

**Adaptive Histogram  
Equalization for Automatic Contrast  
Enhancement of Medical Images**

*Technical Report 86-010*

*February, 1986*

*Stephen M. Pizer, John D.  
Austin, John R. Perry, and Hal D. Safrit*

**The University of North Carolina at Chapel Hill  
Department of Computer Science  
New West Hall 035 A  
Chapel Hill, N.C. 27514**



*To appear in Proceedings of the 1986 SPIE Medicine XIV/PACS IV Conference,  
Newport Beach, CA*

# Adaptive histogram equalization for automatic contrast enhancement of medical images

Stephen M. Pizer<sup>\*+</sup>, John D. Austin<sup>\*</sup>, John R. Perry<sup>+</sup>, Hal D. Safran<sup>+</sup>,  
John B. Zimmermann<sup>†</sup>

Department of Computer Science<sup>\*</sup> and Department of Radiology<sup>+</sup>  
University of North Carolina  
Chapel Hill, North Carolina  
27514

Department of Computer Science<sup>†</sup>  
Washington University  
St. Louis, Missouri  
63130

## Abstract

With the large number of images that will be viewed simultaneously in a medical picture archiving and communication system (PACS) system in the diagnosis of a particular patient, image by image interactive contrast enhancement, at present by intensity windowing, becomes unacceptably time-consuming. Furthermore, windowing has disadvantages of being non-reproducible and providing adequate contrast primarily in selected image regions. The method of adaptive histogram equalization (ahe) appears to provide a solution to these problems. It is reproducible, automatic, and simultaneously provides contrast in all image regions.

After summarizing the basic method, this paper will 1) describe a new contrast limited form of ahe that appears to allow its application to a wide variety of medical images, 2) present a VLSI machine design that will allow the calculation of ahe in a fraction of a second per megapixel, and 3) report the results of a study demonstrating that for chest CT images, ahe provides no measurable loss of diagnostic performance compared to the now standard windowing.

### 1. Clipped adaptive histogram equalization

Adaptive histogram equalization (ahe) is a contrast enhancement method that has produced excellent results in medical imaging. Both controlled studies and clinical use suggest that over a wide range of medical image types all contrast available in the image data can be perceived in the single image that is its automatically produced result (see figure 1). As a result it is a strong candidate to be part of a general medical image display system, such as would be part of a PACS system, since time-consuming windowing of large numbers of images could be avoided. Without such a contrast enhancement technique each medical image to be displayed simultaneously, presumably from a range of medical imaging modalities, would have to be separately windowed by interactive means as a step preliminary to image storage. Alternatively, the inconvenience of interactive on-line windowing of each of many images on the screen would have to be incurred for each presentation, and the cost of providing the capability of applying a different windowing for each of a number of portions of the display screen would have to be incurred.

The basic form of ahe was invented independently by Ketcham [1976], Hummel [1977], and Pizer [1981]. In this basic form the method involves applying to each pixel the histogram equalization mapping based on the pixels in a region surrounding that pixel (its *contextual region* - see figure 2). That is, each pixel is mapped to an intensity proportional to its rank in the pixels surrounding it.









Figure 1. Windowing vs. clipped ahe on various medical images: a) chest CT, b) abdomen CT, c) spine CT.







Figure 1 cont'd. d) renal DSA, e) digitized film angiogram of brain, f) digitized portal film of head and neck.




Figure 2. A pixel(\*) and its contextual region

Experience has shown that the size of the contextual region can be set to either 1/16 or 1/64 of the image area for essentially all medical images, with the smaller region chosen only when the feature size of interest is quite small. With a smaller contextual region the contrast becomes too sensitive to very local variations and in particular to image noise. The oversensitivity to local variations can cause artifacts, which have never been experienced with the preferred contextual region sizes.

Although ahe (as described so far) frequently produces excellent results, in certain cases noise becomes disturbingly obvious. In particular, this occurs where the image includes relatively homogeneous regions or a poor signal to noise ratio (see figure 3). The large noise enhancement is disturbing despite the fact that signal contrast is enhanced proportionately. The following modification of the method, called *clipped* or *contrast-limited ahe (clahe)*, avoids this overenhancement of noise.

Contrast enhancement can be defined as the slope of the function mapping input intensity to output intensity (see figure 4). We will assume that the range of input and output intensities are the same. Then a slope of 1 involves no enhancement, and higher slopes give increasingly higher enhancement. Thus the limitation of contrast enhancement can be taken to involve restricting the slope of the mapping function.

With histogram equalization the mapping function  $m(i)$  is proportional to the cumulative histogram:

$$m(i) = (Display\_Range) * (Cumulative\_Histogram(i) / Region\_Size).$$

Since the derivative of the cumulative histogram is the histogram, the slope of the mapping function at any input intensity, i.e. the contrast enhancement, is proportional to the height of the histogram at that intensity:

$$dm/di = (Display\_Range / Region\_Size) * histogram(i).$$

Therefore, limiting the slope of the mapping function is equivalent to clipping the height of the histogram.




Figure 3. (a) Unclipped vs. (b) clipped ahe on a CT scan of the brain

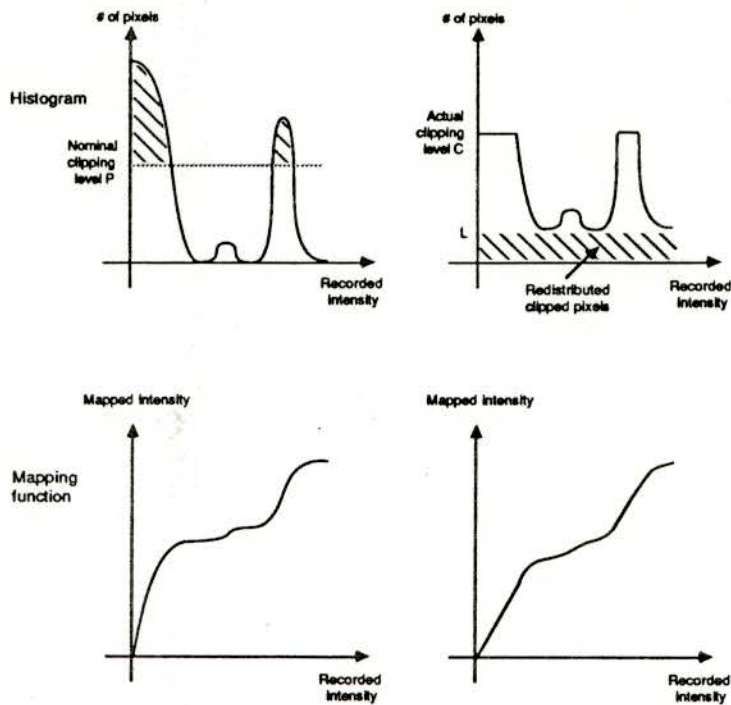


Figure 4. Contrast mapping functions and their associated histograms: (a) original histogram, and (b) clipped histogram

High peaks in the histogram are normally caused by nearly uniform regions. In such a case, with the mapping due to ordinary histogram equalization a limited range of input intensity values is mapped to a wide range of output intensity values, perhaps overenhancing the noise. But enforcing a maximum on the contents of the histogram bins will limit the amount of contrast enhancement and thus the enhancement of noise.

When contrast enhancement is reduced at one location it must be increased in other areas so that the entire input intensity range will be mapped to the entire output intensity range. This corresponds to renormalizing the histogram after clipping so that its area returns to its original value. We think of this as redistributing the clipped

pixels. The most effective way that we have found to redistribute these pixels is to distribute them uniformly in all histogram bins. This can be thought of as adding the contrast mapping due to the clipped histogram to a linear mapping that achieves just the height at the maximum image intensity such that the height of the sum is equal to the intensity range in the original image (see figure 4).

Incorporation of histogram clipping into the existing clahe algorithm is straightforward. After the computation of each histogram, that histogram is modified by clipping. Then the intensity of each displayed pixel is proportional to its rank in the clipped form of the histogram of the contextual region of the pixel.

The clipping limit  $C$  is determined by specifying the limiting slope  $S$ .  $C$  can be shown to be  $S$  times the average histogram bin contents. This relation can be made intuitive by noting that a slope of 1 corresponds to all bins having the same (average) value and that slope is proportional to the value of a histogram bin.

Adding a uniform level  $L$  to the clipped histogram will push the clipped histogram again above the clipping limit, so the original histogram needed to be clipped at a lower limit  $P$  such that  $P + L(P)$  is equal to the clipping limit ( $L$  is written as a function of  $P$  because it depends on  $P$ ). The value of  $P$  that satisfies this equation can be found by a binary search for the largest integer intensity such that the area of the histogram above this level is not greater than the product of the total number of histogram bins and  $C$  minus this level.

Having found the desired  $P$  and the related  $L$ , equal to the clipping limit minus  $P$ , the modified histogram value  $v$  in any bin is calculated from original value  $v_{orig}$  by

$$v = \begin{cases} v_{orig} + L & \text{if } v_{orig} < P \\ C & \text{if } v_{orig} \geq P \end{cases}$$

## 2. A hardware design to speed up clahe

Clahe as described so far is slow, requiring many tens of minutes for a 512 x 512 image on an ordinary superminicomputer. If the method is to be part of a general medical image display system, it must be applicable in a few seconds or less, and it must be applicable to images of as large as 4K x 4K. Therefore, various algorithms for the method or approximations have been developed that increase its speed on various processors [Pizer, 1986]. For a general-purpose processor the fastest method known [Pizer, 1984; 1986] involves computing the mapping at sample pixels and interpolating the mapping elsewhere. That is, for any pixel  $P$  each of the mappings associated with nearby sample pixels are applied to the intensity at  $P$ , and interpolation between these mapped values are made using interpolative (e.g. bilinear) weights depending on the relative position of  $P$  and the sample pixels used. The contextual region used in this interpolative form is chosen so that the total number of pixels affecting the final mapping of any pixel is the same as the desired contextual region for the original, uninterpolated form of clahe.

A version of the algorithm exists (see the algorithm for feedback processors in [Pizer, 1986]) which allows the reasonably efficient application of this interpolative method to images 32 times larger than would be possible if it were necessary to provide space for the whole of the input and output image data. However, even these improvements speed the method only to approximately 10 minutes per megapixel on an ordinary superminicomputer. While application of pipelined processors could speed this up

by an order of magnitude, the objective of a few seconds per megapixel necessary for routine clinical implementation would not be reached. As a result we have designed the architecture of a VLSI-based machine calculated to apply ahe or other image processing operations such as arbitrary convolutions in a fraction of a second per megapixel. The machine is called *MAHEM* – Multiprocessor Adaptive Histogram Equalization Machine.

*MAHEM*'s architecture involves a small processor at each pixel and the ability to broadcast a value to all pixels simultaneously [Austin, 1985]. It should cost no more than the feedback architectures, and at the same time compete equally or favorably with them on a wide range of other image restoration or enhancement operations. However, this design must be extended or modified to do clipped ahe and to handle larger images than can be expected to fit in the local memory of the device.

*MAHEM* consists of a control unit and a pixel store with processors (see figure 5). The control unit interfaces to the host computer, broadcasts data and instructions to the pixel-processors, and receives data from a selected pixel-processor. Each pixel contains memory to hold a small number of initial, temporary, and intermediate or final computed values at that pixel, as well as a 1-bit adder and comparator to zero. It also holds two 1-bit enabling tokens, a transmit-enable token to allow the pixel data to be transmitted to the control unit and a receive-enable token to allow the processor to receive pixel data or execute instructions from the control unit. Either of the enabling tokens can be passed to its right row-wise neighbor or set by a computation of the pixel's processor.

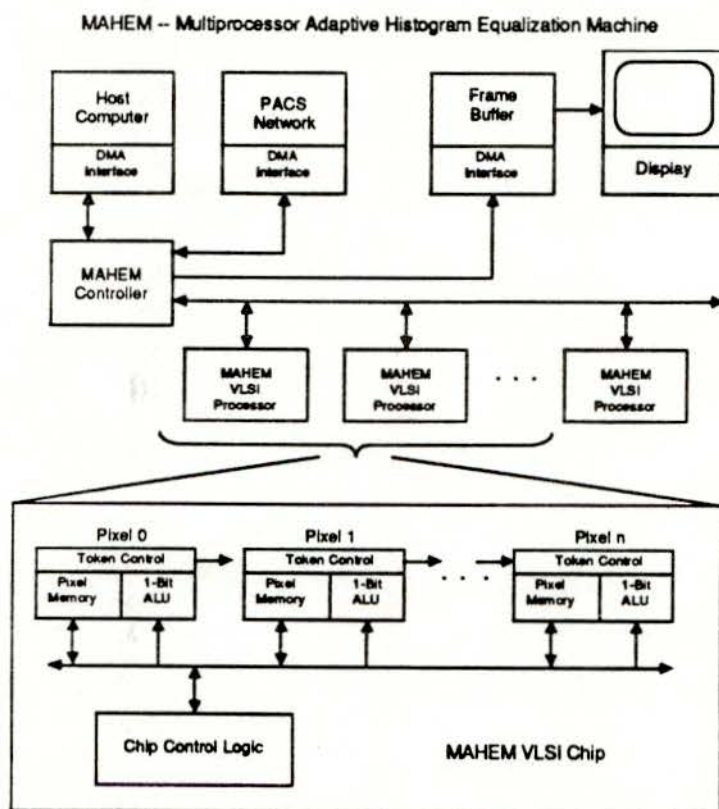


Figure 5. MAHEM system and pixel architecture



All pixels have the ability to communicate data along a common bus to/from the control unit. The basic means of operation is that a pixel with the transmit-enable token sends its initial (or some intermediate) value to the control unit, and from there it is broadcast to all pixels. Those with their receive-enable token set to the on state compute in parallel a new intermediate value based on this "affecting" pixel value that has been broadcast. The receive-enable tokens are shifted between pixel broadcast cycles to allow the receiving region to be related in a fixed way with the location of the broadcasting pixel.

Ahe is computed by this engine as follows. Every pixel holds a result which is its rank in its contextual region, with this value initially set to zero. When a pixel is broadcast, all pixels in whose contextual region it is are enabled, and each compares the broadcast intensity to its own using its own processor in a bit-serial fashion. If the broadcast intensity is less than the stored intensity, the rank field in the pixel's memory is incremented by one, using its own processor. The receive-enable tokens for the affecting region as well as the send-enable token are then shifted one position (with special procedures for image edges), and the cycle is repeated. When all pixels have had their intensity broadcast, the rank fields hold the ahe result.

Clahe can be computed by MAHEM if pixels are able to produce an increment of less than a full integer when some fraction of the pixels at that intensity are to be effectively clipped away. We can approximate the effect by using the broadcast mechanism to allow each pixel to count the number of its neighbors in a surrounding region that have the same intensity, and then use the ratio between that number and the clipping level to determine the incremental fraction that will need to be broadcast with the pixel intensity during the rank calculation process. The approximation arises because not all of these will be in the affecting region of any given result pixel. Nevertheless, it appears clear that such an approximation could give the major effect of clahe. We must carry out simulations to see whether this is the case, and also develop and evaluate other possibilities for implementing clahe or an approximation to it.

MAHEM can compute other image enhancement and restoration results such as filtering, by using its enabling mechanism to specify affecting regions of convolution kernels or their nonlinear generalizations, such as median window filters. Multiple passes with different increments per pass allow variable weighting as a function of distance to take place. Algebraic and logical operations on multiple images are also possible. These possibilities as well as system design issues that follow from the various possibilities must be considered and evaluated via simulation.

### 3. Evaluation of adaptive histogram equalization

A Receiver Operating Characteristic (ROC) rating experiment was used to compare ordinary (unclipped, interpolative) ahe to interactive windowing as means of contrast enhancement [Zimmerman, 1985]. A set of test images was prepared from normal computed tomography (CT) images of the chest. In each normal image, four sites were chosen for the insertion of simulated Gaussian lesions, two sites in the lungs and two in the mediastinum. For each field (lungs and mediastinum), three simulated lesions were prepared, the linear size being chosen to be appropriate for the given field. The intensity of the lesions was chosen so that when a lesion was inserted into an image it was very subtle. From each normal image, twelve additional images were generated, each image having one of the lesions inserted at one of the selected sites. From each of these images, another image was prepared by processing the given image with ahe. The complete test set consisted of the normal images, the images with the lesions inserted, the normal

images processed with ahe, and the images with the lesions inserted processed with ahe. In order that there be equal numbers of images with and without the artificial lesions, multiple copies of the images without lesions inserted were included. The test set was then presented in random order to three trained radiologists.

For those images not processed with ahe, the observers were asked to perform interactive linear min-max windowing on the appropriate image field (lungs or mediastinum). After they had completed the windowing task, they were presented with a replica of the lesion and the exact location of the prospective lesion was indicated with a removable crosshairs. For those images processed with ahe, the observers were shown the lesion replica and the crosshairs immediately. In all cases, they were asked to rate their confidence that a lesion of the type depicted was present in at the indicated site.

The data from the observer studies was analyzed using the CORROC program developed by Metz and his collaborators for correlated ROC data. Observer performance was computed separately for each site in the image and for each simulated lesion, and when the results for the sites within a field showed no significant difference, then pooled across these sites. The areas under the ROC curves and their standard deviations were calculated and used to compare the two modalities for a given site and lesion; the modalities were assumed to have no difference in their diagnostic power if the difference in the areas of the two ROC curves was less than a preset criterion. The results were evaluated for differences of 1.5 and 2.5 standard deviations.

After pooling across sites within each field there were results for three lesion types times two fields for each of the three observers, with results as follows. In the lungs there was no significant difference in observer performance between the two modalities for the more stringent criterion and one such difference in favor of each of the two contrast enhancement methods by one observer for the very lax criterion. It was found that the observers performed slightly better using windowing than with ahe in the mediastinum field. In one of the nine classes there was significant advantage to windowing according to the more stringent criterion, and in four of the remaining classes there was an advantage to windowing according to the lax criterion. We suggest (and the observers reported) that in the relatively homogeneous mediastinum field overenhancement of the noise was a problem. We believe on the basis of anecdotal experience that clipping, invented after the study was in progress, would remove whatever small advantage to windowing might have been suggested by the study, though a controlled study has not yet been done to demonstrate this hypothesis. We are just beginning such a study, with clinical images.

Two other results from the study were of interest. First, little difference was seen in performances of the individual observers. Second, a time series analysis of the data done to determine if the observer performances improved over the course of the experiment showed no such improvement.

The above results on the effectiveness of ahe have been confirmed in a smaller study with prewindowed abdomen CT images by ter Haar Romeny et al [1985]. An ongoing study at the Academic Hospital of Utrecht, Netherlands, as well as clinical experience, suggests that clipped ahe is even more successful.

The results suggest that ahe will be a most useful means of contrast enhancement in clinical studies. It has the important advantages over windowing that it is fast and reproducible, requires no manual intervention, and is able to depict contrast in many areas of the image simultaneously. And even in the case of far lower contrast lesions than are clinically significant, there is little or no significant advantage in diagnostic

performance using *interactive* windowing, and likely no advantage when the clipped form is used or the clinically common preset windowing is used.

The clipped ahe technique has been applied to a wide variety of medical images, including CT scans, MRI images, digital radiographs, digitized x-ray angiograms, digital subtraction angiograms, digital radiotherapy portal films, and scintigrams. The results (e.g. see figure 1) show that the method not only provides contrast in all organs simultaneously, but also the ability to lessen the effects of nonuniform sensitivity, as in MRI images from surface coils. These results to date have been very supportive of the hypothesis that this method could be used as a standard display technique in medical imaging. It is clear that a different clipping level must be chosen for different studies, but it appears that this level can be fixed for any study type (e.g. CT scan of the abdomen with injected contrast) and thus the method can be applied without human intervention.

### Acknowledgements

We thank E. Philip Amburn, Ari Geselowitz, Bart ter Haar Romeny, and Karel Zuiderveld for help in the development of the clipped form of ahe. We thank the latter two and Paul van Waes, Julian G. Rosenman, Edward V. Staab, Bradley Brenton, and William McCartney for their collaboration on the evaluation of ahe. We are indebted to Joan Savrock for manuscript preparation, to Bo Strain and Karen Curran for photography, and to Robert Cromartie for computer assistance.

### References

Austin, J. D. and S. M. Pizer, "An Architecture for Fast Image Processing", Technical Report, Dept. of Computer Science, University of North Carolina, 1986.

Hummel, R. A., "Image Enhancement by Histogram Transformation", *Computer Graphics and Image Processing* 6, pp. 184-195, 1977.

Ketcham, D. J., R. W. Lowe, and J. W. Weber, "Real-time Image Enhancement Techniques", *Seminar on Image Processing*, Pacific Grove, California, Hughes Aircraft Company, pp. 1-6, 1976.

Pizer, S. M., J. B. Zimmerman, and E. V. Staab, "Adaptive Grey Level Assignment in CT Scan Display", *J. Comp. Ass. Tomo.* 8(2), pp. 300-305, 1984.

Pizer, S. M., "Intensity Mappings for the Display of Medical Images", *Functional Mapping of Organ Systems and Other Computer Topics*, pp. 205-217, Society of Nuclear Medicine, 1981.

Pizer, S. M., B. M. ter Haar Romeny, K. Zuiderveld, J. B. Zimmerman, P. Amburn, A. Geselowitz, P. F. G. M. van Waes, and A. de Goffau, *Recent Developments in Adaptive Histogram Equalization*, Scientific Exhibit, 71st Scientific Assembly and Annual Meeting - Radiological Society of North America, Chicago, Illinois, November 17-22, 1985.

Pizer, S. M., E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. M. ter Haar Romeny, J. B. Zimmerman and K. Zuiderveld, "Adaptive Histogram Equalization and Its Variations", Technical Report, Dept. of Computer Science, University of North Carolina; also submitted for publication to *Computer Vision, Graphics, and Image Processing*, 1986.

Zimmerman, J. B., "Effectiveness of Adaptive Contrast Enhancement", Ph.D. Dissertation, Department of Computer Science, University of North Carolina, 1985.