IMAGE ENHANCEMENT BY CONTROLLED ADJUSTMENT OF LOCAL STATISTICS

REHAUSSEMENT D'IMAGES PAR REGLAGE CONTROLE DES STATISTIQUES LOCALES

Robert C. Vogt, John M. Canning

THOMSON-CGR 48, rue Camille Desmoulins 92130 ISSY LES MOULINEAUX (FRANCE)

RESUME

On décrit deux méthodes de rehaussement de contraste des images radiologiques, l'une linéaire et l'autre non-linéaire, qui utilisent les statistiques à l'intérieur d'une fenêtre glissante locale. Les deux algorithmes ont pour contrainte de produire des images ayant approximativement la même moyenne et le même écart-type que ceux de l'image originale normalisée globalement. Les avantages et les inconvénients de chaque méthode sont discutés, et les résultats sont évalués pour leurs vitesses de calcul et l'effet des paramètres d'entrée incluant la taille de la fenêtre.

SUMMARY

Two methods for contrast enhancement of X-ray images are described, one linear and the other non-linear, which both make use of the statistics inside a sliding local window. Both algorithms are constrained to produce output images which have approximately the same mean and standard deviation as a globally normalized original image. The advantages and disadvantages of each method are discussed, and the results evaluated with respect to speed of computation and the effect of input parameters including window size.

Robert C. Vogt, John M. Canning

IMAGE ENHANCEMENT BY CONTROLLED ADJUSTMENT OF LOCAL STATISTICS

REHAUSSEMENT D'IMAGES PAR REGLAGE CONTROLE DES STATISTIQUES LOCALES

INTRODUCTION

In the treatment of X-Ray images, it is valuable to enhance local variations in contrast in order to make the internal structures of the body more visible. These local variations are related directly to the differences between absorbtion coefficients of different tissues in the body. By enhancing them we can make some of the softer structures visible, and the harder bone structures more recognizable and distinct. A common method of measuring local variations in an image is to compute the statistics within a local window of size N x N, centered on each pixel. Fahnestock and Schowengerdt [1] describe a number of such algorithms. Previous work by our group on this subject was reported in Canning and Benchimol [2]; the work described here represents a number of extensions and improvements to what was presented in that paper.

METHOD

We have developed two algorithms for contrast enhancement based on the use of local windows, one linear and the other non-linear. Each consists of three steps, which for convenience we will call the "global" step, the "regional" step, and the "local" step. The global step is identical for both algorithms, and is essentially a form of pre-treatment. It consists of spreading the pixel distribution of the original image optimally over the available dynamic range, based on the global mean, standard deviation, minimum and maximum of the original. This is a linear transformation which shifts the mean according to the overall shape of the distribution, and which increases the standard deviation so that the new distribution just fills the dynamic range, which in our case runs from 0 to 4095 (12 bits). This step is necessary for the other two steps to have maximum effect. The result of this step is called the "globally normalized" image, or image 'G', and it serves as the starting point for the next two steps.

To perform the regional and local steps we must choose a window size N, where N is odd. We center an NxN window over each pixel in image G, and use it to effectively split G into two images--a "regional" image 'R', and a "local" image 'L'. This is done by taking each pixel of G and dividing it into two parts. The first part (R) is the difference between the local window mean for that pixel and the global mean of image G. The second part (L) is the difference between the value of the pixel and the local window mean for that pixel. Note that in both cases, the result may be negative or positive. Image R thus contains information about the deviation of the window mean from the global mean (the regional deviation), while image L contains information about the deviation of an individual pixel from its regional or window mean (the local deviation). Image L corresponds to the notion of "unsharp masking" as defined in Cocklin and Kaye [3].

Both of our algorithms proceed by first reducing the regional deviations in the regional step, and then increasing the local deviations in the local step; they only differ in the manner by which these transforms are carried out. The linear method uses the raw deviations for these transformations, while the non-linear method uses instead the standard score deviations, as defined below. The final result is then obtained by adding together the transformed images R and L, along with a constant image equal to the global mean of image G.

In order to have result images which can be compared in a controlled manner with what we start, we introduce an additional constraint which forces the final images to have approximately the same global mean and standard deviation as the globally normalized image G. (Another reason for the global pre-treatment step is to insure that we can use as much of the available range as possible, even with this constraint). We enforce the constraint by linking the regional and local steps. That is, the user is allowed to specify a parameter for the degree of regional reduction, and then the program calculates a second parameter which specifies the degree of local enhancement, such that the mean and standard deviation in the final image will be nearly the same as those for image G. This requires keeping some additional statistical information about the images R and L, and then solving a quadratic equation for the desired parameter. The end result is that we enhance image G not by increasing its variance, but by spatially redistributing its variance in a way consistent with its local deviations. This method reduces the possibility of oversaturating the final image, or of failing to make full use of the dynamic range.

The linear and non-linear algorithms differ, theoretically, in how they modify image R and image L. In practice, however, they differ only in the local step--in how they modify image L. If we call G(i,j) a pixel in image G, and G'(i,j) a pixel in the resultant image, then the linear algorithm (after the global step) is expressed as Equation (1) :

$$(1)\quad G'(i,j) = K_L*(G(i,j) - M_W(i,j)) \quad \text{[local part]}$$
$$+ K_R*(M_W(i,j) - M_G) \quad \text{[regional part]}$$
$$+ M_G \quad \text{[global part]}$$

$M_W(i,j)$ is the mean of the NxN window centered at (i,j), and M_G is the global mean of image G. $(G(i,j) - M_W(i,j))$ is the value of image L at (i,j), and $(M_W(i,j) - M_G)$ is the value of image R at the same point. K_R is a constant, less than 1.0, specified by the user, which reduces the regional deviations of image R. K_L is a constant, greater than 1.0, computed by the program, which increases the local deviations of image L. The sum of the first two terms over all pixels must be zero for G' to have exactly the same mean as G. The sum of squares of these terms and their covariance terms over all pixels must equal the sum of squares of all the deviations in G, for G' to have the same variance as G.

The non-linear algorithm is expressed similarly by Equation (2) :

$$(2)\quad G'(i,j) = SD_L*F((G(i,j)-M_W(i,j))/SD_W(i,j)) \quad \text{[local]}$$
$$+ SD_R*(M_W(i,j)-M_G)/SD_G \quad \text{[regional]}$$
$$+ M_G \quad \text{[global]}$$

SD_G is the global standard deviation of image G. $SD_W(i,j)$ is the standard deviation of the NxN window centered at (i,j). $(M_W(i,j)-M_G)/SD_G$ is the standard score of the deviation at point (i,j) of image R, relative to the global standard deviation. The corresponding expression in the local part is the standard score of the deviation at point (i,j) of image L, relative to the standard deviation of the window centered at (i,j). SD_R is a standard deviation, necessarily less than SD_G, entered by the user, which reduces the regional deviations of image R. SD_L is the standard deviation, calculated by the program, that is used to increase the local deviations. It can be easily seen by comparing the two equations that since SD_G and SD_R are constant, with SD_R less than SD_G, the results of the regional steps for both the linear and non-linear approaches are algebraically identical--the only real difference is in the local step.

F is a filtering function that is used to handle the problems associated with uniform regions of the image, where the standard deviation inside the window (SD_W) approaches zero. F returns unchanged any standard score between -6 and +6 SDs. Values outside of +/-6 SD are considered to be noise within a uniform region, and are accordingly set to 0 (the mean).

Finally, for both the linear and non-linear algorithms, values for G'(i,j) which go outside of the available range (0 to 4095) are simply clipped to the minimum or maximum value.

RESULTS

Figures 1 and 2 show two series of images generated by the two algorithms, from a head and chest X-ray, respectively. In both series, image 'a' is the original, and image 'b' is the globally normalized image (G). The pairs following these two show the final result of the linear and non-linear algorithms for windows of different sizes. For each window the linear result is on the left, and the non-linear on the right. The images shown all have 512x512 pixels and are coded on 12 bits. The original images were generated by a CGR CE 10000 computer-aided tomograph using radiographic mode at a resolution of 1024x1024 pixels. Figure 1a was obtained by taking the mean of each 2x2 pixel square in order to get a 512x512 image; figure 2a was obtained instead by taking a 512x512 subsection from the 1024x1024 image.

One advantage of having two closely related linear and non-linear algorithms is that the effect of each of these approaches can be easily compared. In general, the non-linear algorithm is better able to pull out variations which are small in absolute magnitude, but large relative to their surroundings. In the case where such variations represent small or soft body structures, this represents an improvment over the linear method. Unfortunately, this property also means that noise will be amplified, especially with smaller windows and in regions which are relatively smooth such as the background. This effect is quite clear in both series of images.

The advantage of the linear algorithm is that the local variations in contrast are amplified proportionally to their arithmetic values. Weak variations are amplified only weakly. This means first of all that noise is not amplified so strongly, and second of all that the final result is less confusing and easier to interpret than in the nonlinear case. Dense structures show up brighter than soft ones, which corresponds to the expectation of the eye, and so one is more sure of what one is seeing. In addition, the linear algorithm is faster because it doesn't require computation of the window standard deviations, nor special handling when these approach zero. The linear results continue to improve as the window size decreases, while the non-linear results show some degradation when the window size is very small.

A second important effect which can be seen in figures 1 and 2 is that of window size by itself. In general, very large windows have means which are statistically closer to the global image mean, so that the regional step has less variance to reduce, no matter that parameter is used. Our constraint on the final mean and variance restricts the amount of local enhancement which can occur in this case. Smaller windows allow more contrast "depth" to be recovered in the regional step by eliminating more of the regional variance. This recovered depth can be exploited for enhancement by the local step. The result is that we see a steady improvement of local enhancement as window size decreases, down to about 11x11. When the window size goes below this, the images begin to appear grainy because the stability of the local statistics between neighboring windows begins to decline.

An additional effect of decreasing window size concerns the dark "ringing" artifacts which appear at strong edges as a consequence of enhancing the local variations. These rings have a width which is proportional to the window size. They are thus less visible as window size decreases, and this is more pleasing to the eye. In fact, when they are very small they are taken simply to be shadows, and they give more of a bas-relief appearance to the image. An added advantage of the small windows is that their local statistics can be computed faster, which is important since this represents the most time consuming part of both algorithms. Depending on the method used, the computation time varies directly with the diameter of the window or with its area.

The third important determinant of the results of our algorithms is the regional step parameter, K_R or SD_R/SD_G. When this value is close to 1.0, most of the regional mean information is being retained, very little depth is recovered for the local enhancement steps, and so the image doesn't change by very much. As this parameter is reduced, more and more depth is recovered for the local enhancement step, and the resulting images show increased detail. When it is set to zero, all of the regional mean information is removed, and only the local deviations remain. In this case the final image appears overly different from the original, which is not desireable. Accordingly, we keep a small percentage of the regional mean information; in all of the images shown here, the value of K_R and SD_R/SD_G is 0.10. We have not illustrated the effect of this parameter due to space considerations, and because it is difficult to distinguish different values of K_R when it is less than 0.50.

CONCLUSION

We have described two closely related algorithms, one linear and the other non-linear, for the enhancement of images by the use of statistics inside a sliding window. The non-linear one was found to be more sensitive, and is therefore able to extract smaller structures, but it is also more susceptible to noise. The linear method gives an image which is easier for the eye to interpret, and which doesn't degrade as quickly due to noise when the window size is very small. In general, smaller windows gave better enhanced images, down to a size of around 11x11, at which point the images begin to appear grainy. We suggest that the best way to use both algorithms is in conjunction with each other, since each one gives slighly different information. Keeping window size fixed, the non-linear algorithm is best able to make small structures visible, while the linear algorithm gives an image which is cleaner, showing the relative densities of structures more clearly, and in which the effects of noise can be more easily separated from the structures in the object.

REFERENCES

[1] J.D. Fahnestock and R.A. Schowengerdt "Spatially variant contrast enhancement using local range modification" OPTICAL ENGINEERING, May/June 1983,Vol.22,N°3

[2] J.M. Canning and C. Benchimol "X-Ray image enhancement using local statistical information" Proceedings--9th GRETSI Conference on Signal Processing and its Applications, May 1983, p.623-628

[3] M. Cocklin and G. Kaye "Image enhancement of chest radiographs using local statistical information" Proceedings--ISMII'82, IEEE Computer Society Press, 1982, p. 82-85.

Figure 1a
ORIGINAL IMAGE

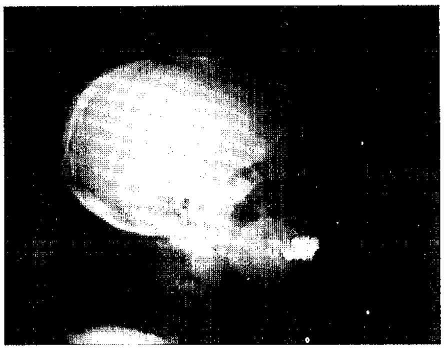

Figure 1b
GLOBALLY NORMALIZED IMAGE (G)

LINEAR RESULTS

NON-LINEAR RESULTS

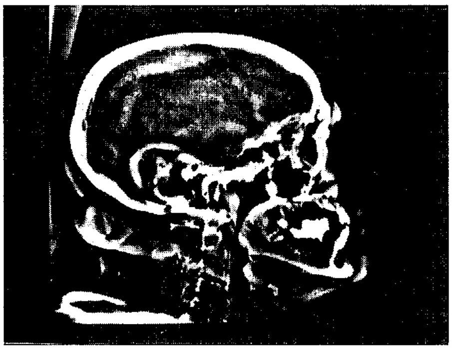

Figure 1c
41 X 41 WINDOW

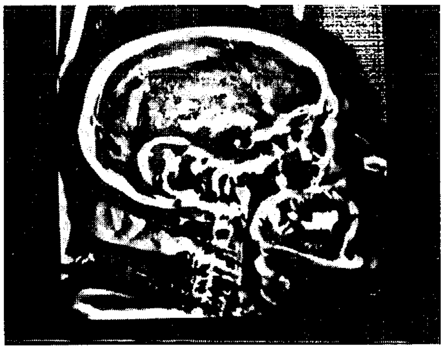

Figure 1d
41 X 41 WINDOW

Figure 1e
21 X 21 WINDOW

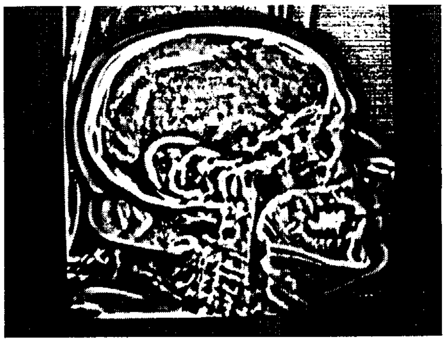

Figure 1f
21 X 21 WINDOW

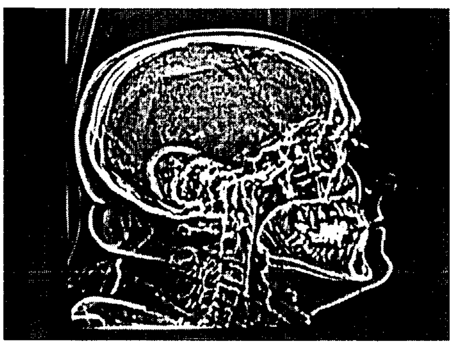

Figure 1g
11 X 11 WINDOW

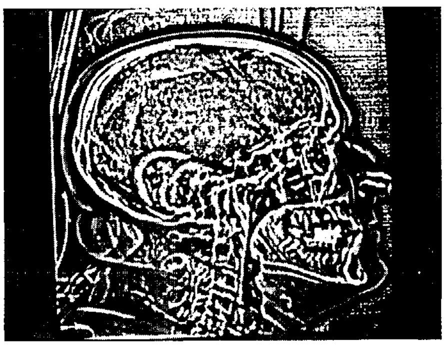

Figure 1h
11 X 11 WINDOW

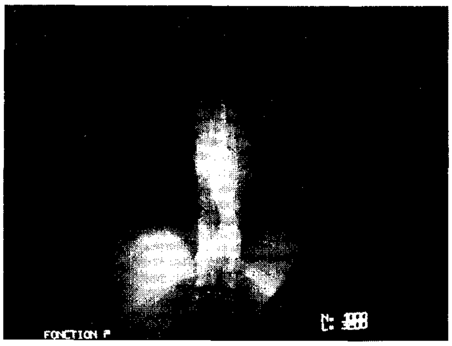

Figure 2a
ORIGINAL IMAGE

Figure 2b
GLOBALLY NORMALIZED IMAGE (G)

LINEAR RESULTS

NON-LINEAR RESULTS

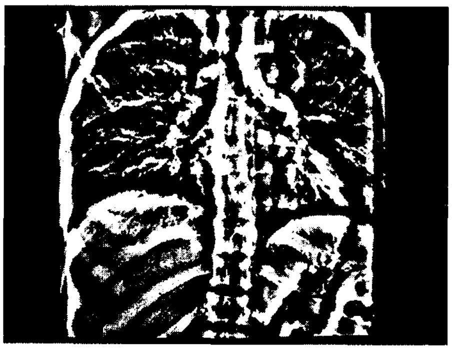

Figure 2c
41 X 41 WINDOW

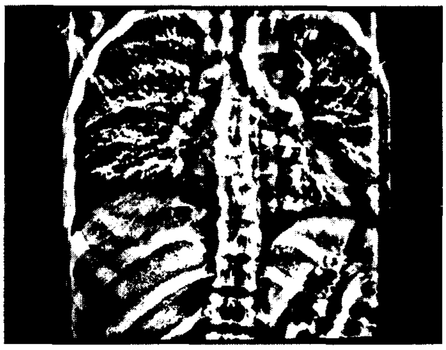

Figure 2d
41 X 41 WINDOW

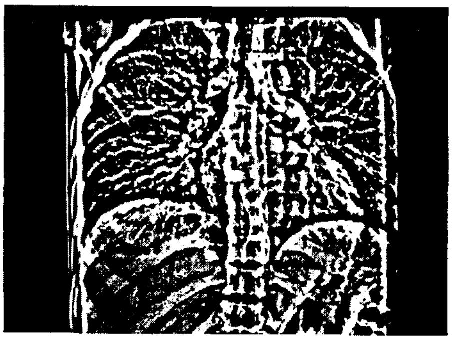

Figure 2e
21 X 21 WINDOW

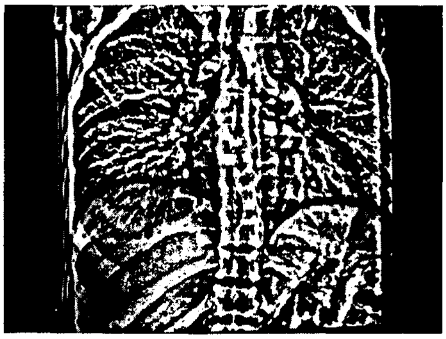

Figure 2f
21 X 21 WINDOW

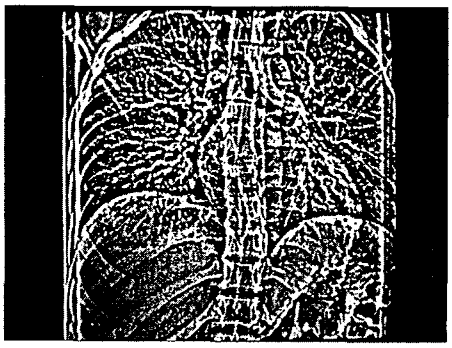

Figure 2g
11 X 11 WINDOW

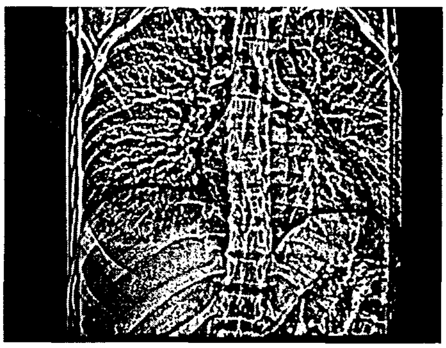

Figure 2h
11 X 11 WINDOW