Reprinted from

Eleventh International Symposium

Machine Processing of

Remotely Sensed Data

with special emphasis on

Quantifying Global Process:
Models, Sensor Systems, and Analytical Methods

June 25 - 27, 1985

Proceedings

Purdue University
The Laboratory for Applications of Remote Sensing
West Lafayette, Indiana 47907 USA

MULTISPECTRAL CHANGE DETECTION USING DEFFERENCE CLASSIFICATION AND BITEMPORAL CLASSIFICATION

GEORGE M. HALEY

Santa Barbara Research Center
A Subsidiary of Hughes Aircraft Company
Santa Barbara, California

ABSTRACT

The problem of detecting changes in boundaries while ignoring simple spectral changes is introduced, and a general approach to this problem is examined in detail. Two methods of boundary change detection are explained and compared in terms of accuracy and computational requirements.

I. OBJECTIVE AND APPLICATIONS

The specific objective of this study is to develop a method of detecting the creation or elimination of borders in a scene from one time to another while ignoring simple spectral changes (those not involving borders) which may occur. Examples of valid changes in this context are:

- an object which is present at one time and not at the next
- an object which has grown or changed in shape
- an object which has moved.

Examples of changes which are to be ignored are:

- an object increasing or decreasing in brightness relative to its background
- an object changing in spectral characteristics relative to its background.

Figure 1 illustrates these situations and the proper outputs of a boundary change detection algorithm.

The uses for such a technique can be best illustrated by enumerating specific applications. These include:

- detection of new roads in agricultural areas where there is an abundance of uninformative spectral variation (caused by crop rotation, seasonal vegetation changes, etc.)
- detection of new buildings
- military surveillance
- observation of glacier movement or snow melt
- any other situation in which general change detection is used where simple spectral changes are not of interest.

II. GENERAL APPROACH

Before we proceed, a few contextual definitions are in order to prevent ambiguity. A *scene* is the actual area being sensed to produce an *image,* which is the data obtained from sensing. The image will have spatial and spectral characteristics analogous to those in the scene. An *object* is a subarea of a scene having uniquely distinguishing spectral (spectral signature) and spatial (size, shape, location) characteristics, while a *region* is a similarly defined subsection of the image. Therefore, every distinct object in the scene will be represented as a distinct region in the image.

A *border* is the one-dimensional space separating two contiguous objects in a scene, and a *boundary* is a one-dimensional space in an image separating two regions. Furthermore, to allow us to make stronger statements about the accuracy and precision of boundary locating techniques, we will assume (somewhat arbitrarily) that the separation of two objects occurs along a curve having the property that all of its points are local maxima of the magnitude of the spatial gradient of the scene in the direction **d** orthogonal to the line t tangent to the curve at that point $\mathbf{x}_b$. This is illustrated in the following diagram:

Scene $g(\mathbf{x}, \lambda)$
(where λ is wavelength)
having spatial gradient g_x

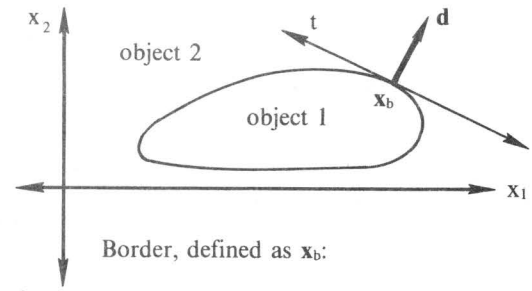

Border, defined as $\mathbf{x}_b$:

$$\left| \int g_x(\mathbf{x}_b, \lambda)\, d\lambda \right| \geq \left| \int g_x(\mathbf{x}_b - c\mathbf{d}, \lambda)\, d\lambda \right|$$
$$\forall \; -\delta \leq c \leq \delta, \text{ where } |\delta| > 0$$

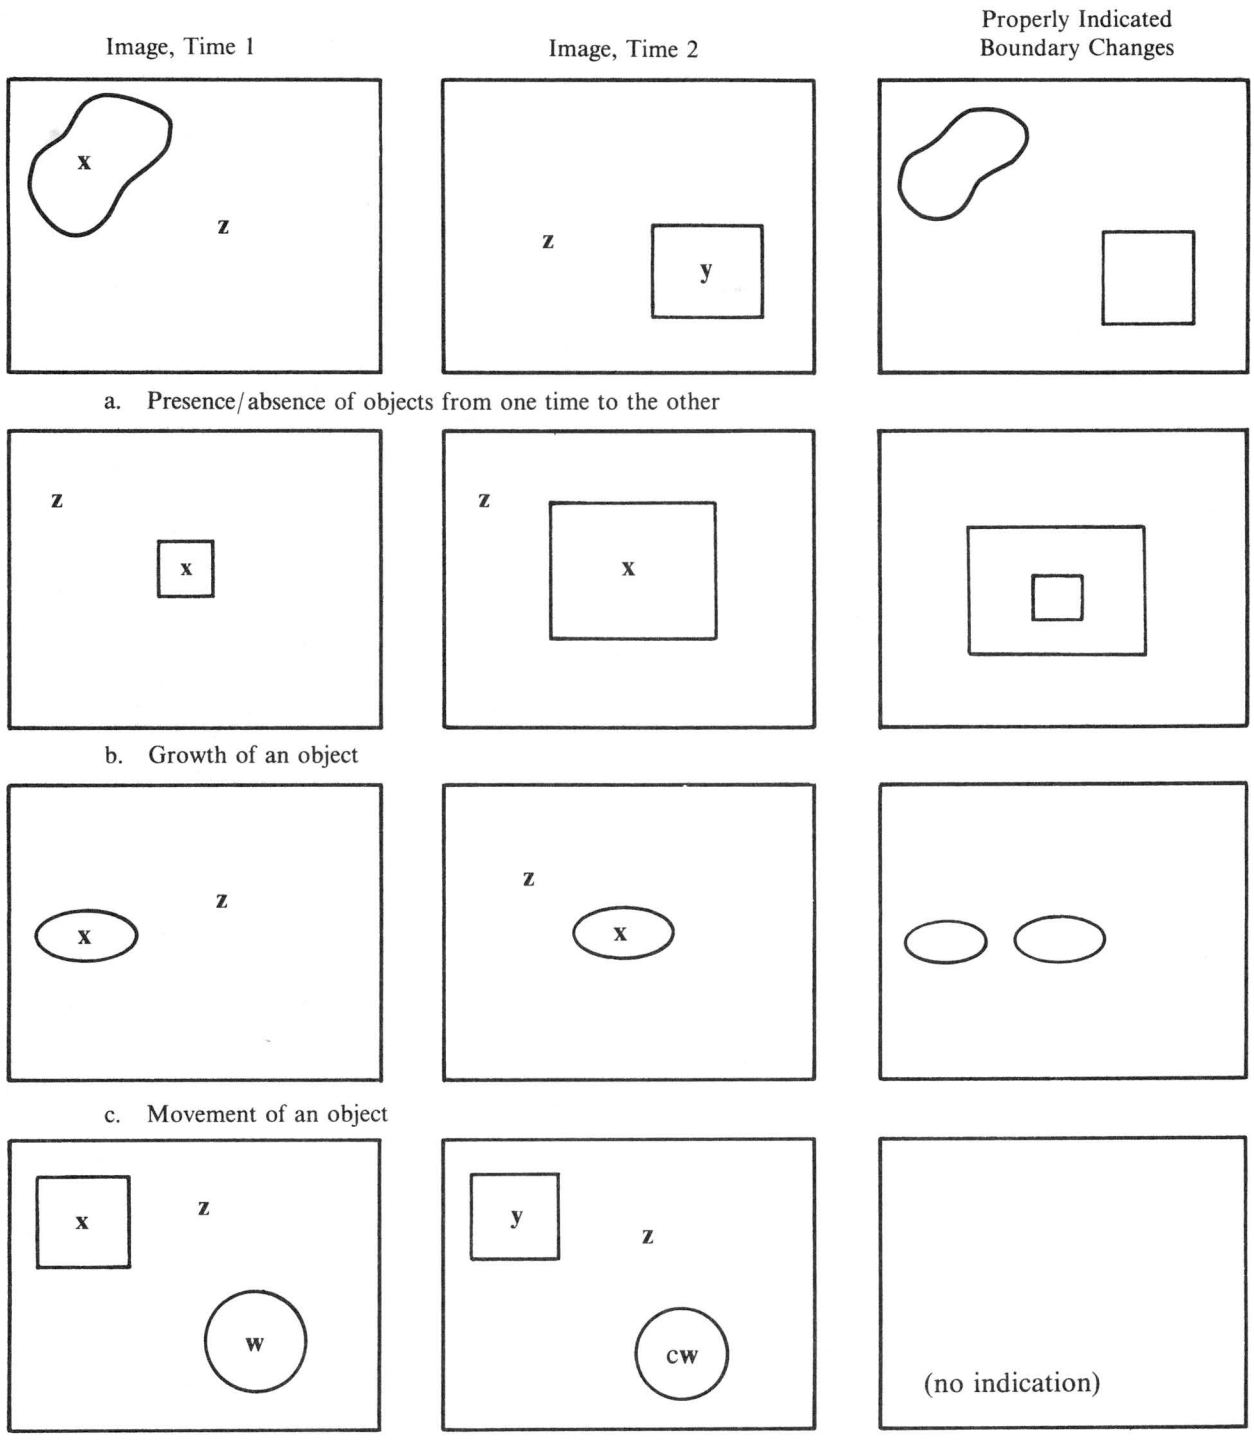

Figure 1. Examples of spectral changes and their proper boundary change indications.

While this leads to a reasonable border, it should be realized that other reasonable borders can be obtained using a different rule.

A *boundary differential* is a boundary with the added dimension of intensity, which is determined by the difference in or ratio of intensities between the regions it separates. An *edge* is an approximation to the spatial gradient having three dimensions—length, width, and intensity. In the classification procedures to be discussed, a *cluster* refers to a group of contiguous points having a single spectral vector, whereas a *class* is a set of spectral vectors having a single nominal vector typifying the spectral characteristics of the points and regions which "belong" to that class. Finally, the notation will be as follows: k and l are spatial coordinates in an image, n is the spectral band number, N is the number of spectral bands in an image, and each vector, matrix, and tensor will be denoted by a boldface letter with a subscript (e.g., i,j,1,2) to indicate the scene to which it corresponds. The significance of these terms and definitions should become clearer as the discussion progresses.

The general approach to boundary change detection amounts to indicating the boundaries, $\mathbf{B}_1$ and $\mathbf{B}_2$, in the images from both times, $\mathbf{I}_1$ and $\mathbf{I}_2$, and then performing the equivalent of an exclusive OR (XORing). These two steps will now be detailed.

A. INDICATION OF BOUNDARIES

There are two qualities important in a boundary: its accuracy in approximating the border of two objects, and its precision in doing so. Two methods of boundary indication and their relative accuracies and precisions will now be discussed.

1. Gradient Methods. By convolving an image with a function approximating a gradient and taking its magnitude, an edge is obtained. A single edge E(k,l) may be formed from a function of the edges in the N spectral bands:

$$E(k,l) = f\,[E(k,l,1),\ E(k,l,2),\ \ldots,\ E(k,l,N)],$$

where f might be the norm, the sum, or the maximum over N bands.

Now, assuming that the images are accurate in their representation of the scene, the values of E(k,l) will be relatively high not only on the border, but within a distance Δw of the border. This is principally due to three factors: the size of the convolution function used to approximate the gradient, the effect of the sensor MTF, and the fact that the transition between "objects" is not always distinct. Thus, while it may be said with certainty that the border will exist solely within areas where E(k,l) is relatively high and that the boundary should be located in such an area, it is difficult to actually calculate with precision where the boundary should be located to yield the highest accuracy. Instead, it is much simpler to determine a "general boundary" $\beta(k,l)$ which can be regarded as a two-dimensional approximation to B(k,l), by thresholding E(k,l) as follows:

$$\beta(k,l) = \begin{cases} \text{TRUE, if } E(k,l) \geq e; \\ \text{FALSE, if } E(k,l) < e, \end{cases}$$

where e is chosen to maximize the number of boundary points included in $\boldsymbol{\beta}$ (maximizing accuracy) while minimizing the number of "TRUE" values in $\boldsymbol{\beta}$ which do not correspond to boundary points (maximizing precision). In actual images, it is virtually impossible to choose a threshold such that $\beta(k,l)$ will correspond exactly, or even closely, to B(k,l) for all k,l. In practical situations, if high accuracy is desired, the precision to which a boundary will be specified using gradient methods will be relatively low, and if high precision is desired, it will not be achieved without gross inaccuracy.

2. Classification. The idea behind classification is to isolate regions and reduce them into clusters representing distinct objects. An ideal classifier will create exactly one cluster per distinct region in the image. The spectral value of each cluster should contain the essential distinguishing spectral characteristics of the associated region. The classified image **C**, at least in this context, will consist solely of clusters, each having a single spectral value determined by the spectral class to which the cluster belongs. Assuming that every object in the scene is adequately represented both spatially and spectrally in the classified image, it may be said that no essential information is lost in the classification process, or that the classifier is, loosely speaking, lossless. In this study, losslessness will be an important virtue of an ideal classifier.

To illustrate classification, consider a scene containing, among other objects, a field of corn with a gravel road through it. A good classifier will create at least two classes, one for corn and one for gravel, and using a rule it will decide which points belong to the class of corn and which to the class of gravel. The resulting classified image will contain one cluster representing the corn field and one cluster representing the road. The spectral values of these clusters should be typical of "corn" and of "gravel," respectively.

The choice of a boundary in a classified image is immediately obvious. The only reasonable assignment of boundaries is along the curves separating the clusters. In a digitized image, boundary pixels could be assigned as follows:

$$B(k,l) = \begin{cases} \text{TRUE, if both of the following conditions hold:} \\ \quad \text{—}C(k,l,n) \neq C(k+1,l,n),\ C(k-1,l,n), \\ \quad\ C(k,l+1,n) \text{ or } C(k,l-1,n) \\ \quad\ \text{for all n, and} \\ \quad \text{—the size of the cluster to which} \\ \quad\ C(k,l) \text{ belongs is larger than its} \\ \quad\ \text{neighboring cluster} \\ \text{FALSE, otherwise.} \end{cases}$$

It should be noted that the second condition of B(k,l) being true ensures single-pixel-wide connected boundaries, and that other conditions may be used instead. In the following classification techniques, the precision of the boundaries will be limited only to the size of the pixels, while the accuracy will be dependent on the particular classifier used.

The boundary differential **Y** can be computed at the boundary pixels as a function f of the differences between C(k,l,n) and a different neighbor C(u,v,n) in all spectral bands as such:

$$Y(k,l) = \begin{cases} f\{[C(k,l,1) - C(u,v,1)], [C(k,l,2) \\ \quad - C(u,v,2), \ldots, [C(k,l,N) - C(u,v,N)]\}, & \text{if } B(k,l) = \text{TRUE} \\ O, & \text{if } B(k,l) = \text{FALSE} \end{cases}$$

The boundary differential will be constant along the boundary between any two clusters, and it will represent the nominal spectral difference between contiguous regions in the original image. Using our previous example and assuming that the cluster of "corn" is larger than that of "gravel," the boundary differential will be located along the pixels of the "corn" cluster adjacent to the "gravel" cluster and will be described as:

$$Y(k,l) = \begin{cases} f[\textbf{corn–gravel}] & \text{if } B(k,l) = \text{TRUE} \\ O, & \text{if } B(k,l) = \text{FALSE} \end{cases}$$

where **corn** and **gravel** are the N-dimensional class values of corn and gravel as determined by the classifier, and f is a function of those vectors.

Two methods of unsupervised classification are particularly appealing in the context of boundary change detection: a global histogram classifier, or GHC, which works from an N-dimensional histogram, and a region growing classifier, or RGC, which operates directly on the N-band image.

a. Global Histogram Classification. The salient differences between the GHC and the RGC can be illustrated by use of our previous example of the cornfield and the road. The GHC would look at the histogram of the whole scene and subdivide it into several distinct classes according to the histogram statistics. Each class will have a nominal spectral vector, again determined by the statistics. Among these classes we would hope to find two whose spectral characteristics were similar to those of the corn and the gravel in our field. If there were several fields of corn and several gravel roads, an ideal GHC would create one class for all corn and one class for all gravel within the scene having the nominal values **corn** and **gravel**, respectively. Each cluster of corn would consist of pixels having the spectral value **corn**, and each cluster of gravel would consist of **gravel**-valued pixels. Now, if our cornfield is "typical" corn and if our road is made of "typical" gravel, the representation will be accurate. But what if, say, our field of corn is not "typical," that for some reason it is distinctly different in spectral characteristics than the rest of the corn in the scene. A GHC would do one of the following:

- create a new class for our particular region of "corn" (which would be good),
- classify the region as typical corn and assign the cluster the value **corn** (which would be bad), or,
- classify it as something else and assign it a class that does not reflect its spectral characteristics well (even worse).

The tradeoff involved in a GHC, insofar as it relates to our purposes, is between the number of classes created and the spectral accuracies to which the regions are represented by their clusters. In general, the more classes created, the higher the spectral accuracy, but as an undesirable side effect, there will be superfluous clusters created. With remotely sensed data, it is virtually impossible for a GHC to create exactly one cluster per object, and even if this were possible, the spectral accuracy would not necessarily be high, since individual regions are not best described by global spectral characteristics, at least in our context. Thus, even an ideal GHC cannot be regarded as lossless in general.

As far a boundary accuracy is concerned, since the GHC does not use spatial data, it cannot be expected to "know" where the gradients are high, and as a result, the boundaries created do not always correspond to the points of maximum gradient in the image. However, when the spectral characteristics assigned to two adjacent clusters are accurate, the boundary usually will be accurately, or at least reasonably, located. The main problem with the accuracy of the GHC is that, depending on the number of classes it produces, it may create superfluous clusters, fail to create clusters for some distinct regions, or both. Thus, superfluous boundaries may be created, and some borders may not be represented by boundaries if no cluster has been created for the object in the classified image. In general, the boundaries created by a GHC cannot be regarded as accurate representations of borders and, furthermore, even if the

boundaries were spatially accurate, the lack of spectral accuracy will lead to inaccurate boundary differentials.

b. Region Growing Classification. Some of the aforementioned problems with the GHC are avoided when using a RGC. Consider again our cornfield. An ideal RGC will create exactly one class and one cluster for each distinct region. Since each class is derived from a single region, the cluster representing it can be expected to have accurate spectral characteristics. Thus, the cluster of "corn" representing our particular region of "corn" and the region itself will be very similar in spectral characteristics. This will be true regardless of global spectral characteristics. Each cornfield and each gravel road in the scene will have exactly one cluster with spectral characteristics representative of that particular field or road. In general, it is possible for a RGC to be lossless.

The fact that in an ideal RGC there is a one-to-one correspondence between distinct regions, classes, and clusters gives a RGC two fundamental advantages over a GHC in our context: first, there will be one-to-one correspondence between boundaries and borders in the scene, allowing a much greater potential for spatial accuracy and, secondly, since each region and its cluster will have a unique class, the potential for spectral accuracy is much higher. These advantages make a RGC more attractive in terms of our objective, the detection of changes in boundaries.

B. XORing

Having indicated the boundaries in both $\mathbf{I}_1$ and $\mathbf{I}_2$, the general boundary change detection scheme is as follows. If in $\mathbf{I}_i$ there is a boundary which is not present in $\mathbf{I}_j$, then a boundary change has occurred. If the boundary is present in both $\mathbf{I}_i$ and $\mathbf{I}_j$, then, although there might have been spectral changes in the regions involved, the boundary between them remained invariate and therefore no boundary change was involved and there will be no boundary change indication. The process by which boundary changes are indicated is seen to be an exclusive OR of the boundaries in $\mathbf{I}_i$ with those in $\mathbf{I}_j$ (refer back to Figure 1 for examples).

Unfortunately, the problem of XORing boundaries is not as trivial as it might seem, especially when the images are Thematic Mapper (TM) data. The first obstacle to confront is intertemporal, and even interband, misregistration. A prerequisite to satisfactory performance of any type of XOR requires that the two images be registered correctly, i.e., that each pixel in each band and each time correspond to exactly the same spatial coordinates in the scene. If $\mathbf{I}_1$ and $\mathbf{I}_2$ aren't correctly registered, the boundaries of an object present in both scenes will differ, and the algorithm will indicate a change in boundaries. Another problem especially apparent in Landsat images is that one image may be slightly warped compared to another. In this case, it is usually possible to register a portion of the images correctly while maintaining reasonably good, though not perfect, registration in the rest of the image. Resampling one image using control points established by the other image ("warping" or "rubber sheeting") can further reduce inaccuracy in registration, but this is a complicated and computationally inefficient process. Furthermore, the establishment of control points is difficult without human intervention. Even with rather sophisticated techniques, anomalies will still exist, though they will most likely be limited in size to a fraction of a pixel.

To compute an XOR, we obtain boundaries from the two times, $\mathbf{B}_1$ and $\mathbf{B}_2$. The boundaries indicated for a given object present in both times can be expected to differ slightly from one another both spatially and spectrally. The spatial difference (assuming it is small) can be attributed to misregistration and to the operation of the classifier or to the inherent inaccuracies in a general boundary. Also, any changes in illumination or spectral characteristics of either the background or the object may lead to slightly different boundaries. Thus, it is to be expected that the object's boundaries will not overlap perfectly from $\mathbf{I}_1$ to $\mathbf{I}_2$. If an XOR were performed for the boundaries of the object in both scenes, the cancellation would be imperfect and there would be a broken edge as an unwanted remainder. To remedy this, a neighborhood XOR could be used as such: Assuming k,l is a boundary pixel in $\mathbf{I}_i$, if there are any boundary pixels within a specified neighborhood around k,l in $\mathbf{I}_j$, they are considered part of the edge of a common object and are cancelled. The larger the neighborhood, the more complete the cancellation. But the larger the neighborhood, the higher the chances of the boundary points of one object being cancelled out by the boundary points of a different object, or of a valid boundary point in one time becoming cancelled by a mistakenly indicated boundary point in the other time.

This leads to the development of a technique by which an object is assigned only one boundary whether it is present in either or both times, as such:

- $B_i(k,l) = B_j(k,l) = \text{TRUE}$, if the boundary is present in both times,
- $B_i(k,l) = B_j(k,l) = \text{FALSE}$, if the boundary is present in neither time,
- $B_i(k,l) = \text{TRUE}$ and $B_j(k,l) = \text{FALSE}$ if the boundary is present at time i and not at time j.

Having defined boundaries thusly, XORing is a much more straightforward procedure.

III. METHODS

Two methods of boundary change detection will now be presented, one which uses gradient operators and difference classification, and another which uses bitemporal classification. Both of these methods represent boundaries in the advantageous fashion described above, though they are quite different in their approaches.

A. METHOD A

As discussed earlier, the general boundaries obtained from gradient methods are unsatisfying in terms of accuracy and precision. Furthermore, when using an XOR process general boundaries tend to be especially intractable since they are spatially two-dimensional and may be quite different in both times. Thus, general boundaries do not seem to lend themselves readily to our purpose. However, there is an important simplification which may be made. First, it should be noted that for any boundary change to occur, there must be a corresponding spectral change. The boundaries that have changed will be a subset of the boundaries of the regions undergoing spectral change. All other boundaries are not of interest. Therefore, we may restrict our attention to the boundaries of the regions having spectral changes. These boundaries may be precisely determined by classifying the difference **D** between $\mathbf{I}_1$ and $\mathbf{I}_2$. Thus, the XOR need only be performed along the boundaries of the classified difference image, **S**. This is equivalent to projecting $\mathbf{E}_1$ and $\mathbf{E}_2$ onto the spatially single-dimensional boundary space and then XORing this projection. The method proceeds in the following steps:

1. Preprocessing. First, $\mathbf{I}_1$ and $\mathbf{I}_2$ are registered to within one pixel of each other. Then their intensities are adjusted such that the standard deviations of their histograms are identical. This accounts for the difference in illumination between scenes.

2. Edge Creation. $\mathbf{E}_1$ and $\mathbf{E}_2$ are created from $\mathbf{I}_1$ and $\mathbf{I}_2$ as described in II.A.1.

3. Image Differencing. The images from the two times are differenced on a pixel-by-pixel basis in each band, and the absolute value of those differences is taken. This difference image, **D**, will have N bands.

4. Difference Classification. The difference image **D** is classified to produce **S** using any one of a variety of techniques:

a. Thresholding. The magnitudes of the spectral vectors in **D** may be taken and then thresholded. The result will be a classified image with two classes, one for regions in which spectral change has occurred, and one for regions in which no spectral change has occurred. Due to the simplicity of this technique, anomalies can arise in the presence of certain complex spectral changes. These will be explained in the following section.

b. Global Histogram Classification. A GHC elaborates upon thresholding by making it a multispectral process. It is desired that one class be created for every type of spectral change greater than a certain magnitude. Consider our cornfield once again. Assume that in the first sensing time ($\mathbf{I}_1$), there was no road and that the field was not of corn, but of squash. There would be two pertinent difference vectors, **sc** (the N-band difference between **squash** and **corn**) and **sg** (the difference between **squash** and **gravel**). The magnitudes of both of these vectors will be relatively high, though they will be markedly different. Using a simple thresholding procedure as in the preceding section, no boundary would be detected since the magnitude of both **sc** and **sg** will surpass the threshold at all points and, therefore, there would be no boundary for the difference between the road and the squash. Using a GHC, this problem could be avoided by creating separate classes for both of the difference vectors, **sc** and **sg**.

c. Region Growing Classification. Although a GHC will yield good results in general, a RGC will allow better classification of differences of marginal magnitudes since there will be less likelihood that the region will be broken up.

5. Boundary Indication. The boundaries of the classified difference image are determined as prescribed in II.A.2.

6. XOR Decision Processes. The final step in this method is to project $\mathbf{E}_1$ and $\mathbf{E}_2$ onto the boundaries and then to compare these projections. Since the values of $\mathbf{E}_1$ and $\mathbf{E}_2$ are not binary, a decision process somewhat more sophisticated than an XOR must be used. An example of such a process is as follows:

$$R(k,l) = \begin{cases} \text{TRUE, if } | E_i(k,l) - E_j(k,l) | > \delta_1, \text{ and} \\ E_i(k,l) < \delta_2, \text{ (where } \delta_2 \text{ is the nominal minimum} \\ \text{value for an edge point to be possible boundary} \\ \text{point)} \\ \text{FALSE, otherwise.} \end{cases}$$

Since the values of $\mathbf{E}_1$ and $\mathbf{E}_2$ will vary along any given boundary, it is likely that for any threshold chosen, there will be some boundaries in the result, **R**, which will be broken, and **R** will therefore contain inconsistent boundaries. In cases such as these, further decision processes are required to determine whether or not the boundary has actually changed.

B. METHOD B

Method B, while intuitively more straightforward than method A, requires more computational time due to its more complex classification scheme. The idea behind method B is to detect objects which were present in one time but not in another or have otherwise changed their spatial characteristics (shape, location). Since this is the ultimate goal of boundary change detection, method B can be regarded as the most direct attack on the problem.

The key concept in this method is that if an object were indeed present in both sensing times and did not change in spatial characteristics, its spatial characteristics could be represented adequately by a single cluster. If two clusters were produced for the object in separate scenes, the boundaries so obtained would differ somewhat due to a number of factors (especially subpixel misregistration and the idiosyncracies in the operation of the classifier), making a direct XOR of the boundaries impossible. Thus, it would be reasonable to create only one cluster per object, whether it exists in only one or both scenes. This is accomplished with bitemporal classification, which produces a single 2N-band classified image from two temporally distinct N-band images. Since there will be only one cluster per object, it is useful to think of two N-band spectral classes existing for each cluster, and to regard $\mathbf{C}_i$ as the N-band classified image with the classes corresponding to $\mathbf{I}_i$.

Assuming that $\mathbf{C}_1$ and $\mathbf{C}_2$ are so produced and their (identical) boundaries indicated, the question arises, which of the boundaries are real, and which are not? If an object is present in both scenes, its boundaries in both $\mathbf{C}_1$ and $\mathbf{C}_2$ will be real, i.e., correspondent to actual borders in their respective scenes. We will call such boundaries true boundaries. Now, if an object is present in scene i and not in scene j, the classifier will nevertheless create a cluster which will exist in both $\mathbf{C}_i$ and $\mathbf{C}_j$, and the boundaries of that cluster will be true in $\mathbf{C}_i$. But if there aren't corresponding borders in scene j, the boundaries of that cluster in $\mathbf{C}_j$ will not be true in the sense described above. Boundaries which do not correspond to borders in their respective scene will be termed false boundaries. Only if a boundary is true for one scene and false for the other do we wish to indicate a change. Since we are interested only in indicating a changed boundary, it is apparent that boundary change detection can be accomplished by determining which boundaries are true and which are false and then performing an XOR of the boundaries in $\mathbf{C}_1$ with those in $\mathbf{C}_2$.

The four distinct steps in method B are:

1. Preprocessing. The images are preprocessed as in step 1 of method A.

2. Bitemporal Classification. The images $\mathbf{I}_1$ and $\mathbf{I}_2$ are treated as a single 2N-band image in the classification process. Either a GHC or a RGC may be used for this step, though a RGC may yield more useful results, as previously discussed. Note that at this point, a general change detection procedure could be implemented using the difference between $\mathbf{C}_1$ and $\mathbf{C}_2$.

3. Boundary Indication and Calculation of Boundary Differentials. The boundaries $\mathbf{B}_1$ and $\mathbf{B}_2$, and their differentials $\mathbf{Y}_1$ and $\mathbf{Y}_2$, of the classified image are found using the techniques of II.A.2.

4. Decision and XOR. The decision as to whether a boundary is true or false can be determined by comparing $Y_1(k,l)$ and $Y_2(k,l)$ to a minimum threshold. Assuming proper operation of the classifier, there will be at least one true boundary per indicated boundary. Therefore, if a boundary is indicated but is determined to be false in either $\mathbf{Y}_1$ or $\mathbf{Y}_2$, that boundary would be indicated as a change, i.e., R(k,l) = TRUE.

Unfortunately, the above decision process will be unsatisfactory in most cases since an ideal classifier is difficult to construct. Also, in general a false boundary will have a nonzero boundary differential which in some cases might be higher than the boundary differential of a true boundary. Here, a single threshold cannot be selected to yield correct results. To reduce the possibilities of this occurring, it is useful to introduce another threshold based on the difference between $\mathbf{Y}_1$ and $\mathbf{Y}_2$, which is to be used in a decision scheme such as that in section III.A.2. If the difference is minimal, then the likelihood of the boundary being either true or false in both times is very high. If the difference is high, say, $Y_i(k,l) >> Y_j(k,l)$, it would seem to indicate that the likelihood of $B_i(k,l)$ being true is greater than the likelihood of $B_j(k,l)$ being true.

While taking this into account, it should be noted that, in the case of, for instance, an object increasing in brightness relative to its background, $\mathbf{Y}_2$ will have a higher magnitude than $\mathbf{Y}_1$ for that object. Therefore, care must be taken not to automatically consider a boundary to be false simply because its $\mathbf{Y}_1$ will be much lower in comparison. By using a careful combination of these two types of thresholds, satisfactory results may be obtained in many circumstances.

IV. CONCLUSION AND RESULTS

The approaches to boundary change detection presented in this discussion have a high theoretical potential for accuracy. In general, their accuracy will be limited only by the performance of the classifiers used. While the

approach used in method B may be regarded as more direct and potentially more accurate (since the boundaries are calculated directly from the images, not their differences), it is more sensitive to the performance of the classifier and requires more computational time. Method A, while less accurate in general, seems more suited for an on-board process, since it involves a much simpler classification scheme and fewer computational requirements.

Further research into this topic should provide more information on the optimal classification scheme and, we hope, an adaptive means of choosing optimal thresholds for the processes described herein.

George M. Haley (BS, University of Notre Dame) is on the Technical Staff at Santa Barbara Research Center and will receive a Master's degree in electrical engineering from the University of California at Santa Barbara this year. His current interests include digital signal processing, image data compression, and optimal hardware configurations for signal processing algorithms. He is a member of IEEE and Eta Kappa Nu.