COMPUTER VISION, GRAPHICS, AND IMAGE PROCESSING 33, 129–137 (1986)

Adaptive Gray Scale Mapping to Reduce Registration Noise in Difference Images*

THOMAS F. KNOLL

Computer, Information, and Control Engineering, University of Michigan, Ann Arbor, Michigan

AND

EDWARD J. DELP

School of Electrical Engineering, Purdue University, West Lafayette, Indiana

Received August 25, 1984; revised February 19, 1985; accepted June 26, 1985

Difference images are used in various image processing applications such as change detection, radar imaging, remote sensing, and biomedical image analysis. The difference image, or difference picture, is found by subtracting one image from another. One practical problem with difference images is that, if the images are not in perfect spatial registration before subtraction, their difference image will contain artifacts caused by incomplete cancellation 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 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, with less degradation of the actual differences between the images.

1. INTRODUCTION

The difference image can be simply thought of as an image that indicates change between two images. It 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 gray levels are often offset so that zero (no difference) is represented as middle gray. If the images are not in perfect alignment before subtraction, their difference image will contain artifacts caused by incomplete cancellation of the unchanged background objects. These artifacts will be referred to here as registration noise. We shall assume in this paper that there is no appreciable illumination difference between the two images. In many applications one would have to adjust for illumination effects before the difference image is found.

Difference images have been used to detect changes in aerial photographs [7], radar images [5], Landsat images [6], and X rays of bones [8]. They have also been used in digital subtraction angiography [1, 4]. We detected the spatial distribution of hyaluronate across a tissue section by computing the difference image between a control tissue section and an adjacent digested (to remove any hyaluronate) tissue section [3]. All of these applications rely on the fact that any feature that has not changed between the two images is canceled in the difference image. This cancellation is useful, either because it greatly reduces the number of features that need to be

*This research was partially sponsored by the Biomedical Research Council and the Office of the Vice-President for Research, University of Michigan. T. F. Knoll is supported by an Office of Naval Research graduate fellowship. Send reprint requests to Edward J. Delp, School of Electrical Engineering, Purdue University, West Lafayette, Indiana 47907.

0734-189X/86 $3.00

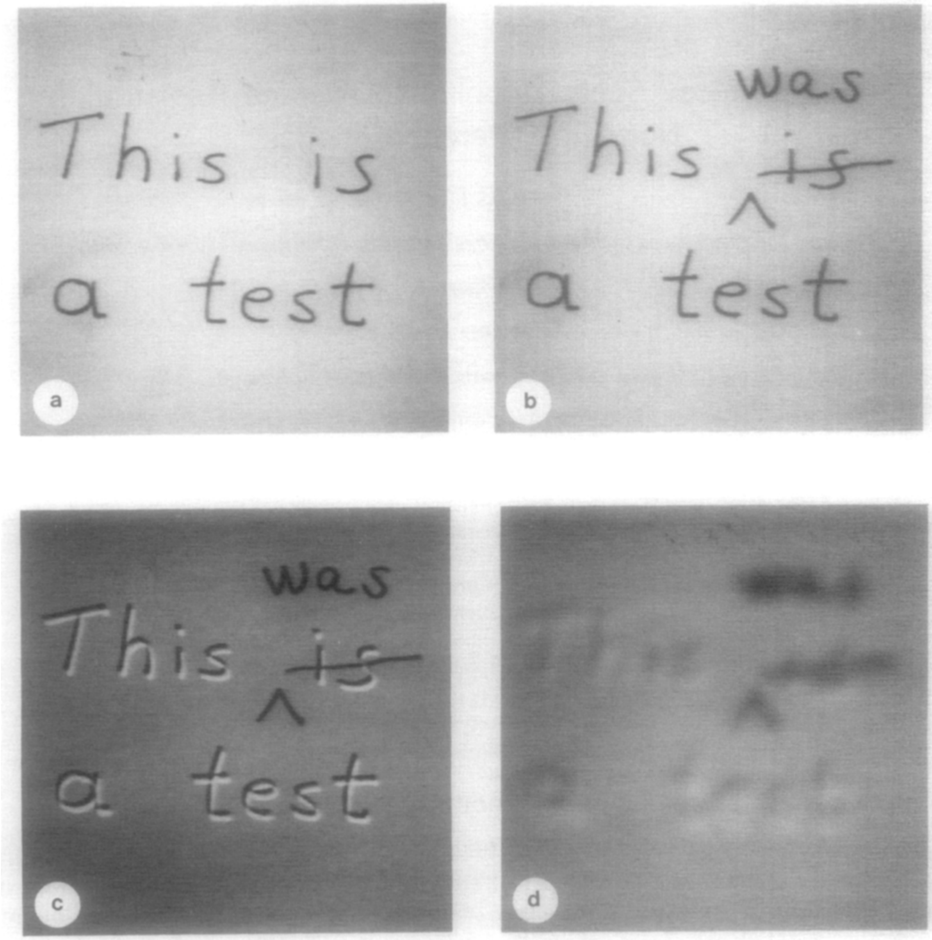

FIG. 1. (a) A 350 by 350 image of handwritten text. (b) Same as (a), except with edit marks. (c) Difference image between (a) and (b). Notice the light and dark areas of registration noise. (d) Image from (c) after smoothing. The registration noise is reduced but the edit marks are blurred. (e) Image from (c) after adaptive gray scale mapping using $\Delta = 15$. Notice the reduction in registration noise. (f) The edit marks can now be extracted by simple thresholding of (e).

considered (by human or machine) in later analysis, or because features that were obscured in the background are revealed in the difference image.

One other application of difference images is motion detection [2]. This application is different from the above applications in that, instead of features or objects appearing or disappearing (or changing size), they are moving. There will be areas in the difference image at the leading and trailing edges of moving objects, where the objects are covering or uncovering the background. These areas present a problem, as they look exactly like registration noise. It is impossible for any local operation to remove registration noise from these types of difference images without destroying the motion information. This paper will not consider motion detection.

One practical problem with difference images is that, if the two images are not in perfect spatial registration before subtraction, their difference image will contain registration noise. Objects in the two images will not line up perfectly, causing

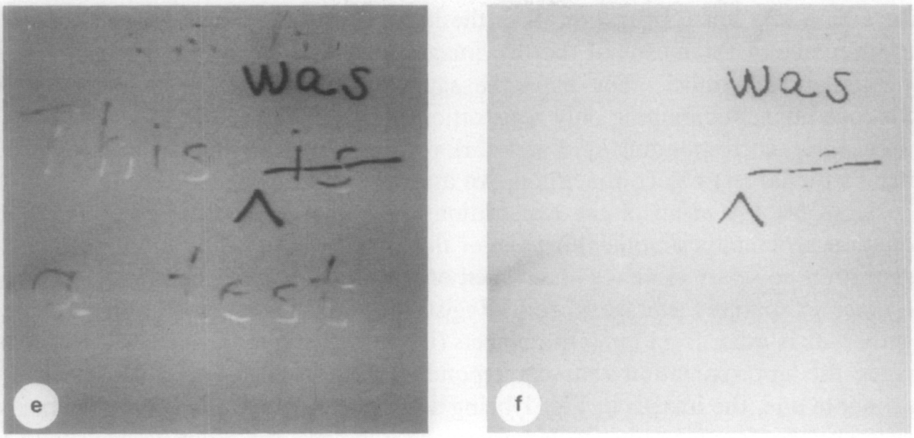

FIG. 1—*Continued.*

extraneous light and dark areas in their difference image. (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 1a and b show two images of handwritten text, before and after editing. Figure 1c shows their "raw" difference image. Ideally the difference image would show only the edit marks; however, because the images were not in perfect alignment before subtraction, there is a great deal of registration noise in the difference 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 Fig. 1d). Another common method is to threshold the difference image, mapping all pixels in a range around middle gray to middle gray. This 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 mean of a difference image containing only registration noise. Let $g(i, j)$ be the difference image between a discrete image $f(i, j)$, and a shifted version of that same image:

$$g(i, j) = f(i, j) - f(i - \Delta i, j - \Delta j). \tag{1}$$

Subtraction is a linear operation, thus the mean of the difference image is equal to the difference of the means of the two images. Because the images are only shifted versions of each other, they have the same mean. Therefore the mean of the difference image (containing only registration noise) is zero.

In general, corresponding light and dark regions can be far apart, separated by as much as the largest object dimension. An averaging neighborhood would have to be very large for the mean of the registration noise in it to be near zero. 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 reasonable one, the images in Fig. 1 being a good example. For many other images this is not the case (see Fig. 2). However, these images can often be dealt with by applying an edge detector (such as the Sobel operator) to each image before computing the difference image; corresponding light and dark regions of registration noise will then be close to each other.

3. CURRENT METHODS USED TO REDUCE 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 [5, 7]. 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 image.

The simplest way to reduce registration noise is to smooth the image by replacing each pixel by the average gray level of pixels in its neighborhood. 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 1d shows the the raw difference image in Fig. 1c after smoothing. The registration noise is reduced, but the real differences have been so blurred that they are hardly visible.

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 Fig. 1c is thresholded, its noise is almost totally unaffected.

4. ADAPTIVE GRAY SCALE MAPPING TO REDUCE 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 1e shows Fig. 1c after adaptive gray scale mapping. The noise is reduced just as much as in Fig. 1d, but the real difference image is not blurred. The edit marks can now be extracted by simple thresholding (Fig. 1f).

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 affected, 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 is approximated mathematically as follows. First, two conditional local moments are computed for each pixel in the difference image, the brightness moment

$$b(i,j) = \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} (f(m,n) - r) I\{f(m,n) > r\} \tag{2}$$

and the darkness moment

$$d(i,j) = \sum_{m=i-\Delta}^{i+\Delta} \sum_{n=j-\Delta}^{j+\Delta} (r - f(m,n)) I\{f(m,n) < r\} \tag{3}$$

where I is the indicator function, r is the reference gray level representing no difference between the two images, and Δ is the local neighborhood radius. Pixels outside the image are assumed to be at the reference level. The brightness moment

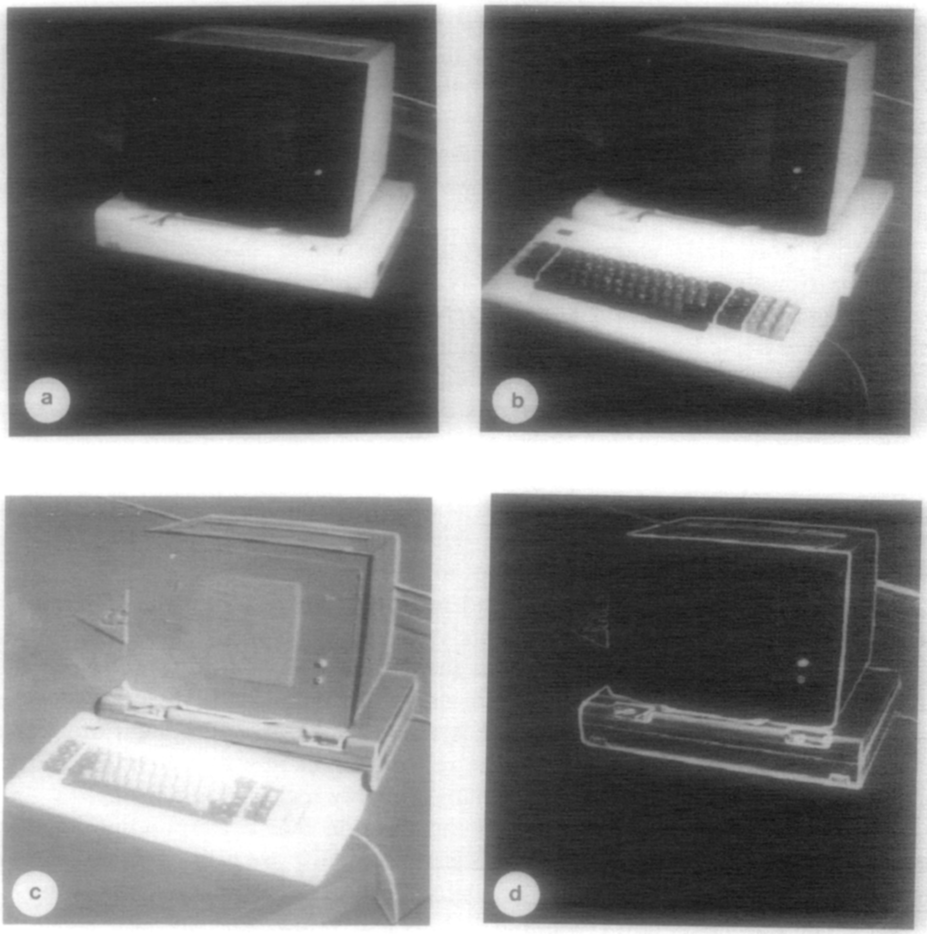

FIG. 2 (a) A 460 by 460 image of a terminal without a keyboard. (b) Same as (a), except with a keyboard. (c) Difference image between (a) and (b). Notice that, because of the size of the objects, not all of the corresponding light and dark areas of registration noise are near each other. (d) Image from (a) after edge detection using the Sobel operator. (e) Image from (b) after edge detection using the Sobel operator. (f) Difference image between (d) and (e). The corresponding light and dark areas of registration noise are now near each other. (g) Image from (f) after adaptive gray scale mapping using $\Delta = 10$. Notice the reduction in registration noise. (h) Image from (g) thresholded to detect the new edges.

corresponds to the amount of sand in the hills near each pixel, while the darkness moment corresponds to the amount of negative sand, or air, in the valleys.

The parameter Δ specifies the size 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 not to be affected by the filter, it must be isolated from regions of opposite color by at least the neighborhood radius. 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.

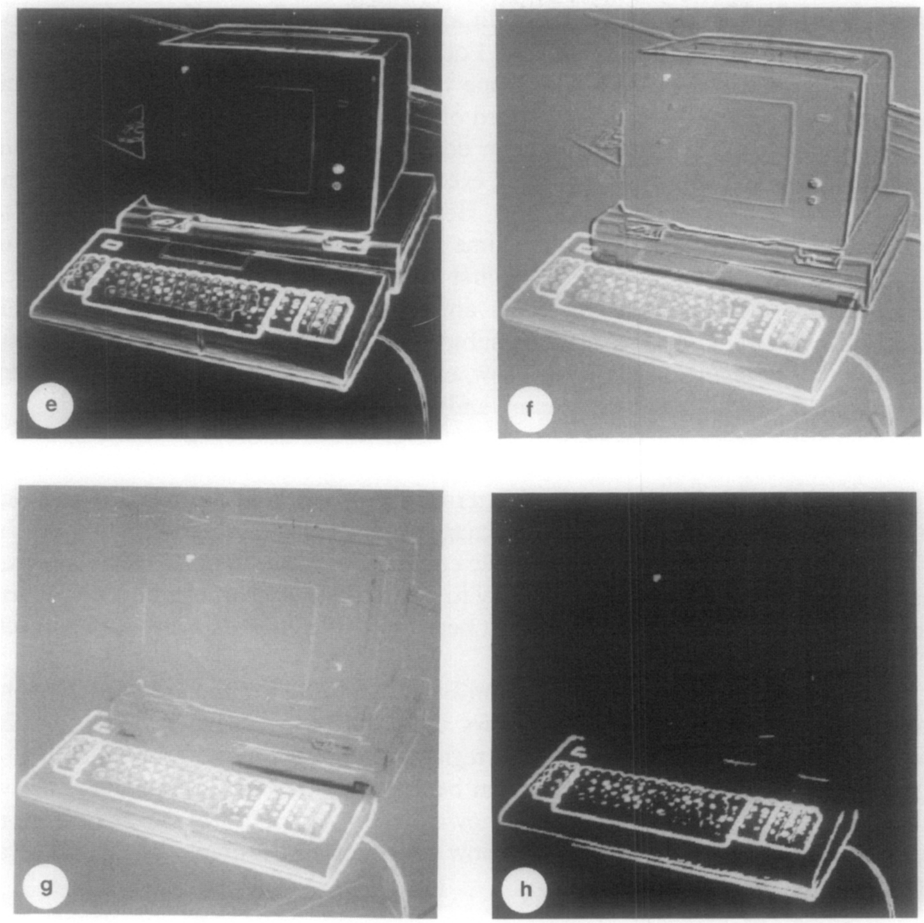

FIG. 2—*Continued.*

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

$$
\begin{aligned}
g(i, j) &= r + (f(i, j) - r)\frac{b(i, j) - d(i, j)}{b(i, j)} \\
&\text{if } f(i, j) > r \text{ and } b(i, j) > d(i, j) \\
&= r - (r - f(i, j))\frac{d(i, j) - b(i, j)}{d(i, j)} \\
&\text{if } f(i, j) < \text{ and } b(i, j) < d(i, j) \\
&= r \\
&\text{otherwise.}
\end{aligned}
\tag{4}
$$

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.

5. EXAMPLE

Figures 2a and b show images of a terminal, one without a keyboard and one with. Figure 2c shows the difference image between Fig. 2a and b. Because the images were not aligned perfectly, Fig. 2c contains registration noise. Adaptive gray scale mapping cannot be applied directly to Fig. 2c because corresponding light and dark regions of registration noise are not near each other. This is because the objects in the image are large.

Figures 2d and e show the images from Figs. 2a and b after the application of the Sobel edge detector, and Fig. 2f shows their difference image. Unlike Fig. 2c, corresponding light and dark regions of registration noise are near each other in Fig. 2f, and adaptive gray scale mapping can be used. The results of applying adaptive gray scale mapping to Fig. 2f are shown in Fig. 2g. Notice the reduction in registration noise. The new edges can now be extracted by simple thresholding, as shown in Fig. 2h.

6. CONCLUSIONS

Difference images are useful in detecting change between two images. One problem with difference images 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 image. We have presented a new algorithm using adaptive gray scale mapping, that reduces registration noise from difference images. It removes noise as well as or better than existing methods, while not blurring the real differences between the images.

ACKNOWLEDGMENT

We would like to thank Dr. Linda L. Brinkley of the Department of Anatomy and Cell Biology, University for Michigan, for suggesting the problem in [3] and providing technical support and advice.

REFERENCES

1. C. K. Chow, S. K. Hilal, and K. E. Niebuhr, X-ray image subtraction by digital means, *IBM J. Res. Dev.* **17**, 1973, 206–218.
2. R. Jain, Extraction of motion information from peripheral processes, *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-3**, 1981, 489–503.

3. T. F. Knoll, L. L. Brinkley, and E. J. Delp, Difference picture algorithms for the analysis of extracellular components of histological images, *J. Histochem. Cytochem* **33**, 1985, 261–267.
4. R. S. Ledley, L. M. Arminski, L. S. Rotolo, T. J. Golab, and C. M. Park, TEXAC-based digital fluoroscopy *in* "IEEE Comput. Soc. Conf. Pattern Recognit. Image Process.," 1982, pp. 530–533.
5. R. L. Lillestrand, Techniques for change detection, *IEEE Trans. Comput.* **21**, 1972, 654–659.
6. C. J. Robinove, P. S. Chavez, D. Gehring, and R. Holmgren, Arid land monitoring using landsat albedo difference images, *Remote Sens. Environ. Sci.* **11**, 1981, 133–156.
7. M. S. Ulstad, An algorithm for estimating small scale differences between two digital images, *Pattern Recognit.* **5**, 1973, 323–333.
8. R. L. Webber and U. E. Ruttimann, Digital subtraction of radiographs as an aid for detecting pathologic changes in periodontal bone, *in* "IEEE 1982 Intl. Work Phys. Eng. Med. Imaging," pp. 130–136.