x,y) = \frac{D(x,y)}{s_p(x,y)\sqrt{\frac{1}{n_1} + \frac{1}{n_2}}}$$

Here, a location (pixel)-based calculation of the standard deviation is used. Another approach is to calculate the pooled standard deviation across the image (image-based) and arguments for using the image-based standard deviation are given by Worsley, et al. (1992). Typically, the number of observations under each condition is small, less than a dozen, and therefore the effective degrees of freedom for $T(x,y)$ is generally small and needs to be used in the transformation of the t-field to a standard normal Gaussian field.

The cumulative probability of a t-statistic is found from the T distribution function with the appropriate degrees of freedom. This probability is then used with the inverse of the Gaussian distribution function to get the z-score value:

$$P(Y \leq y) = T(y;n)$$
$$z = G^{-1}(P(Y \leq y))$$

The resulting fields are now multiGaussian SPM's and the anomaly detection algorithms developed for SPM analysis can be applied

Regression Models

There are two goals for the imagery analysis in this study: Anomaly detection from historical data and online anomaly detection. To accomplish this second goal, it is necessary to investigate? anomalies in spatial-temporal data sets as those data sets become available. One widely used approach to anomaly detection is to predict the expected value of an observation prior to taking that observation and then examine the residual between the predicted and observed values. Here, non-linear regression models are used to make predictions of future values based on previously observed values.

In particular, we examine the following simple periodic model to predict values of a variable Z the next time step $(t + 1)$ from a data vector collected at previous times, t:

$$Z(t+1) = \beta_0 + \beta_1(t) + \beta_2 \cos\left(\frac{t}{dd}\right) + \beta_3 sin(\frac{t}{dd})$$

where *dd* is the period of the periodic process and the vector of regression coefficients, β's, are determined through solution of the linear system:

$$\hat{\beta} = XZ^{-1}$$

where $\boldsymbol{X}$ is a design matrix containing the operations in the regression model applied to all times *t* through *(t-P)* where *P* is the number of terms in the regression model and $\boldsymbol{Z}$ is the vector of the *P* previously observed values (data vector). The coefficients, β_0 through β_3, modulate the intercept, linear trend, cosine and sine components of the model, respectively.

Here, the regression model is utilized to predict the next time slice of the observed spatial-temporal process. Parameters of the regression model are identified independently for each spatial index (*i,j*) within the image at each time step. It is also possible to make the β vector a function of location and employ a model of spatial correlation between the regression coefficients to potentially improve the regression estimates, but that approach is left to a future study.

Calculation of the t-statistic is based on the assumption of independent samples within the vector *X*. Here, we will calculate the t-statistics using a regression model that exploits the temporal correlation within a one-dimensional data vector to estimate the value at the next time step. The regression estimate is a mean of the previous *n* observations, but these observations are not independent. In order to account for correlation within the data set, the standard error, s_e, within the t-statistic calculation is replaced with (Gilbert, 1987, p. 184):

$$s\left[\frac{1}{n}\left(1 + 2\sum_{i=1}^{n-1}\hat{\rho}_i\right)\right]^{1/2}$$

where the $\hat{\rho}_i$ are the values of the autocorrelation of the data set for each of the *i*th lag spacings. This replacement decreases the size of the resulting t-statistic to account for correlation inherent in the data set. The degrees of freedom, *v*, remains unchanged at *n-1*.

Anomaly Detection

Anomaly detection here is broken into two types: Omnibus and localization. The former indicates that an anomalous feature is present, but does not provide information on the location of the feature. The latter provides an approximate location for the anomaly.

Omnibus Detection

Omnibus detection is defined here as the ability to detect a change in an image relative to other images without any localization (i.e., change detection without determining the location of the change). Omnibus tests as defined here are single tests that pertain to the entire image. This definition is consistent with that of Fox et al. (1988) and Worsley et al. (1992) but not that of the more locally specific definition proposed by Friston et al. (1995).

Omnibus detection approaches are applied to the SPM, not to the images directly. Several approaches to omnibus change detection have been developed in the brain imaging literature and two of them are evaluated here: 1) the significance of the maximum value in the image, T_{max}, relative to a threshold value; and 2) the number of distinct regions in an image that are above a threshold value.

In this discussion, the calculated value of T_{max} and the resulting p value can be calculated for either a t-statistic field or a z-score field (standardized Gaussian). In the absence of any activation signal that would change the nature of the field, the T field can be well approximated by a mean-zero Gaussian field with unit variance and second-order stationarity. The maximum value of an mG field is defined by the Gaussian distribution and the level of spatial correlation in the field relative to the domain size. The probability that the maximum of an mG field exhibiting spatial correlation is above a threshold, t, where t is relatively high (e.g., ≥ 3.0), has been the subject of study by a number of researchers. Both Adler and Hasofer (1976) and Adler (1981) have developed general expressions for the p-value of the maximum of an mG field in multiple dimensions. For the 2D case of concern here, Worsley et al. (1992) provide an expression for this p-value as:

$$P(T_{max}^{S} > t) \approx R^{S}(4\ln(2))(2\pi)^{-1.5}te^{-0.5t^2}$$

where T_{max}^{S} is the maximum of the 2D "slice" in the nomenclature of Worsley et al. (1992). The number of resolution elements, or *Resels*, R, is a measure of the spatial correlation of the field calculated as the number of FWHM values that fit within the study domain:

$$R = S/(FWHM_x \times FWHM_y)$$

where S is the area of the study domain. R provides the degrees of freedom in the T test. The p value is interpreted as the probability of obtaining T_{max}^{S} under two different null-hypothesis scenarios: 1) there is no difference between the map predicted by the regression model and the observed image; and 2) the case where there is no difference between the two average maps each calculated from sample maps under the two different conditions.

Isolated Regions of Activation

A second approach to omnibus detection is comparison of the number of regions within a map that exceed a given threshold level ("regions of activation") to a reference model of the expected number of such regions. Truncation of a Gaussian field at a threshold *t* defines the t-level excursion set:

$$X_t = \{ x \hat{1} \ R^D : Y(x) \, 3 \ t \}$$

A large body of literature on the properties of excursion sets (regions of exceedence) in Gaussian random fields is available (e.g., Adler, et al., 2009; Friston et al., 1994; Lantuejoul, 2002). From this area of research, Friston et al. (1994) characterize three related properties of excursion sets in truncated Gaussian random fields:

N the number of pixels above the truncation threshold, *u*,
m the number of distinct regions (inclusions) above the threshold, and
n the number of pixels in each region,

with expectation relationship *E[N] = E[m]E[n]*. For threshold value, *u*, the number of cells above that threshold, *N*, is provided by the Gaussian cdf and the size of the domain, *S*:

$$\mathrm{E[N]} = S \int_t^{\infty} (2\pi)^{-1/2} e^{-z^2/2} \, dz$$

A measure of the number of isolated regions above the threshold can be obtained from the Euler Characteristic, EC. In two dimensions, the EC represents the number of connected excursion sets in the domain minus the total number of holes within those sets. Therefore EC goes to 0.0 at t = 0 and EC becomes negative when t < 0.0 as the truncated field represents a single domain-spanning set containing a large number of holes. In 2D, and at relatively high truncation thresholds, EC is equivalent to the number of regions above the threshold, E[m].

$$\mathrm{E}[m] = EC = \left| (2\pi)^{-(D+1)/2} W^{-D} S t^{(D-1)} e^{t^2/2} \right|$$

where *W* is an alternative measure of the spatial correlation of the mG field defined as a fraction of the FWHM:

$$\mathrm{W=FWHM}/\sqrt{4 \ln(2)}$$

For a given threshold, u. the average area of the individual regions is found from the expectation relationship:

$$\mathrm{E[n]=E[N]/E[m]=E[N]/[EC]}$$

Figure 1, from McKenna et al (2011) compares a direct calculation of EC using the Matlab Image Processing toolbox (Matlab, 2009) with estimates made using the Euler characteristic equation above across a range of *u* values increasing from left to right. The corresponding binary fields (500 × 500 cells) are also shown for several representative threshold values. Note, that here we are interested in both extreme ends of the graph corresponding to *t* values (truncation thresholds) with absolute values of 2.5 or greater.

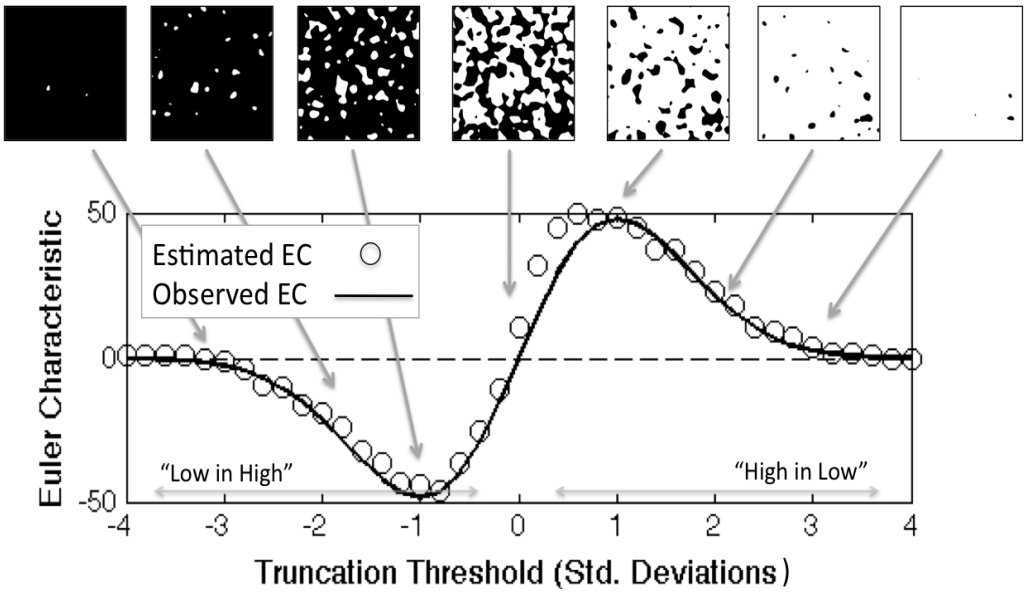

Figure 1. Observed (calculated) and estimated Euler characteristic for a truncated mG field as a function of the truncation threshold, *t*. The corresponding excursion sets for $t > 0$ are black regions in the binary fields at the top of the image (after McKenna, et al., 2011).

For each difference map and threshold *t*, the quantities *N, m* and *n* (Worsley, et al., 1994) are calculated directly and compared to the theoretical values as determined for an mG field with a FWHM as observed in the difference map. Additional omnibus tests are then simple comparison of the observed and theoretical values of these quantities:

$$N_{obs}(t, FWHM) > N_{theo}(t, FWHM)$$
$$m_{obs}(t, FWHM) > m_{theo}(t, FWHM)$$
$$\max(n_{obs}(t, FWHM)) >> n_{theo}(t, FWHM)$$

If the observed values of the number of pixels, Nobs, exceeding the threshold and/or the number of excursion regions, mobs, is larger than their theoretical counterparts, that serves as additional evidence for an anomaly. The third test is comparison of the maximum observed excursion size against the theoretical mean excursion size. No significance levels are applied to these tests here. Some work has been done on determining p-values for

these quantities (Worsley, et al., 1994) and other authors have used Monte-Carlo simulation to define approximate confidence intervals for these quantities. Here, we simply examine the ratios of the observed to theoretical quantities as indicators of anomalies.

Localized Anomaly Detection

Further analysis of the excursion sets leads to information on the location of the detected anomalies. The excursion set maps themselves can be examined to see where the excursions are occurring. The location of the centroid of the largest excursion is recorded to define the location of the anomaly.

Process Validation

Two fundamental aspects of the algorithmic approach defined above, calculation of the FWHM from an image and calculation of the p-value of the maximum t-statistic, are tested here to determine their range of applicability.

Validation of the FWHM Estimation

The relationships outlined above depend on accurate knowledge of the spatial correlation of the mG field as defined by the FWHM value. In the case of a constructed field, the FWHM is a known parameter. In the case of an observed field, the FWHM must be estimated directly from the field (image). Here, those estimation approaches are tested against fields created with known FWHM values.

A set of mG fields are created using the convolution approach for each level of increasing FWHM values. These mG fields are then used as input to the pixel differences approach to calculating FWHM and the actual and estimated FWHM values are compared (**Figure 2**).

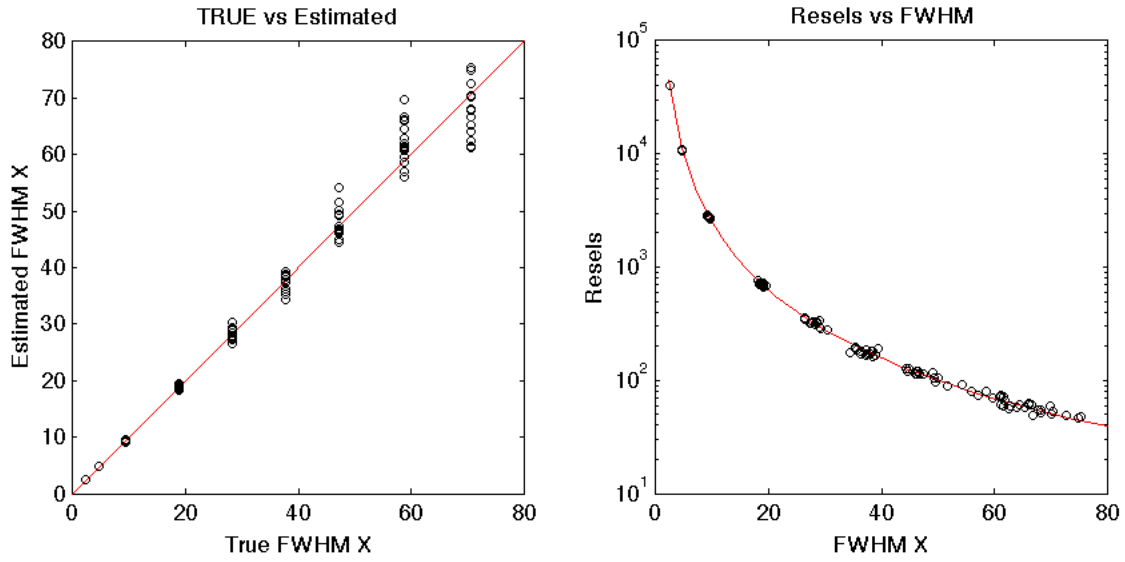

Figure 2. Results of testing the estimation of FWHM and the resulting number of resels (the red line in the right image defines the true resel values). Each circle represents a unique result from a generated mG field.

A range of σ values from 1.0 to 30.0 were evaluated. These σ values correspond to FWHM values ranging from 2.35 to 70.6 units. For each FWHM value, 15 realizations were run. The domain used in these evaluations is 500x500 units square. At FWHM values of 20 or less (4 percent or less of the domain size) the estimated FWHM values are accurate and very precise relative to the true FWHM values. At FWHM values larger than 20.0, the estimates become less precise, but they remain relatively unbiased. The increase in variation is due to ergodic effects as the correlation length (FWHM) of the field increases to become a larger fraction of the field size.

The right image of **Figure 2** shows that the increasing variation in the estimated FWHM at larger sizes is also reflected in increasing variation in the estimated value of *R* (resels). However, the right image of **Figure 2** shows that the estimates of *R* are robust across the range of FWHM values.

The same values of σ were evaluated again with anisotropic Gaussian fields. The level of anisotropy was set such that the FWHM in the X-direction is 5 times that of the FWHM in the Y-direction. The results of these validation runs are shown in **Figure 3**. The true FWHM values are accurately estimated in both directions. The true value of *R* is also estimated accurately; although, at the smallest FWHM values the true value of *R* is under-estimated.

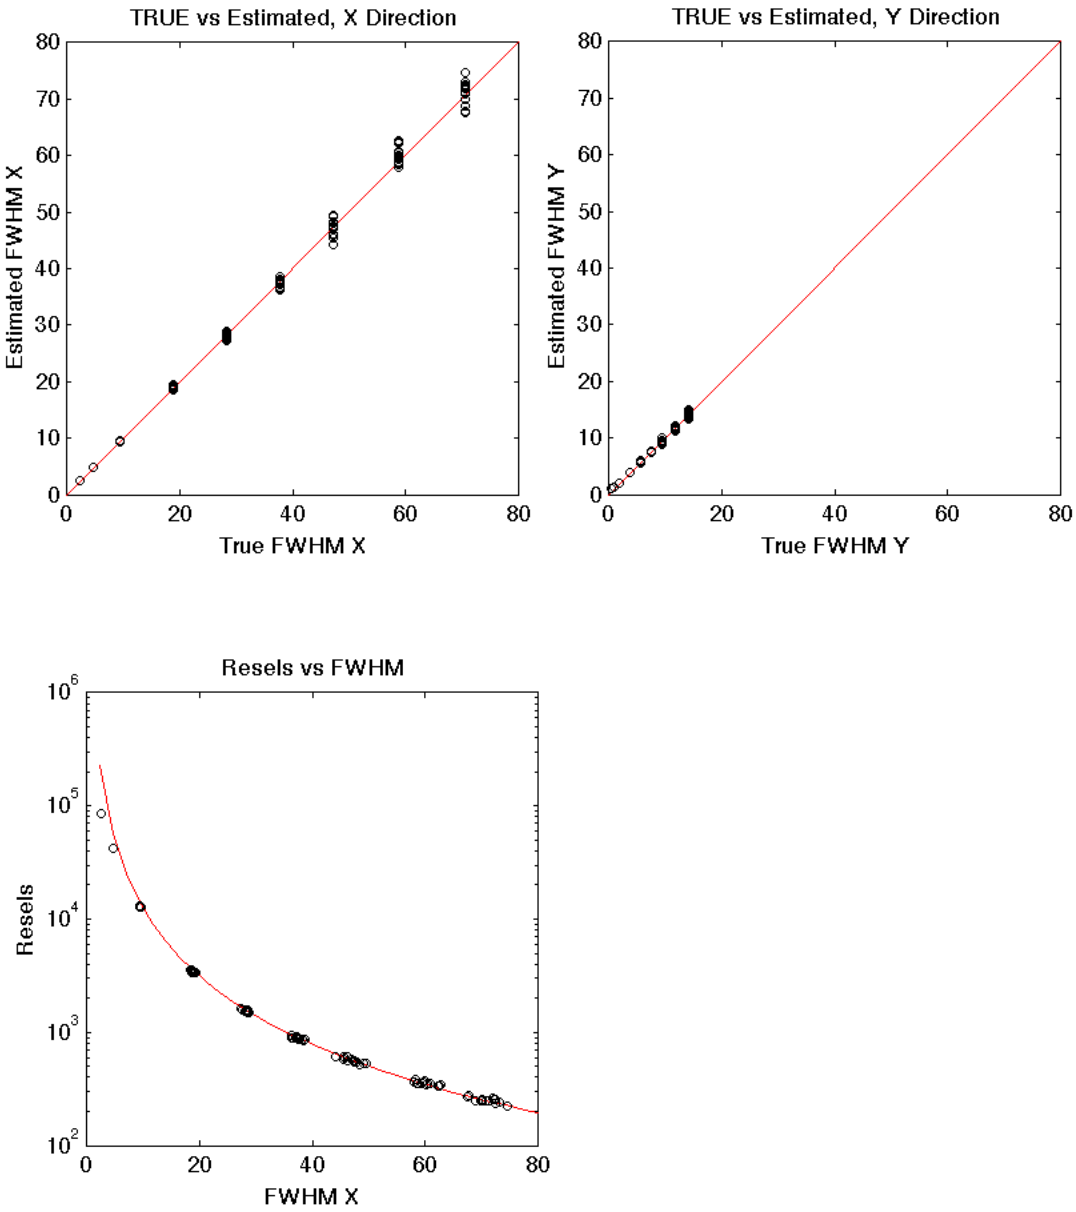

Figure 3. Predicted and estimated FWHM values for the X (left top image) and Y (right top image) directions for the anisotropic case. The bottom image shows the estimated resels compared to the true values. Each circle represents a unique result from a generated mG field.

Validation of the P(T_{max}>t) Estimation

The equation for calculation of the significance of the maximum t value ($P(t_{max} > t)$ given above (Worsley, et al., 1992) provides an estimate of $P(T_{max} > t)$. This estimator was derived to apply to fields with spatial correlation. The accuracy of this estimator is checked by creating Gaussian fields with a given level of spatial correlation and determining the proportion of these fields where the observed T_{max} exceeds the threshold value t. The proportions of fields for which the maximum exceeds the threshold are compared directly to the calculations from the equation in **Figure 4**. For each set of simulations, the proportion of fields where $T_{min} < t$ is also calculated and because the T distribution is symmetric, the $T_{min} < t$ values can be used as another observation at each threshold value. For each image in **Figure 4**, 1000 Gaussian fields were simulated to calculate the observed exceedance proportions. **Figure 4** shows that the accuracy of the $P(T_{max} > t)$ estimates improves as the FWHM increases. The calculations are also most accurate for $P(T_{max} > t)$ values ≤ 0.20, which is expected as this calculation is most often applied to higher thresholds where the chances that the maximum T value exceeds the threshold are smaller.

Another solution to estimating the probability of a value exceeding the threshold is tested. Following the discussion by Brett et al. (2003), consider a single t-test with 40 degrees of freedom that produces a t-statistic of 2.42. Comparison to the t distribution shows a 1% probability of a value this high or higher occurring by chance from background (null hypothesis) conditions. In an omnibus test applied to image analysis, the question is not the probability of whether one location (pixel) is from the background distribution, but the probability that the entire image could have come from the background distribution. The acceptable risk of assuming the t-statistics in the image came from the background distribution when in fact they did not is called the family-wise (or group-wise) error (FWE). Control of the FWE is set by identifying a threshold that takes into account the number of tests within the family. For example, if there are n = 1000 individual t-tests each having 40 degrees of freedom, then 10 of those tests should result in t-statistics of at least 2.42. The goal is to redefine the threshold such that there is a 1% chance of at least one t-statistic being at or above the threshold and the Bonferroni Correction has been developed to calculate this threshold.

For small α, the approximation $P^{FWE} <= n\alpha$ is valid□□□□□□a threshold for the single pixel probability, α, can be found that produces the desired P^{FWE}: $\alpha = P^{FWE}/n$. The value of P^{FWE} is the probability of any t value exceeding the threshold across all n pixels.

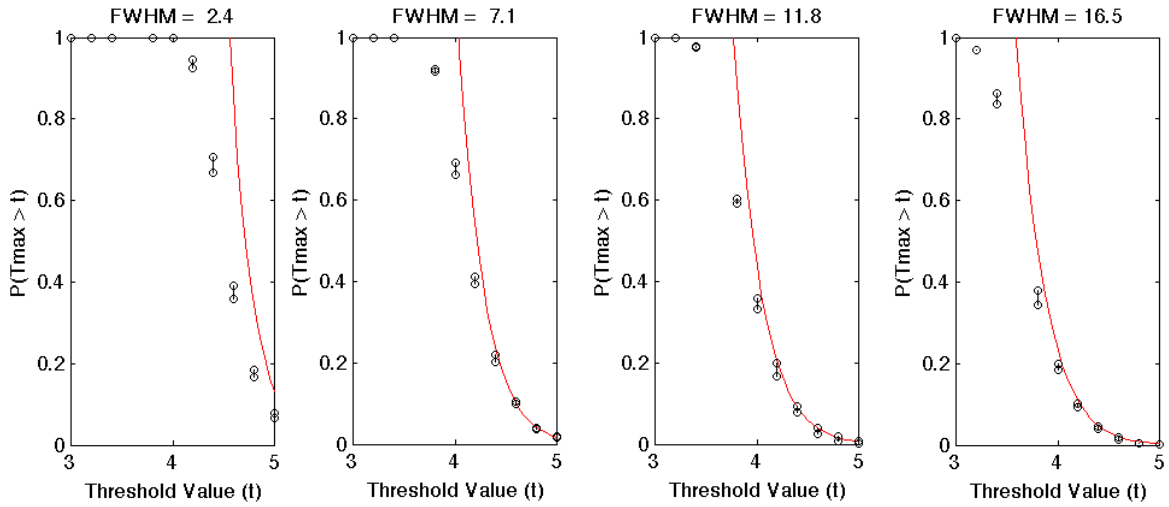

Figure 4. Validation results for calculation of $P(T_{max}>t)$. Each graph shows results for a different FWHM value as shown in the titles. The theoretical result is shown by the red line and the empirical results (black circles) are the proportion of times that $Tmax > t$ over 1000 fields. There is a separate black circle for $T_{max} > t$ and $T_{min} < t$ for each value of t.

This correction is demonstrated here with an example. For a regression model with 40 time steps (degrees of freedom = 39) and a threshold value, t, of 4.0, the cumulative probability of the T distribution is 0.9999 and the probability of exceeding the threshold is 1.37E-04. This is the α value of the test. For a relatively small image with 2500 pixels, there will be 2500 t-tests and in order to get a FWE equivalent to a single t-value exceeding 4.0, the modified α value for each individual test is now α/n = 1.374E-04/2500 = 5.48E-08. Going back to the null distribution, this corresponds to an adjusted single test threshold value of 6.45.

The Bonferroni correction assumes independence between the statistical tests and this is some of the reason why the resulting values are overly conservative for t-tests. The coordinate pair from the example calculation (6.45, 5.48E-08) would not even show up in the images of **Figure 4** that contain the observed $P(T_{max} > t)$ for t values expected to be used in this work. If the size of the image increases to more practical levels (e.g., 10,000 or more pixels), the adjusted α value decreases further and the corresponding threshold increases. Based on these calculations, the Bonferroni correction to $P(T_{max} > t)$ is not considered further.

Example Applications

Two data sets are used as example applications for the techniques described above, collectively referred to here as "SPM". Each example data set represents a spatial-temporal process. The anomaly detection focus here is on the maximum and minimum SPM values, either t-statistics or z-score values, and the statistical significance of these values and on the size of the maximum regions of excursion above or below a threshold level. For the second application to the NDVI data, the observed levels of spatial correlation, the FWHM, are also examined.

Data Set 1: Simulated Data

The first example applications are demonstrated using a synthetic data set. Such a data set allows for complete knowledge of what is background variation and what are the true anomalies of interest.

The synthetic data set is constructed as the linear combination of two independent, second-order stationary mG fields using the deformation approach developed by Hu (2000). The combination is done as a function of a non-dimensional time, *v*, that allows for the temporal evolution of the combined field to be dependent on previous states of that field.

For each time step, *Y(v)* is a Gaussian random function defined in 2D.

$$Y(v) = Y_1 cos(v) + Y_2 \sin(v)$$

If Y_1 and Y_2 are zero-mean Gaussian functions, then *Y(v)* is also a zero mean function and the spatial covariance model of *Y(v)* is that of Y_1 and Y_2. The *cos* and *sin* operators act as weights to normalize the variance and the resulting *Y(v)* is also a Gaussian function due to its construction as the linear combination of two other Gaussian functions.

Moving from the random functions to specific realizations of these functions, the relationship above provides a chain of realizations (time slices) that vary smoothly as function of *v*:

$$y(v) = y_1 cos(v) + y_2 \sin(v)$$

and the smaller the time step, Δv, the smoother the transition from one time slice to the next will be. A full cycle of *v* from 0 to 2π creates all possible fields. In order to continue to simulate unique data, some level of measurement noise as well as a trend, here linear, can be applied to the simulated values.

As an example, two Gaussian fields are constructed by convolving two uncorrelated mG fields with an isotropic Gaussian kernel having a FWHM value of 10.0 (σ = 4.25). These two fields are the initial conditions for creation of time varying field using the Gaussian deformation approach. The time step is set such that there are 62.667 time steps per cycle ($\approx 10*2\pi$). The simulation is run for three complete cycles (188 time steps) and both noise and a decreasing linear trend are added to the simulated data. The spatial domain is

discretized into 100x100 pixels, and for each pixel and for each time step, a random Gaussian variable is drawn and added to the simulated data. This Gaussian noise distribution has a mean of zero and a standard deviation of 0.10. The linear trend is also added to the simulated value to decrease them by 0.01 per time step. **Figure 5** shows an example resulting combined field for 9 time steps.

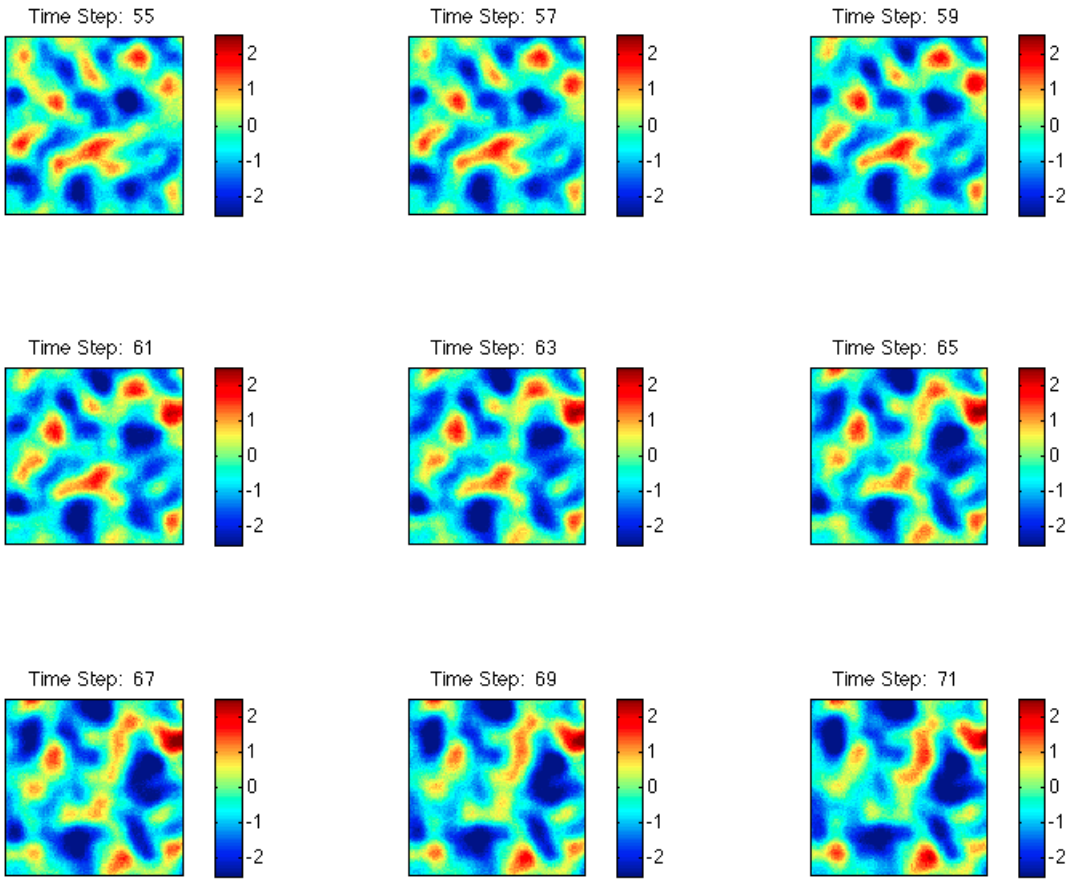

Figure 5. Example time slices from the simulated data set.

The online anomaly detection approach is applied to simulated data sets created using the Gaussian deformation algorithm. The simulations use the deformation approach described above in a domain that is 100x100 and has an isotropic FWHM of 10.0 (correlation length is $1/10^{th}$ of the domain size). The simulated values are drawn from a standard normal multiGaussian distribution. The deformation approach is periodic and from a given pair of starting realizations, any simulated values separated by one period in time will be exactly the same. To eliminate the repeating values, both noise and a linear trend are added to the simulated values. A decreasing linear trend with a slope of -0.01 is added to the simulated values after which noise is added as a zero-mean, uncorrelated Gaussian process with a standard deviation of 0.10. The Gaussian deformation process is run to simulate 188 time steps where each time step is 0.10.

The regression model defined previously is used here with a window size (P) of 50 time steps and a fixed period of 1.0 (10 time steps). The regression model is used to predict the values of the spatial-temporal process for each time steps from time step 101 through 125. Maps of the t-statistic derived from the differences between the regression model estimate and the observed values at each time step are then analyzed for anomalies. It is noted that the SPM approach typically converts the t-maps into Gaussian z-score maps prior to analysis, but here the t-statistic maps are used directly.

A number of tests are completed in which an anomaly is placed into the data sets at time step 110. Different combinations of anomaly size, strength and shape are tested and the following performance measures are recorded through time: the number of individual pixels exceeding the threshold (N), the Euler characteristic (m, or the number of regions exceeding the positive value of the threshold) and the minimum and maximum t-statistic values. **Figure 6** shows example results for circular, square and kernel anomalies of an intensity of 5.0 and a size of 6.0. The intensity value is added directly to the background data for the circular and square anomalies and defines the maximum increase above background at the center of the kernel shaped anomaly. The size defines the radius, the side length or the standard deviation (σ) of the Gaussian kernel for the circular, square and kernel anomalies, respectively. These definitions lead to kernel anomalies being wider and having a lower average intensity relative to the circular and square anomalies. In all cases, the anomaly is added to the background data for a single time step 110.

Results in **Figure 6** (left column) show extreme values of the maximum and minimum t-statistics at early time steps. A number of these values exceed the Y-axis limits of the graphs (-5.0, 5.0), especially at the earlier time steps. The probability of the maximum, or minimum, observed t-statistic exceeding an absolute value of 5.0 for a 500x500 field with a FWHM of 10.0 is 0.008 (see **Figure 4**). The reason for the large numbers of values that exceed 5.0 is not clear and it is possibly related to the start up of the regression model. From the graphs below, the clearest evidence of the anomaly is seen for the kernel shape where the maximum t-statistic jumps off the graph at time step 110.

The right column in **Figure 6** shows the observed and estimated Euler characteristic. All three graphs present different results with strong changes in both EC measures at time step 110 for the circular and square anomalies, but little if any change for the kernel anomaly. The estimated EC is derived from the FWHM values calculated on the image at each time step. The results below indicate that the FWHM values on the image must change significantly in the presence of the anomalies added here.

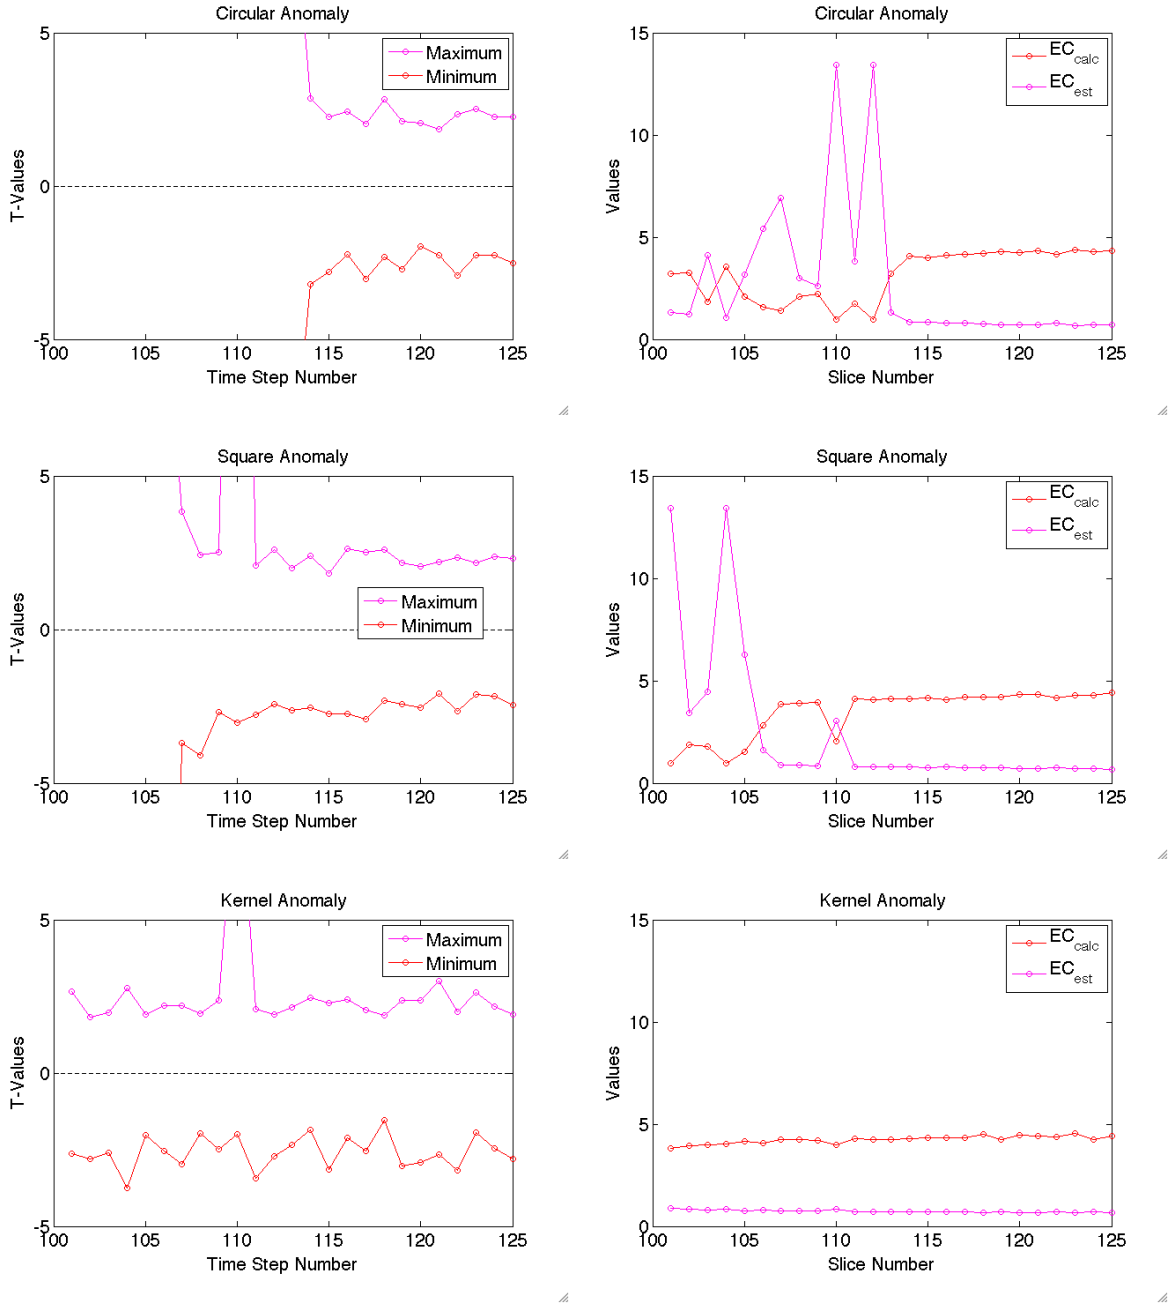

Figure 6. Example anomaly detection measures for simulated data. The left column shows the minimum and maximum T values at each time step and the right column shows the observed Euler characteristic ("EC_{calc}") and the Euler characteristic estimated from the spatial structure of the field and the threshold value ("EC_{est}"). All results are for a threshold of 3.5 an anomaly size of 6.0 and an anomaly intensity of 5.0.

Increasing the intensity of the anomaly results in increasing numbers of pixels exceeding the threshold *t* in the difference maps. Results for *t* =2.50 and a kernel anomaly (**Figure 7**) show that an increase in the anomaly intensity from 1.0 to 10.0 results in approximately a 5X increase in the number of pixels exceeding the threshold. For all intensities examined, the number of pixels exceeds the expected value of pixels ("Nexceed Est"). The variation in the trend is due to each point on the graph being the result of a single simulation.

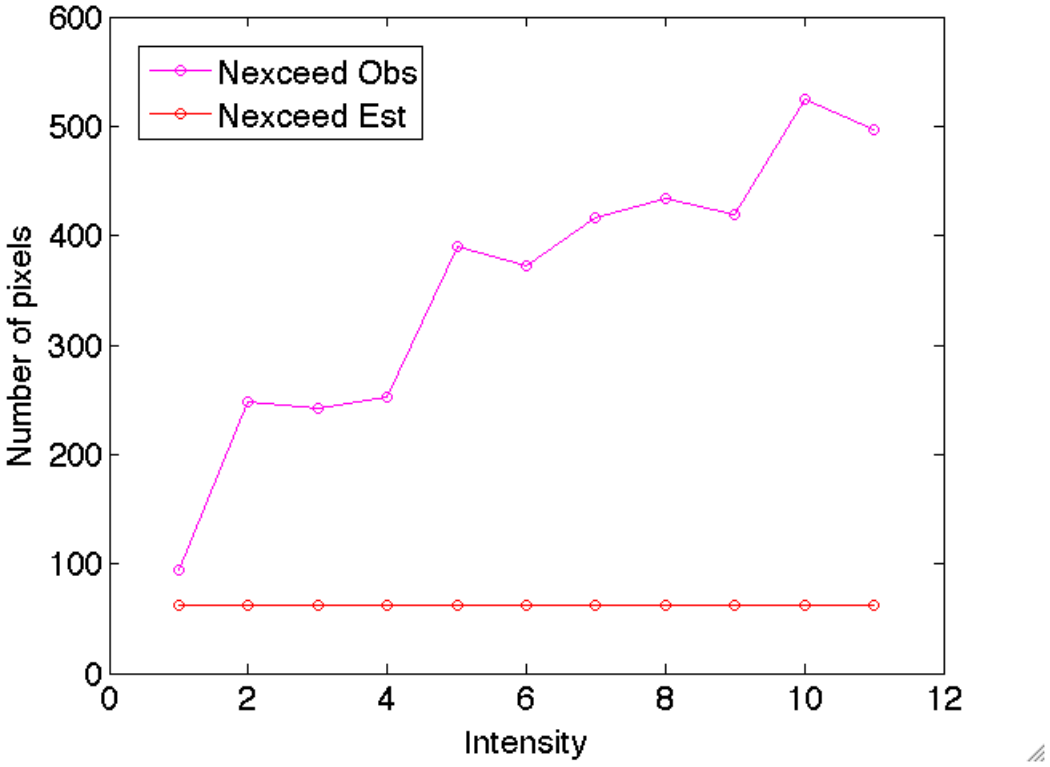

Figure 7. Number of pixels exceeding the threshold *t* as a function of increasing anomaly intensity for a kernel anomaly. For these simulations, the threshold, *t*, is set to 2.50.

Adding discrete anomalies to simulated spatial-temporal data allows for testing of the algorithms used here. However, additional testing is needed with multiple runs of the simulated background data and the added anomalies before any strong conclusions can be made.

Data Set 2: Normalized Difference Vegetation Index (NDVI) Data Set

The NDVI data set provides total global coverage of the earth on a weekly sampling frequency for the past 30+ years. The Wikipedia site: http://en.wikipedia.org/wiki/Normalized_Difference_Vegetation_Index provides a high-level overview of NDVI and its application. Data obtained for this project run from 1981 into early 2011, although much of the 1994 data are missing. For that reason, the analyses presented here are for the 1995-2011 data.

NDVI is calculated from multi-spectral satellite measurements and provides a normalized measure of the amount of green vegetation for a region on the earth's surface. NDVI is calculated as the difference between the near infrared (NIR) and visible regions (VIS) normalized by the sum of these two quantities:

$$NDVI = \frac{(NIR - VIS)}{(NIR + VIS)}$$

Calculated values of NDVI range from -1.0 to 1.0. NDVI is based on the ability of chlorophyll in plant leaves to strongly absorb visible light (wavelengths between 0.4 and 0.7 μm) to drive photosynthesis. Contrary to light in the visible range, plant leaves are strong reflectors of light in the near infrared range (0.7 to 1.1 μm). In general, increasing amounts of green vegetation in an image pixel results in increasing amounts of reflected light in the NIR wavelengths and decreasing amounts of reflected light in the VIS wavelengths leading to increasing NDVI values.

Typical values of NDVI in areas of dense vegetation are 0.3 to 0.8. Regions of cloud cover or snowfields, with little to no absorption of visible light will result in negative NDVI values. Water bodies tend to have low reflectance in both NIR and VIS wavelengths resulting in NDVI values near zero. Bare soils have smaller positive NDVI values up to 0.2. Direct relationships between NDVI and photosynthetic capacity and energy absorption of plant communities have been developed (Sellers, 1985; Myneni et al., 1995). An example image of NDVI for the 40th week of 1995 for North America is shown in **Figure 8**. Week 40 is nearly 77% of the way through the year, early October, and the harvest has already occurred in the Great Plains resulting in relatively low NDVI values. NDVI values are extremely low in the northern latitudes, upper one-third of the image, as would be expected for this time of year. NDVI values are still relatively high in the Appalachians and in the southern US as well as Caribbean islands and in Central America. The image shown below is 424 x 835 pixels and is one of 835 weeks of available imagery in the NDVI data set.

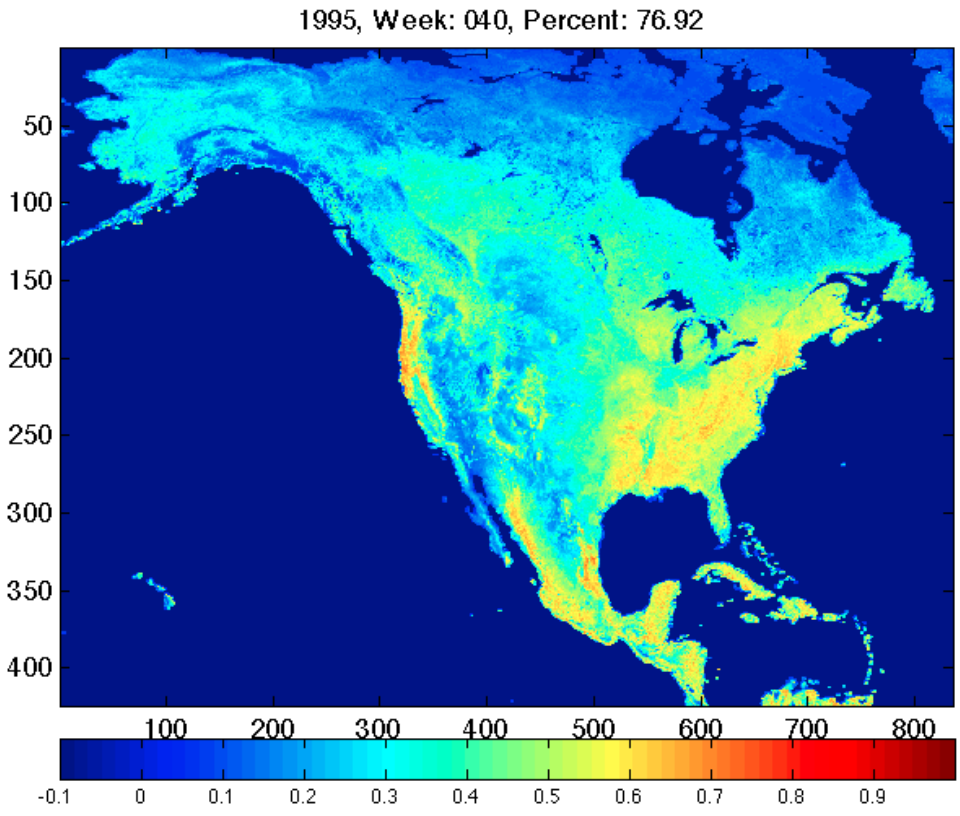

Figure 8. Example NDVI imagery for North America. This image is for early October in 1995.

Online Detection (Regression)

The NDVI data are also analyzed using the online anomaly detection mode. The regression model is paramterized with a window-size (P) of 40 weeks and a period of 52 weeks. In contrast to analysis of the simulated data sets, here the t-statistics are transformed to mG z-score fields prior to the anomaly detection. Ten years (520 weeks) of weekly data are available for analysis. Two locations within the tropics are chosen for analysis: Amazon Basin and Southern India. Details on these two regions are provided in Table 1.

Error! Not a valid bookmark self-reference.

	Amazon Basin	**Southern India**
Latitude Limits	-72 to -58 degrees	-80 to -74 degrees
Longitude Limits	-10 to 0 degrees	-21 to -14 degrees
Dimensions in pixels (rows and columns) and total	69 x 97 (6693)	49 x 42 (2058)
Approximate Area (km^2)	1.5E+06	4.6E+05

Figure 9 shows three consecutive example weeks from the end of 1999 to 2000 of NDVI data (left column) collected for the Amazon Basin with the corresponding SPM in units of z-scores as calculated through the regression-based approach (right column). A similar view of consecutive weeks of NDVI data collected in late 2000 is shown for the Southern India region in **Figure 10**.

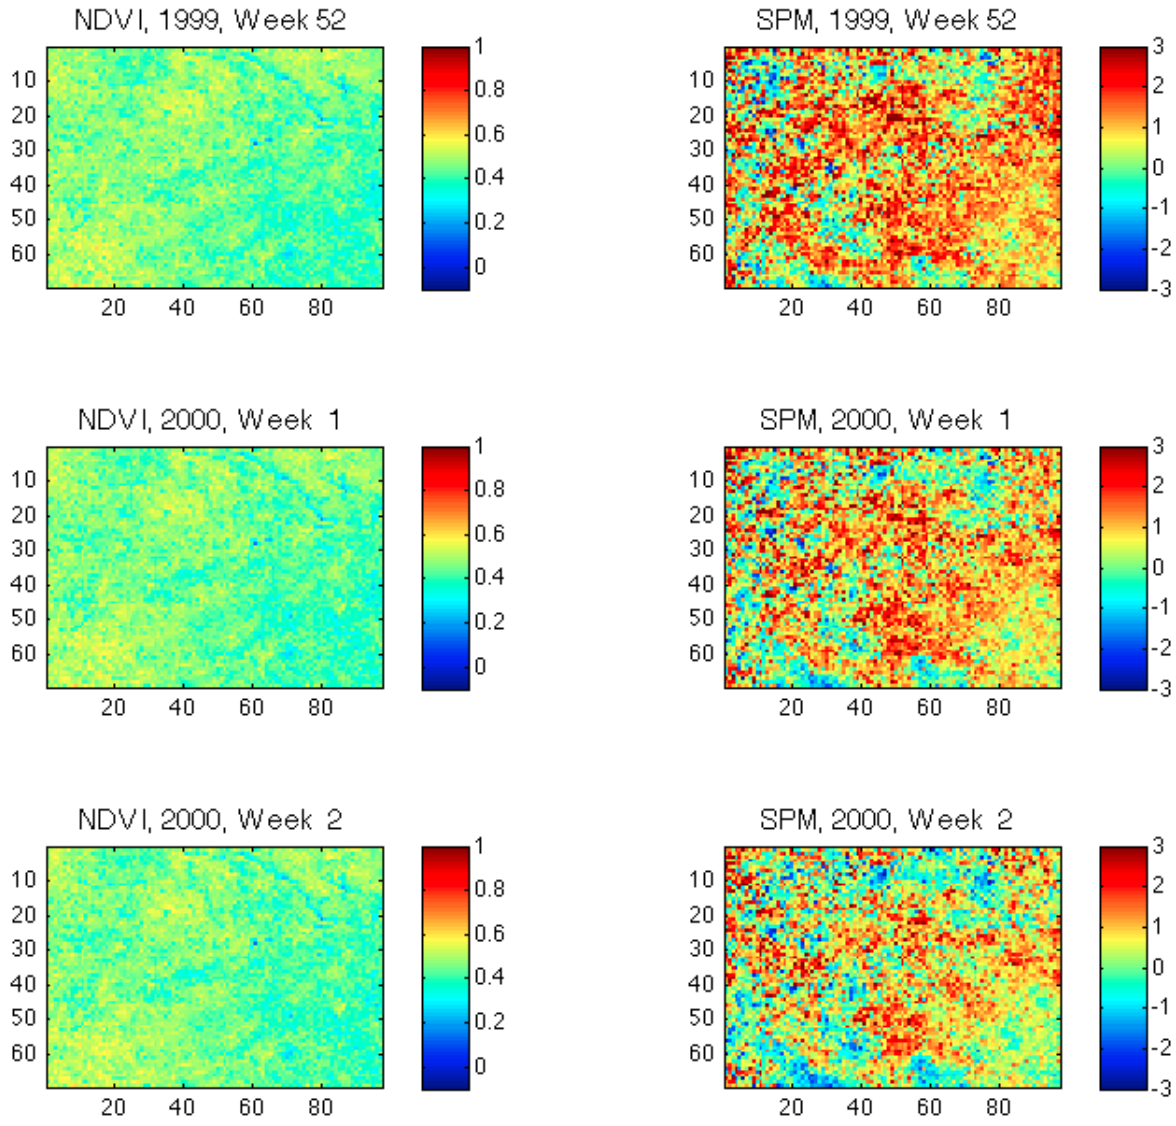

Figure 9. Example images from the Amazon Basin region showing three weeks of NDVI data (left column) and the corresponding SPM maps (right column).

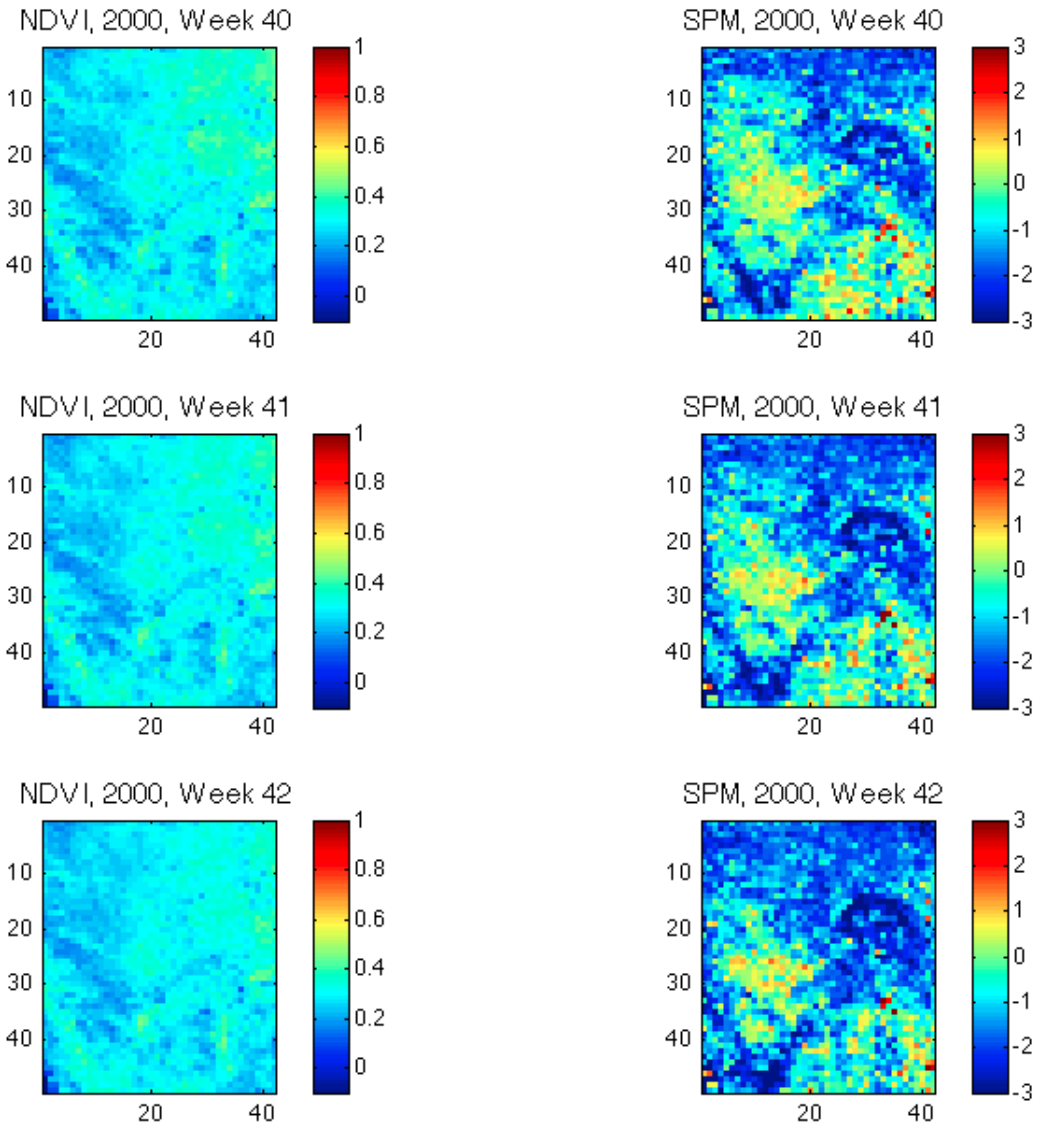

Figure 10. Example images from the Southern India region showing three weeks of NDVI data (left column) and the corresponding SPM maps (right column).

Figure 11 shows a comparison of the calculated FWHM values and indicates that the SPM's are isotropic at both locations (the values for the X and Y directions are approximately the same) and that the FWHM values are somewhat smaller for the Amazon Basin (**Figure 11**, left) than those for Southern India (right).

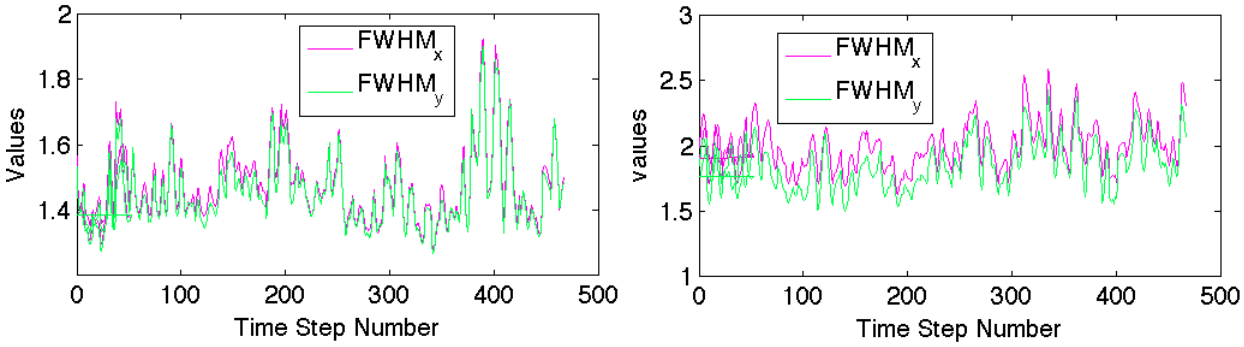

Figure 11. Comparison of the calculated FWHM values at each time step for the Amazon Basin (left) and South India (right).

Figure 12 compares the maximum and minimum observed SPM values, here z-score values, for the Amazon Basin (left) and Southern India (right). Using the area of each region in pixels and the approximate average FWHM values from the previous figure, the values of *P(Tmax>t)* for t = 5.0 are calculated as 0.0098 and 0.0017 for the Amazon Basin and Southern India, respectively. The Amazon Basin results show only a few excursions of the SPM values outside of the -5.0, 5.0 bounds while the Southern India results are much more variable and show many excursions outside these bounds.

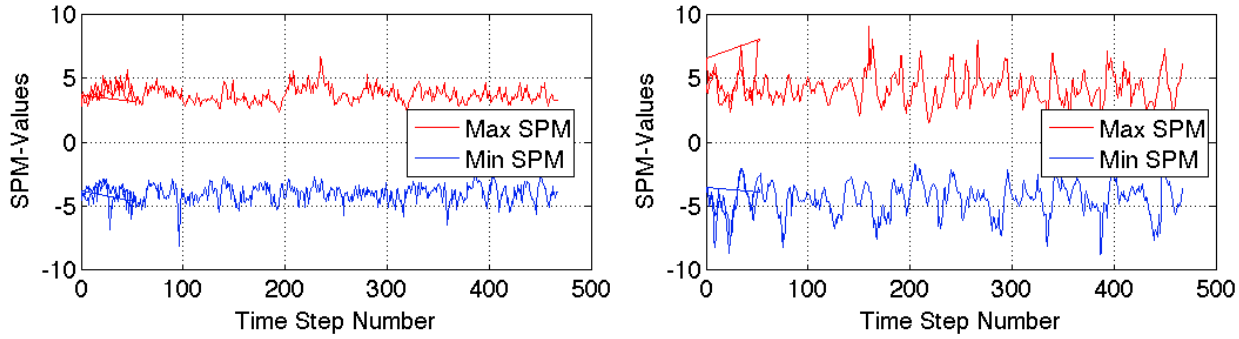

Figure 12. Comparison of the maximum and minimum observed SPM (z-score) values at each time step for the Amazon Basin (left) and South India (right).

Figure 13 shows the maximum size of the region that exceeds a threshold t = 4.0 for the Amazon Basin (left) and Southern India (right). Large areas of exceedence are indicative of anomalous NDVI conditions and the locations of these regions can be determined. The Amazon Basin has a maximum excursion region of 3 pixels, while Southern India has two time steps where the maximum excursion exceeds 30 pixels. Both of these are negative excursions indicating that the regression model underestimated the amount of NDVI at these times.

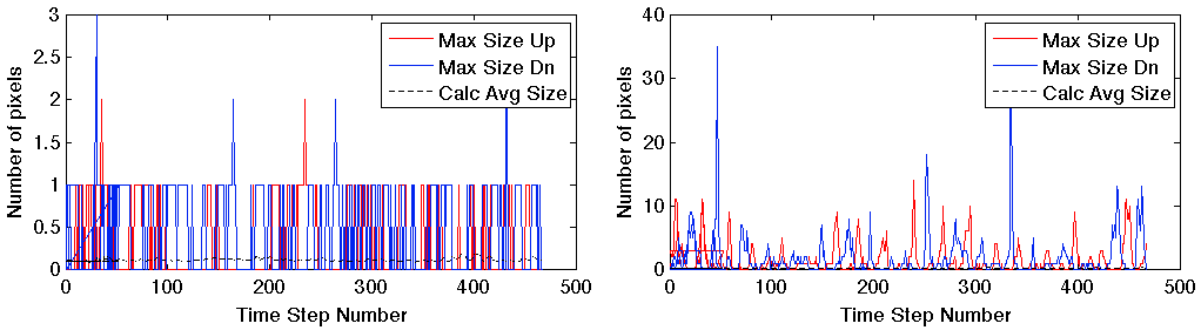

Figure 13. Sizes in pixels of the maximum excursions above and below a threshold of 4.0 (-4.0) for the Amazon Basin (left) and Southern India (right).

Conditional Differences

In addition to the online, regression-based anomaly detection, the NDVI data can also be analyzed in a manner that is more similar to the brain scan studies for which the SPM approaches have been developed. In particular, external stimuli that should produce a rapid response in the subject, the earth, as measured with NDVI are available. Here we consider the el Nino Southern Oscillation (ENSO), more commonly referred to as *el Nino* and *la Nina* as the external conditions that influence the expression of NDVI. Additionally, a third condition, the "background" conditions in the absence of either *el Nino* or *la Nina* conditions can be examined. **Figure 14** shows the occurrence of these conditions over the extent of the NDVI data set.

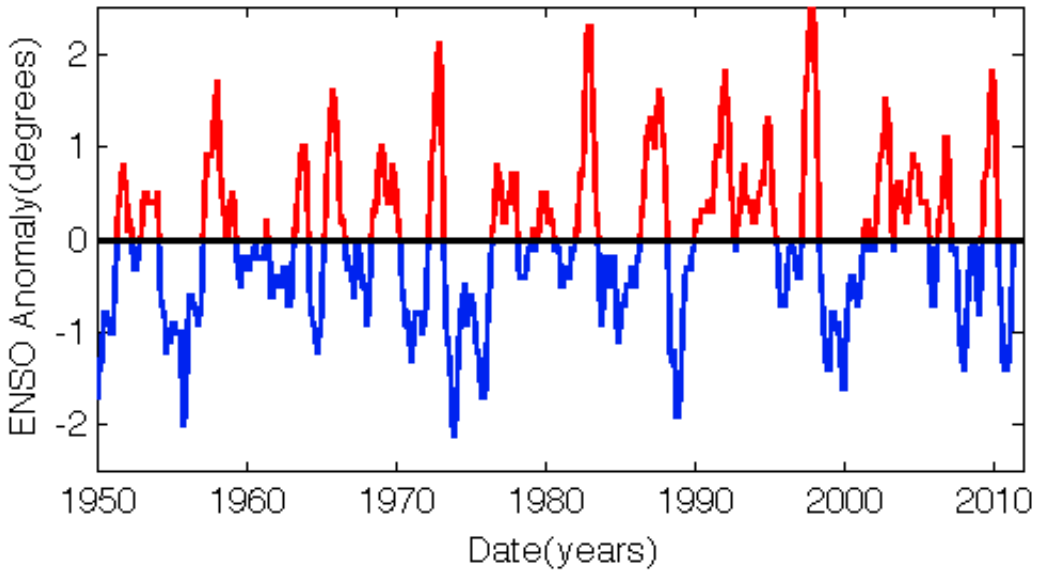

Figure 14. ENSO anomaly in degrees C for the 1950-2011 period. Warmer periods shown in red correspond to *el Nino* events and cooler temperatures in blue define *la Nina* events.

Table 1 shows the years selected as representative of each condition within the 1995-2011 dataset (compare with **Figure 14**). Each year listed in **Table 1** is the start of the "water year" which extends from October 1st through September 30th of the following year.

Table 1. Classification of years since 1995 for use in constructing average conditional data sets.

El Nino Years	La Nina Years	Background Years
1997	1999	1996
2002	2000	2001
2006	2007	2005
2009	2010	2008

The t-statistic is calculated as: (activated state – normal state), and here that calculation is: (NDVI during *el Nino*) – (NDVI during *la Nina*). Therefore, positive t values, and the positive corresponding z-score values, are the case where the activated state (*el Nino*) results in higher NDVI (higher vegetation density) relative to the normal state (*la Nina*). Negative values are due to less vegetation density in the activated condition relative to the background condition.

The FWHM values for this analysis are shown in **Figure 15**. These are similar to the values calculated above in the regression-based anomaly detection with the Amazon Basin (**Figure 15**, left) being near or below 1.5 and the values for Southern India (**Figure 15**, right) generally at or below 2.0.

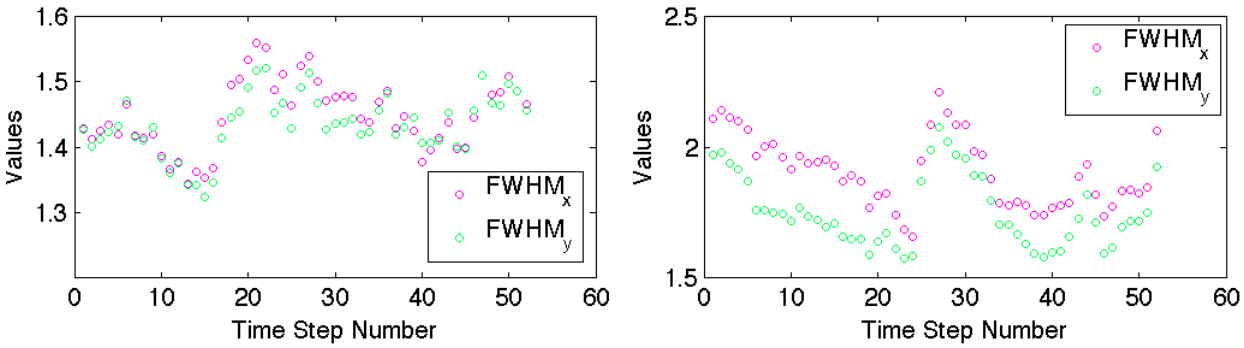

Figure 15. Comparison of the calculated FWHM values at each time step for the Amazon Basin (left) and South India (right). The time steps are for the calendar year and the results are average values for *el Nino* years compared to *la Nina* years.

The maximum and minimum z-score statistics across the calendar year are shown in **Figure 16** for the Amazon Basin (left) and Southern India (right). There are few, if any, points where the maximum t-value exceeds 5.0. However, for both regions, the minimum z-score values exceed -5.0 frequently. This result indicates that NDVI during *la Nina* is higher than NDVI during *el Nino*. The values of $P(Tmax>t)$ for t = 5.0 (-5.0) are the same as above in the regression-based anomaly detection and calculated as 0.0098 and 0.0017 for the Amazon Basin and Southern India, respectively.

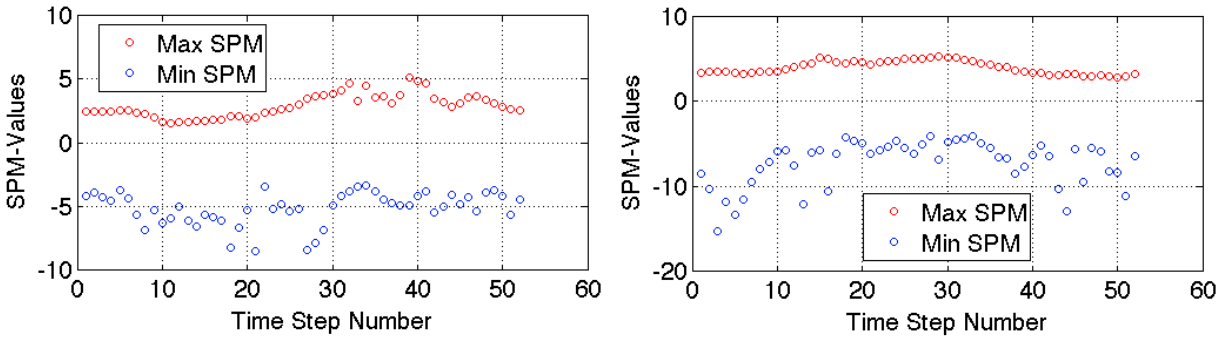

Figure 16. Maximum and minimum SPM values (z-scores) for the Amazon Basin (left) and Southern India (right) under conditional differences.

A somewhat unexpected result in both analyses was the minimal correlation lengths (FWHM values) of the resulting z-score maps. Both regions examined here, as well as several others, produced z-score maps with mean FWHM values in the range of 1.5 to 3.0 pixel lengths. This result is indicative of the regression model functioning as designed and producing relatively accurate predictions of the NDVI values at the next time step. However, the resulting small FWHM values are not optimal for the SPM approach. Recall the validation of the $P(T_{max} > t)$ calculations where smaller FWHM values led to increasing bias in the analytical estimation of $P(T_{max} > t)$.

One interesting result occurs when using *el Nino* as the activated condition and *la Nina* as the "normal" condition. The majority of the significant deviations are negative – meaning those areas have lower NDVI (less vegetation) in *el Nino* years relative to *la Nina* years. This result may be consistent with underlying factors contributing to increased incidence of civil conflict seen with increasing level of temperature in the ENSO cycle (Hsiang et al., 2011). At this point, no direct connection between decreased NDVI during *el Nino* relative to *la Nina* and decreasing levels of agricultural vegetation, has been identified, but NDVI may turn out to be one of several explanatory factors for civil conflict.

Summary
This work summarizes previous techniques developed as statistical parametric mapping, with applications focused on human brain imaging, and interprets those techniques for new applications in anomaly detection within remotely-sensed imagery. Two approaches to analysis are developed: online, regression-based anomaly detection and conditional differences. These approaches are applied to two example spatial-temporal data sets: data simulated with a Gaussian field deformation approach and weekly NDVI images derived from global satellite coverage. Results indicate that anomalies can be identified in spatial temporal data with the regression-based approach. Additionally, compared to *la Nina* years, *el Nino* conditions lead to significant decreases in NDVI in both the Amazon Basin and in Southern India.

Acknowledgements
The authors are grateful to Vineet Yadav for providing the NDVI data set used in this study.

References

Adler, R.J. and A.M. Hasofer, 1976, Level Crossings for Random Fields, *Annals of Probability*, 4, pp. 1-12.

Adler, R.J., 1981, *The Geometry of Random Fields*, Wiley, New York,

Adler, R.J. and J.E. Taylor, 2007, *Random Fields and Geometry*. Springer

Adler, R.J. and J.E. Taylor and K.J. Worsley, 2009, *Applications of Random Fields and Geometry: Foundations and Case Studies,* available at: http://webee.technion.ac.il/people/adler/publications.html accessed, January 2009.

Brett, M., W. Penny and S. Kiebel, 2003, An Introduction to Random Field Theory, Chapter 14, in: *Human Brain Function, 2nd edition*, (R.S.J. Frackowiak, K.J. Friston, C. Frith, R. Dolan, K.J. Friston, C.J. Price, S. Zeki, J. Ashburner, and W.D. Penny, editors), Academic Press.

Fox, P.T., M.A. Mintun, E.M. Reiman and M.E. Raichle, 1988, Enhanced detection of focal brain responses using intersubject averaging and distribution analysis of subtracted PET images, *Journal of Cerebral Blood Flow and Metabolism*, 8, pp. 642-653.

Friston, K.J., K.J. Worsley, R.S.J. Frackowiak,J.C. Mazziotta, and A.C. Evans, 1994, Assessing the Significance of Focal Activations Using Their Spatial Extent, *Human Brain Mapping*, 1, pp. 210-220.

Friston, K.J, A.P. Holmes, K.J. Worsley, J.-P. Poline, C.D. Firth and R.S.J. Frackowiak, 1995, Statistical Parametric Maps in Functional Imaging: A General Linear Approach, *Human Brain Mapping*, 2, pp. 189-210.

Gilbert, R.O., 1987, *Statistical Methods for Environmental Pollution Monitoring*, Van Nostrand Reinhold, New York, 320 pp.

Hu, L.Y., 2000, Gradual Deformation and Iterative Calibration of Gaussian-Related Stochastic Models, *Mathematical Geology*, 32 (1), pp. 87-108.

Hsiang, S. M., Meng, K. C., Cane, M. A., 2011, Civil conflicts are associated with the global climate, *Nature*, 476, pp. 438-441.

Lantuejoul, C., 2002, *Geostatistical Simulation: Models and Algorithms*, Springer, Berlin, 256 pp.

Matlab, 2009, Image Processing Toolbox. Mathworks, Inc., Natick, MA, USA.

McKenna, S.A., J. Ray, B. van Bloemen Waanders, and Y. Marzouk, 2011, Truncated MultiGaussian Fields and Effective Conductance of Binary Media, *Advances in Water Reeources*, 34, pp: 617-626.

Myneni, R. B., F. G. Hall, P.J. Sellers, and A.L. Marshak, 1995, The interpretation of spectral vegetation indexes, *IEEE Transactions on Geoscience and Remote Sensing*, 33, pp. 481-486.

Normalized Difference Vegetation Index
http://en.wikipedia.org/wiki/Normalized_Difference_Vegetation_Index
Accessed August, 2011.

Sellers, P. J., 1985, Canopy reflectance, photosynthesis, and transpiration, *International Journal of Remote Sensing*, 6, pp. 1335-1372.

Worsley, K.J., A.C. Evans, S. Marrett and P. Neelin, 1992, A Three-Dimensional Statistical Analysis for rCBF Activation Studies in Human Brain, *Journal of Cerebral Blood Flow and Metabolism*, 12, pp. 900-918.

Worsley, K.J., S. Marrett, P. Neelin, A.C. Vandal, K.J.Friston, and A.C. Evans, 1996, A Unified Statistical Approach for Determining Significant Signals in Images of Cerebral Activation, *Human Brain Mapping*, 4, pp. 58-73.

Distribution:

1 Technical Library, 08944 (electronic)