THE UNIVERSITY OF MICHIGAN

COMPUTING RESEARCH LABORATORY[1]

ADAPTIVE GRAY SCALE MAPPING TO REMOVE REGISTRATION NOISE IN DIFFERENCE PICTURES

Edward J. Delp and Thomas F. Knoll

CRL-TR-5-84

JANUARY 1984

Room 1079, East Engineering Building
Ann Arbor, Michigan 48109
USA
Tel: (313) 763-8000

engn
UMR1016

Adaptive Gray Scale Mapping to Remove

Registration Noise in Difference Pictures

Edward J. Delp and Thomas F. Knoll

Computer, Information, and Control Engineering

and

Department of Electrical and Computer Engineering

University of Michigan

Ann Arbor, Michigan

Abstract

Difference pictures are used in various image processing applications such as change detection, radar imaging, remote sensing, and biomedical image analysis. The difference picture, or difference image, is found by subtracting one image from another. One practical problem with difference pictures is that, if the images are not in perfect spatial registration before subtraction, the difference picture will contain artifacts caused by incomplete cancel'ation of the unchanged background objects. These artifacts (registration noise) show up as extraneous light and dark regions on either side of the background objects. Usually, this noise is ignored, or is reduced by either smoothing (blurring), or thresholding the difference image. This paper describes a new method to reduce registration noise using adaptive gray scale mapping. This simple digital filter reduces registration noise as well as, or better than, previous methods, without degrading the real difference image.

1. Introduction

The difference picture can be simply thought of as an image that indicates change between two images. The difference picture is found by spatially aligning the two images as well as possible and, pixel by pixel, subtracting one image from the other. To avoid having to process negative gray level values, the difference picture gray levels are often offset so that zero (no difference) is represented as middle gray. We shall assume in this paper that there is no appreciable illumination difference between the two images, so that the difference picture can be used as a measure of what is actually different between the two scenes. In many applications one would have to adjust for illumination effects before the difference picture could be found.

Difference pictures have been used to detect changes in aerial photographs [9], radar images [6], Landsat images [8], and x-rays of bones [10]. They are also used in digital subtraction angiography [1], [5]. All of these applications rely on the fact that any feature that has not changed between the two images is canceled in the difference picture. This cancellation is useful, either because it greatly reduces the number of features that need to be considered (by human or machine) in later analysis, or because features that were obscured in the background are revealed in the difference picture.

Another application of difference pictures is motion detection [2], [3], [7]. This application is different from the above applications in that, instead of features or objects appearing or disappearing (or changing size), they are moving. If the frames are closely spaced in time compared to the object speeds, the objects will have moved only a few pixels between the frames. There will be areas in the difference picture at the leading edges of the objects, caused by the objects covering the background, and areas at the trailing edges, caused by the background being uncovered. These areas present a problem, as they look exactly like registration

noise. (Registration noise is also caused by a displacement between the two frames—it is as if the entire background is moving). It is impossible for any local operation to remove registration noise from these types of difference pictures without destroying the motion information. This situation will not be considered in this paper.

However, if the frames are widely spaced in time compared to the object speeds, the objects will have moved a large distance (several object diameters) between the frames. This can be thought of as simple change detection; objects have disappeared from some locations, and other objects have appeared at different locations. This paper will consider the removal of registration noise from these difference pictures.

We have used difference pictures to detect molecules in tissue samples [4]. Two consecutive slices are used as samples, so they are similar in shape and features. One sample is digested with an enzyme that removes a specific molecule, while the other sample is untreated. Areas that contains the molecule will be darker in the control sample than in the treated sample, whereas areas that do not contain the molecule will have the same gray level in both samples. Thus the difference picture shows the locations and densities of the removed molecule.

One practical problem with difference pictures is that, if the two images are not in perfect spatial registration before subtraction, the difference picture will contain registration noise. Objects in the two images will not line up perfectly, causing extraneous light and dark areas in the difference picture. (We shall assume in this paper that the misregistration is due to changes in camera angle or small translational motion of objects in the scene, and not to changes in camera distance or perspective). Each misregistered object will have corresponding light and dark areas beside it. Which side the light and dark areas are on depends on the direction of the registration error, and on whether the object is lighter or darker than the

background.

Figures 1 and 2 show two slightly different images, and Figure 3 shows their ideal (no registration noise) difference image. Notice how the objects around the edge of the image cancel completely, because they are identical in both images. However, when the image in Figure 2 is shifted slightly before subtracting it from Figure 1, the "raw" difference image is found (shown in Figure 4). Notice the registration noise, caused by the misregistered objects around the edge of the image.

Even a slight misregistration will add noise to the difference image, possibly overwhelming the real differences between the images. Removing this noise will assist human or machine interpretation of the difference images. One method to reduce registration noise is to smooth the difference image; this has the side effect of blurring the real differences (see Figure 5). Another common method is to threshold the difference image, mapping all pixels in a range around middle gray to middle gray. This destroys much of the gray scale information in the picture, and is ineffective if the registration error is greater than one pixel (assuming the images contain sharp edges). Adaptive gray scale mapping (discussed later) is a new, simple digital filter that is effective in reducing registration noise, while preserving the real differences between the images.

2. Properties of Registration Noise

A simple property of registration noise is that it has a zero sample mean. The corresponding light and dark areas on either side of each background object spatially average to zero. To show this, let us compute the spatial average of a difference image containing only registration noise. Let $g(i,j)$ be the difference image between a discrete image $f(i,j)$, containing objects on a constant background, and a shifted version of that same image.

$$g(i,j) = f(i,j) - f(i-\Delta i, j-\Delta j)$$

This difference image will consist totally of registration noise. Let $f(i,j)$ be equal

to a constant c, outside some region $A \in I^2$ (I is the set of all integers). Let region $B \in I^2$ be the region A shifted by $(\Delta i, \Delta j)$. Summing the difference (noise) image over all non-zero values gives

$$\sum_{(i,j) \in A \cup B} \sum g(i,j) = \sum_{(i,j) \in A \cup B} \sum f(i,j) - f(i-\Delta i, j-\Delta j)$$

$$= \sum_{(i,j) \in A \cup B} \sum f(i,j) - \sum_{(i,j) \in A \cup B} \sum f(i-\Delta i, j-\Delta j)$$

$$= \sum_{(i,j) \in A} \sum f(i,j) - \sum_{(i,j) \in B} \sum f(i-\Delta i, j-\Delta j) + \sum_{(i,j) \in \bar{A} \cap B} \sum c - \sum_{(i,j) \in A \cap \bar{B}} \sum c$$

and using the fact that regions $\bar{A} \cap B$ and $B \cap \bar{A}$ have the same area

$$= \sum_{(i,j) \in A} \sum f(i,j) - \sum_{(i,j) \in B} \sum f(i-\Delta i, j-\Delta j)$$

$$= \sum_{(i,j) \in A} \sum f(i,j) - \sum_{(i,j) \in A} \sum f(i,j)$$

$$= 0$$

Thus, the registration noise caused by each object or set of objects will spatially average to zero.

The above proof assumes that the images can be modeled as objects on a constant background. (Non-uniform gray level objects are allowed—they are just treated as multiple, overlapping uniform objects). An image which could not be modeled as objects on a background is an infinite step reference scale. Here the registration noise is monochromatic (all light or all dark). Fortunately, most images can be approximately modeled as objects on a background.

In general, corresponding light and dark regions can be far apart, separated by as much as the largest object dimension. A large averaging neighborhood would be required for the registration noise to average out. The noise reduction problem is simplified if most of the corresponding light and dark areas of registration noise are close together. Then only a small neighborhood is required for the noise to spatially average to zero. Registration noise from small objects is close together, as is noise from long thin objects (if the end effects are ignored). For many images the

approximation that corresponding regions are near each other is a good one, because most images contain many more small objects than large objects. Small objects cause more noise per area of the object than large objects (a long thin object will cause just as much noise as a fat object of equal length).

Both smoothing and adaptive gray scale mapping will use the fact that registration noise is zero mean, and will assume that corresponding light and dark regions are near each other.

3. Current Methods Used to Remove Registration Noise

The best way to reduce registration noise is to register the images accurately before subtracting. Most of the work in the reduction of registration noise has addressed the problem by proposing better registration algorithms [6], [9]. Both linear and higher order geometrical transforms are used to align the images as accurately as possible. However, to completely prevent registration noise, the images must be aligned to sub-pixel accuracy. Even if the images are registered carefully beforehand, some noise is likely to remain in the difference picture.

Sometimes it is not possible to align the images accurately enough before subtraction. One case is subtractive television [7]. Here the subtraction is done by the camera; the only input to the computer is the difference image. Registration error caused by non-repeatability of the camera cannot be corrected beforehand, and must be removed from the difference image.

Another case where it is not possible to align the images accurately enough is our work with consecutive tissue samples [4]. Although the samples are very similar, they are not identical. While the basic outline of the samples can be aligned, small scale features, such as cells, usually cannot. Again, the noise must be removed after computing the difference image.

3.1. Smoothing to Remove Registration Noise

The simplest way to remove registration noise is to smooth the image by replacing each pixel by the average gray level of pixels in its neighborhood. This can be either an unweighted (simple) or a weighted average. If the averaging neighborhood is large enough, it will probably contain both of the corresponding light and dark areas of registration noise. Because the noise is zero mean, it will average out and be removed.

The problem with smoothing is that, if the averaging neighborhood is large enough to reduce the registration noise, it also blurs the real differences between the images. Figure 5 shows the the raw difference image in Figure 4 after simple smoothing. The registration noise is greatly reduced, but the real differences have been so blurred that they are hardly visible.

3.2. Thresholding to Remove Registration Noise

If the registration error is less than one pixel (slightly more if the images contain no sharp edges), thresholding can be used to reduce registration noise. When the registration error is small, the registration noise is likely to be less intense (closer to middle gray) than the real difference image. By setting pixels whose gray level is within a certain distance of middle gray to middle gray, noise can be reduced. Other pixels can either be left alone, or set to the corresponding maximum or minimum value.

One problem with thresholding is that it fails if the registration error is larger than one pixel. Registration noise caused by a sharp edge and an alignment error of only one pixel will be just as or more intense as the real differences between the images. If the image in Figure 4 is thresholded, it noise is totally uneffected.

4. Adaptive Gray Scale Mapping to Remove Registration Noise

Adaptive gray scale mapping is a filter that, for each pixel, uses its neighborhood to estimate whether the pixel is part of the true difference image, or part of the registration noise. A non-linear gray scale map is computed for each pixel, based on the amount of "brightness" and "darkness" in a neighborhood around that pixel. Each pixel is then passed through its own map to produce the filtered difference image.

Adaptive gray scale mapping can work much better than simple smoothing. Figure 6 shows Figure 4 after adaptive gray scale mapping. The noise is reduced just as much as in Figure 5, but the real difference image is not blurred.

Imagine the difference image as a sandbox. Where the difference image is light, there is a hill; where it is dark, there is a valley. Registration noise shows up as corresponding hills and valleys. Because registration noise has a zero sample mean, each hill contains exactly as much sand as is missing from its corresponding valley. In effect, all the registration noise does is scoop up sand from some parts of the image, leaving valleys (dark areas), and dump it in other parts, creating hills (light areas). Adaptive gray scale mapping tries to reverse this process.

To remove registration noise, the sand must be removed from the hills and put back in the valleys. One way to do this is to remove sand from all the hills (in proportion to their height), and dump it into all the valleys (in proportion to their depth). This continues until either there are no more hills, or there are no more valleys. The problem with this method is that all hills and all valleys are treated alike; isolated hills (probably part of the real difference image) are chopped just as much as are hills next to valleys (which are probably registration noise).

The algorithm is refined by using the assumption that corresponding light and dark areas (hills and valleys) caused by registration noise are near one another. Sand from a hill is used to fill only the nearby valleys. Isolated hills or valleys are

not effected, whereas areas that contain both hills and valleys that are near each other are smoothed out. The final image will now contain some areas with only hills, some areas with only valleys, but no areas with both hills and valleys.

This algorithm cannot be implemented exactly as described. The sand moving process is parallel and continuous. Sand is moved from all the hills to all the valleys at the same time. For a given hill and valley, the rate of sand transfer between them will depend inversely on their distance, and their *current* (not *initial*) height and depth. The algorithm would have to be approximated by many iterations, each moving a small amount of sand.

The algorithm is speeded by, instead of actually moving the "brightness" (sand) around, the final gray level of a pixel is estimated from the amount of "brightness" and "darkness" in the neighborhood around that pixel. For example, if a pixel is below the reference level (in a valley), and there is equal brightness and darkness (sand in hills and sand missing from valleys) in its neighborhood, the pixel is mapped to the reference level (the valley is filled in). Or, if there is twice as much darkness as brightness, the pixel is mapped to a gray level halfway between its current level and the reference level (the valley is halfway filled in). A non-linear gray scale map is computed for each neighborhood, based of the amount of brightness and darkness in it, and its center pixel is then mapped to a new value.

4.1. Implementation of Adaptive Gray Scale Mapping

Adaptive gray scale mapping has two parameters. The first is the radius Δ, of the square neighborhood around each pixel that sand could be moved to or from. If the neighborhood is too small, it will not contain both the light and dark regions of registration noise, and the algorithm will not remove the noise. However, for an object in the real difference image is not to be effected by the filter, it must be isolated from regions of opposite color by at least the neighborhood radius. As the radius is increased, more and more of the real difference image will not be isolated

enough, and will be mistakenly dimmed or completely removed. The ideal radius is large enough to reduce noise, but small enough to not excessively dim the real difference image. It must, in general, be found by trial and error for each class of difference images.

The second parameter is the reference gray level r, used to separate hills and valleys. A pixel brighter than the reference is considered to be part of a hill; a pixel darker is considered to be part of a valley. Normally, the reference level is set to the gray level in the difference picture that corresponds to no difference between the two input images. Areas that contain only registration noise will contain equal moments above and below this reference level. Because both moments are equal, the noise will be completely removed. If the reference level is set to some other level, the moments will not be equal, and the noise will be only partially removed.

However, if the background gray level has changed between the two images, a different reference level will achieve greater noise reduction. Registration noise in these areas will not average to zero, but instead to the difference in the background levels. Maximum noise reduction is then achieved when the reference level is set to the difference in background levels.

Two conditional local moments are computed for each pixel in the difference image. The first is the brightness moment

$$b(i,j) = \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} (f(m,n)-r) \, 1\{f(m,n)>r\}$$

where I is the indicator function. Pixels outside the image are assumed to be at the reference level. In the sand model, the brightness moment corresponds to the amount of sand in the hills near the pixel. The other moment computed is the darkness moment

$$d(i,j) = \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} (r-f(m,n)) \, 1\{f(m,n)<r\}$$

This corresponds to the amount of negative sand, or air, in the valleys near the

pixel.

The gray level of each pixel in the difference image is then mapped to a new gray level

$$
g(i,j) = \begin{cases} r + (f(i,j) - r) \dfrac{b(i,j) - d(i,j)}{b(i,j)} & \text{if } f(i,j) > r \text{ and } b(i,j) > d(i,j) \\[2ex] r - (r - f(i,j)) \dfrac{d(i,j) - b(i,j)}{d(i,j)} & \text{if } f(i,j) < r \text{ and } b(i,j) < d(i,j) \\[2ex] r & \text{otherwise} \end{cases}
$$

If a pixel is in the minority, i.e. a dark pixel in a mostly bright area, or a bright pixel in a mostly dark area, it is mapped to the reference level. If the pixel is in the majority, its contrast is toned down with a ramp function whose slope depends on the amount of disagreement in the area. For example, if there is three times as much dark as bright, the ramp will have slope two-thirds. (In the sandbox model, this corresponds to filling the valley one-third of the way).

Adaptive gray scale mapping can be done quickly. If the conditional moment calculations are implemented as written, execution time would be proportional to the image area times the neighborhood area. However, by using incremental updating to compute the moments, the execution time is reduced so that it is proportional to only the image area. This is done by first computing column subtotals, and then summing the subtotals. The moments are incrementally updated as the row is scanned, by adding in the leading subtotals, and subtracting out the trailing subtotals. When changing to the next row, each of the column subtotals is incrementally updated by looking at the leading and trailing rows of pixels.

4.2. Preservation of Neighborhood Mean Gray Level

Because registration noise has a zero mean, it is desirable that a filter that removes registration noise should not change the mean gray level of the difference picture. We shall show in this section that adaptive gray scale mapping approximately preserves the neighborhood mean gray level.

Computing the neighborhood sum of the difference picture around (i,j), assuming $b(i,j) \ge d(i,j)$

$$\sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} g(m,n) =$$

$$\sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} \left[r + \frac{b(m,n)-d(m,n)}{b(m,n)} (f(m,n)-r) I\{f(m,n)>r\} \right]$$

Assuming that neighborhoods of pixels in the neighborhood of (i,j) all have about the same amount of brightness and darkness, i.e.

$$b(m,n) \approx b(i,j)$$

$$d(m,n) \approx d(i,j)$$

This approximation holds well if the neighborhood size is large compared to the characteristic dimension of the texture field. Using this approximation

$$\sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} g(m,n) \approx \left[\sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} r \right] + \frac{b(i,j)-d(i,j)}{b(i,j)} \times$$

$$\sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} (f(m,n)-r) I\{f(m,n)>r\}$$

$$= \left[\sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} r \right] + \frac{b(i,j)-d(i,j)}{b(i,j)} b(i,j)$$

$$= \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} r + \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} (f(m,n)-r)$$

$$= \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} f(m,n)$$

This can also be shown when $b(i,j) \le d(i,j)$. Thus the neighborhood sum, (and thus the neighborhood mean), are approximately preserved by adaptive gray scale mapping.

5. Example: Molecule Detection Using Consecutive Tissue Samples

We have used difference pictures to detect the locations and densities of specific molecules in tissue samples [4]. Two consecutive tissue slices are used. One is kept as a control, while the other is treated with an enzyme that digests the molecule in question. Both samples are stained and photographed. The

photographs are then digitized and registered as well as possible. The treated image is then subtracted from the control image. The resulting difference picture shows, in theory, the locations and densities of the molecule in question.

The problem with this method is that the two samples are not identical. While the basic outline of the samples can be aligned, the smaller scale features, such as cells, cannot. This is because the cell dimensions are of the same order as the slice thickness. Although each sample will have about the same density of cells, the cells will not line up between the two images. The unfiltered (raw) difference picture will contain of lot of registration noise.

Figures 7 and 8 show the digested and the control images before subtraction. Figure 9 shows their raw difference image. Notice how the registration noise (extraneous light and dark regions) makes it difficult to see which areas are really darker in the control image.

The raw difference image was then filtered using adaptive gray scale mapping. The reference gray level was set to the value corresponding to no difference between the input images. The region radius was found be trial and error—starting with a small value and increasing it in steps until enough noise reduction was obtained. The resulting filtered difference image is shown in Figure 10. Notice how it is much easier to see where, and how much, the control sample is darker than the treated sample (and thus the locations and densities of the molecule in question).

In this application, the texture of the tissue areas in the difference images is not important, only their mean gray levels. The texture is an artifact caused by the cells, and should be ignored. This texture could be removed by smoothing the difference image, but that would blur details in other areas of the difference image.

6. Conclusions

Difference pictures are useful in detecting change between two images. One problem with difference pictures is that, if the two images are not aligned perfectly before subtraction, extraneous light and dark regions will appear in the difference image. These artifacts are caused by the misaligned background objects in the two images not canceling completely in the difference picture. We have presented a new algorithm using adaptive gray scale mapping, that reduces registration noise from difference pictures. It removes noise as well as or better than existing methods, while not blurring the real differences between the images. It also preserves the mean gray level, which is desirable because registration noise does not change the mean gray level.

References

[1] C. K. Chow, S. K. Hilal, K. E. Niebuhr: X-ray Image Subtraction by Digital Means, *IBM Journal of Research and Development* Vol. 17, 1973, pp. 206-218

[2] R. Jain: Extraction of Motion Information from Peripheral Processes, *IEEE Transactions of Pattern Analysis and Machine Intelligence* Vol. PAMI-3, 1981, pp. 489-503

[3] R. Jain: Dynamic Scene Analysis Using Pixel-Based Processes, *Computer* Vol. 14, August 1981, pp. 12-18

[4] T. F. Knoll, L. L. Brinkley, E. J. Delp: A Method for the Analysis of Histological Images Using Difference Pictures, Submitted to *Analytical and Quantitative Cytology*

[5] R. S. Ledley, L. M. Arminski, L. S. Rotolo, T. J. Golab, C. M. Park: TEXAC-Based Digital Fluoroscopy, *IEEE Computer Society Conference on Pattern Recognition and Image Processing*, 1982, pp. 530-533

[6] R. L. Lillestrand: Techniques for Change Detection, *IEEE Transactions on Computers* Vol. C-21, 1972, pp. 654-659

[7] M. Onoe, N. Hamano, K. Ohba: Computer Analysis of Traffic Flow Observed by Subtractive Television, *Computer Graphics and Image Processing* Vol. 2, 1973, pp. 377-392

[8] C. J. Robinove, P. S. Chavez, D. Gehring, R. Holmgren: Arid Land Monitoring Using Landsat Albedo Difference Images, *Remote Sensing of Environment* Vol. 11, 1981, pp. 133-156

[9] M. S. Ulstad: An Algorithm for Estimating Small Scale Differences Between Two Digital Images, *Pattern Recognition* Vol. 5, 1973, pp. 323-333

[10] R. L. Webber, U. E. Ruttimann: Digital Subtraction of Radiographs as an Aid for Detecting Pathologic Changes in Periodontal Bone, *IEEE 1982 International Workshop on Physics and Engineering in Medical Imaging*, pp. 130-136
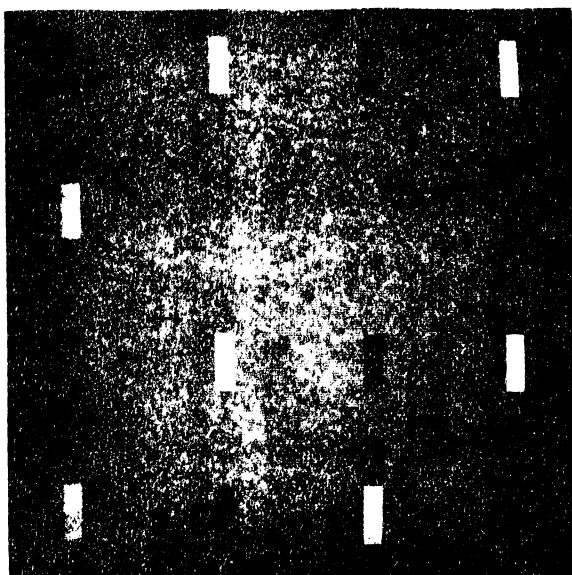
Figure 1: A 256x256 pixel test image that contains both dark and light objects on a constant background.
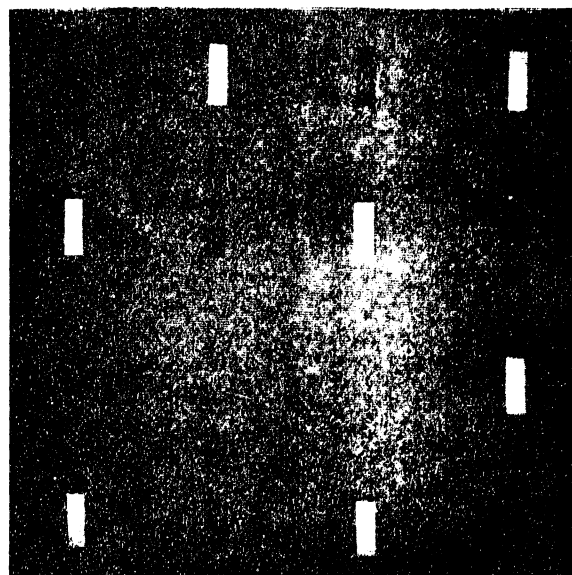

Figure 2: A second test image, almost the same as Figure 1, except for the center objects.
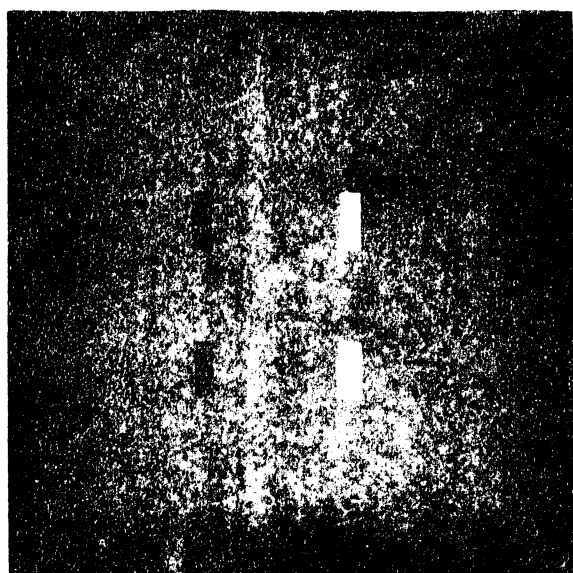

Figure 3: The ideal difference image found by subtracting Figure 2 from Figure 1. Note how the objects around the edge completely cancel out since they are the same in both images.
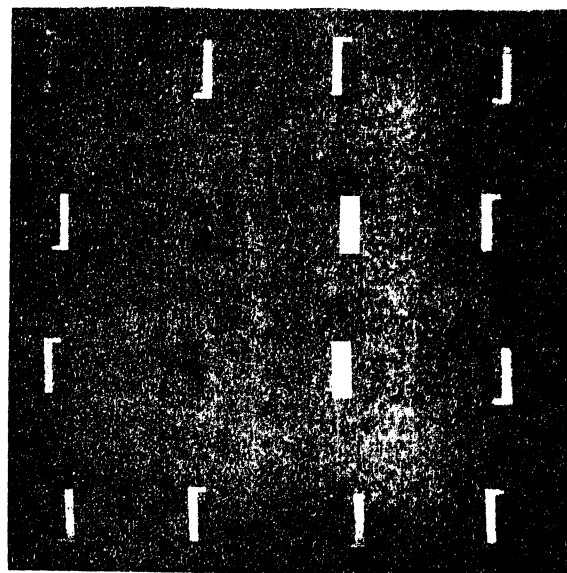

Figure 4: The raw difference image between Figure 1 and a shifted version of Figure 2. Figure 2 has been shifted by 4 pixels in the horizontal direction and 2 pixels in the vertical.
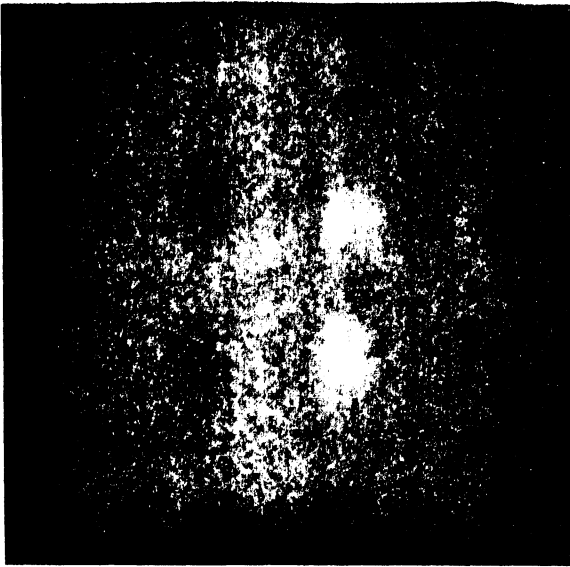
Figure 5: The raw difference image from Figure 4 after simple smoothing with a neighborhood of 25 by 25 pixels. The registration noise is reduced, but the difference image is blurred.
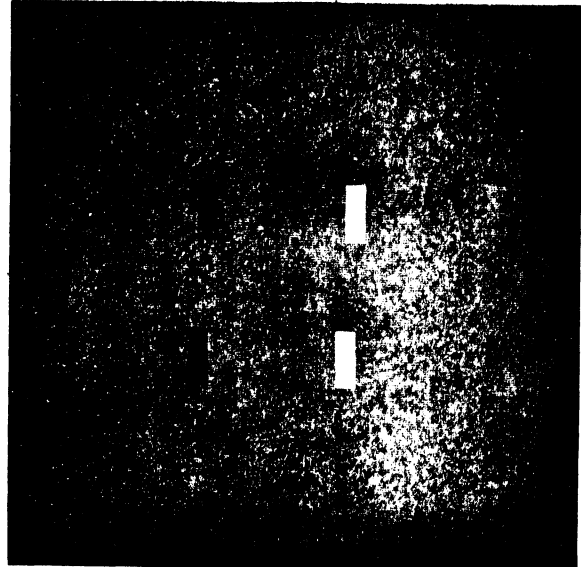

Figure 6: The raw difference image from Figure 4 after adaptive gray level mapping with a neighborhood of 25 by 25 pixels. The registration noise is almost completely removed, and the difference image is left unchanged.
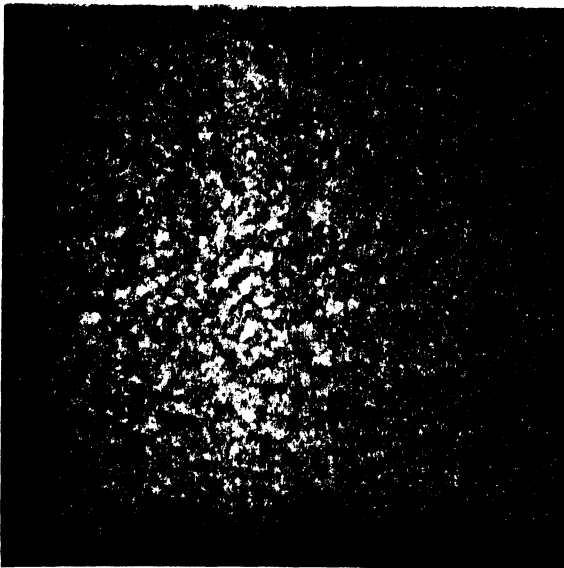

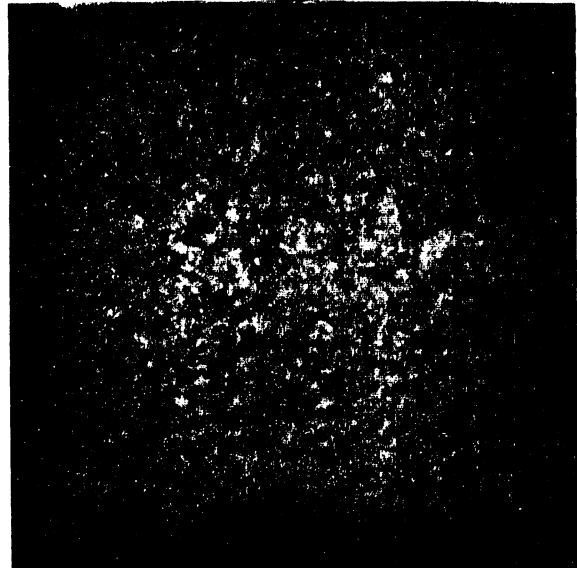**Figure 7**: A 256x256 pixel image of the control tissue sample.

Figure 8: A 256x256 pixel image of the digested tissue sample. This sample slice is consecutive with the sample in Figure 7. This sample has been treated to removed the molecule in question.
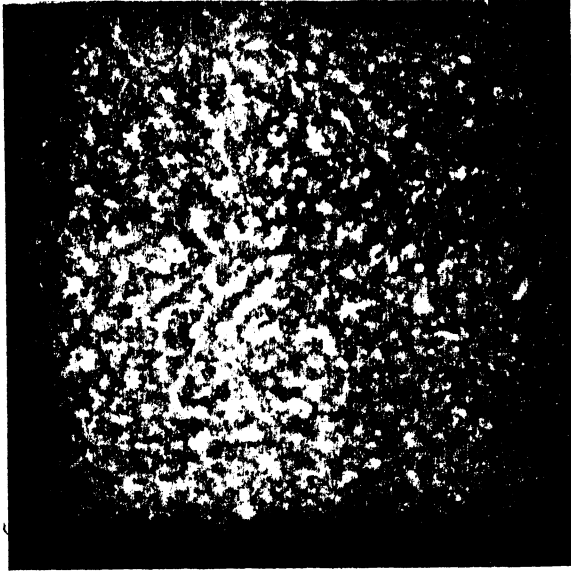
Figure 9: The raw difference image found by subtracting Figure 8 from Figure 7. Notice how the registration noise makes it difficult to see where the images are really different.
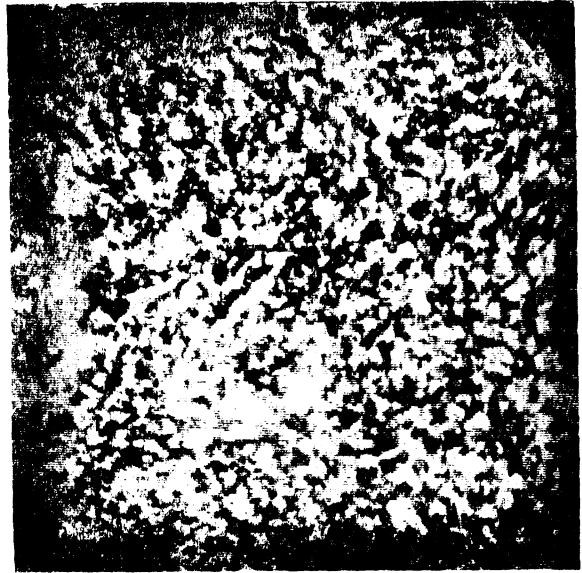

Figure 10: The filtered difference image found by using adaptive gray scale mapping on Figure 9. A neighborhood radius of 7 was used. Notice how it is easier to see the areas where the control image is darker than the digested image.