

Image Processing in Digital Radiography: Basic Concepts and Applications

Gary T. Barnes and Karen Lauro

Digital x-ray images are routinely processed to enhance diagnostic information and to suppress irrelevant detail, and also to extract quantitative information. The basic concepts and terminology of image processing as it applies to x-ray projection radiography are discussed and defined. In general, the processing of an image involves one or more point, local, and/or global operations. Clinical examples of linear and nonlinear gray-scale and algebraic point operations are presented. Examples are also given of local operations. Included in the latter group are distortion corrections, misregistration corrections, linear filtering, and nonlinear filtering.

© 1989 by W.B. Saunders Company.

A DIGITAL IMAGE IS A two-dimensional array of quantized gray-scale values or pixels

$$I(m, n) \quad (1)$$

where m represents the pixel column location (x direction) and n represents the pixel row location (y direction). For an $M \times N$ image matrix or array m can take on values $1, 2, \dots, M$; n can take on values $1, 2, \dots, N$; and for a K -bit deep-image matrix, individual pixels can have integer (gray-scale) values from 0 to $2^K - 1$. For example, a computed tomography (CT) image with a 512^2 matrix will have 512 columns and rows, ie, 512 pixels across in the x direction and 512 pixels vertically in the y direction. If the matrix or pixel depth is 12 bits, the image can have up to 2^{12} or 4096 shades of gray or gray scale values. Likewise, an 8-bit image has 2^8 or 256 gray-scale values, and a 10-bit image, 2^{10} or 1,024 gray-scale values.

Digital image processing involves manipulating one or more digital images. The subject covers a broad range of techniques, and medical imaging applications generally involve the en-

hancement of diagnostically important information (and the suppression of nondiagnostic information). This report reviews basic image processing concepts and their use in medical x-ray imaging.

Operations of practical importance on a single image involve artifact suppression, gray-scale manipulations, distortion corrections, and edge enhancement. Important techniques involving two or more images are (time difference) subtraction, dual energy cancellation of bone or soft tissue structures, and the extraction of flow or organ function parameters. The processing of a single image can be classified as either a point operation, a local operation, or a global operation. A point operation uses a single point (pixel) of the initial or input image to obtain the corresponding point of the final or output image. A local or neighborhood operation uses several pixels in a limited area of the input image to obtain a point in the processed or final image. Likewise, the less common global operation uses information from essentially all, or almost all, of the input image to obtain a pixel in the output image. Frame processing techniques, ie, algorithms involving two or more images, can be similarly classified.

POINT OPERATIONS 1: GRAY-SCALE MANIPULATIONS

Gray-Scale Processing

An important class of image manipulation involves changing the gray-scale values to better display diagnostically important information. This is analogous to manipulating the sensitometric curve of conventional x-ray film to optimize radiographic information content. However, it is a much easier process with a digital image and is also used routinely to suppress artifacts.

Gray-scale manipulations are more generally referred to as point operations, because each pixel value in the input image is mapped into the output image in such a way that each output pixel's value depends only on the value of the corresponding input pixel. Such gray-scale manipulations make a single input image into a single

From the Department of Radiology, University of Alabama Hospitals, University of Alabama at Birmingham, and CT Engineering Group, Picker International, Cleveland, OH.

Address reprint requests to Gary T. Barnes, PhD, Department of Radiology, University of Alabama Hospitals, University of Alabama at Birmingham, Birmingham, AL 35233.

© 1989 by W.B. Saunders Company.

0897/1889/89/0203-0006\$03.00/0

output image. A given point operation is completely specified by a functional relationship that determines the mapping of the input gray levels to the output image gray levels. A look-up table is often used for performing repetitive point operations. Such a table consists of an ordered pair of input and output pixel values, incorporating the functional relationship of the point operation. Thus, rather than calculating the output pixel value from the input pixel value, the output value is retrieved from memory. When one is dealing with large and even moderate size images, this can save a significant amount of time.

The Gray-Level Histogram

Before discussing gray-scale manipulations, it is useful to define the gray-level histogram of an image and to review its properties. The histogram of an image (Fig 1) is a plot of the frequency distribution of the pixel gray levels, ie, the number of pixels with a given gray level versus gray level. It is a useful tool in image gray-scale processing and can be used to optimize the information displayed. It can also be used to determine if the digitization process is linear.

Windowing and Leveling

The linear point operations of windowing and leveling were introduced to medical imaging by EMI Ltd (Hayes, Middlesex, England) in 1972

on their Mark I CT scanner and are used to obtain virtually every CT and digital subtraction angiography (DSA) image that is viewed or filmed. Mathematically, these operations can be characterized by

$$I_{out}(m, n) = c \cdot I_{in}(m, n) + d \quad (2)$$

where $I_{in}(m, n)$ is the input image, $I_{out}(m, n)$ is the output or display image, and the constants c and d are determined by the window and level display controls and the number of output or display image bits. When $c = 1$ and $d = 0$ (Fig 2, line C), the input and output images are identical. Keeping the level the same and decreasing the window increases image contrast and decreases the range of image gray-scale values displayed (Fig 2, line B). Changing the level at the same window setting (Fig 2, line A) selects a different subset of image gray-scale values to be converted to the display gray scale. In practice, gray-scale values that exceed the maximum display gray-scale value are set equal to the maximum value. Likewise, converted image values that are less than the minimum display pixel value are set equal to the minimum value.

The effect of changing the window and level values on the chest image in Fig 1 are shown in Fig 3. To obtain the image in Fig 3A, the window spanned the range of image gray-scale values, and the level was set at the center of the gray-scale range. This corresponds to the output

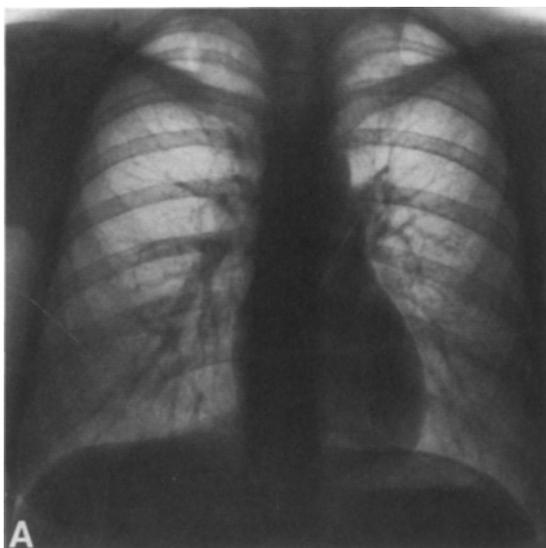


Fig 1. Unprocessed digital chest image (A) and its histogram (B). The image was obtained with the Picker-prototype scanning-slit dual-energy chest unit (Picker International, Inc, Cleveland, OH).

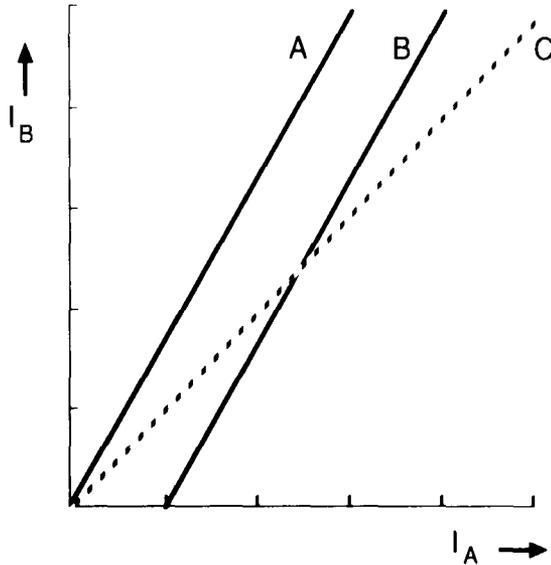


Fig 2. The linear point operations of windowing and leveling.

image equalling the input image (ie, $c = 1$ and $d = 0$; equation 2 and Fig 2, line C). In Fig 3B the window is reduced by a factor ≈ 2.5 to a clinically useful setting and the level is unchanged (Fig 2, line B). In Figure 3C, the level is also reduced to a clinically useful setting (line A in Figure 2).

Nonlinear Gray-Scale Mapping

Often to enhance diagnostically important information it is desirable to use a nonlinear gray scale transformation. As illustrated in Fig 4, a nonlinear operation modifies the shape of the gray-scale histogram in a predictable manner, and can be characterized by the equation

$$N_B \cdot \Delta I_B = N_A \cdot \Delta I_A \quad (3)$$

where N_A is the number of pixels in the input image with gray scale values between I_A and $I_A + \Delta I_A$. N_B and ΔI_B are similarly defined.

One common operation of this type involves flattening or equalizing the histogram. Assuming N_{ave} equals the average number of pixels per unit gray level, the operation that accomplishes this follows.

$$\Delta I_B = (N_A/N_{ave}) \cdot \Delta I_A \quad (4)$$

Figure 5 shows an example of histogram equalization applied to the chest image of Fig 1. In comparing the two images (Figs 1 and 5), it is apparent that a wider range of anatomy (ie, mediastinum as well as lung) is displayed with

the histogram-equalized image. However, low-contrast lung detail is less apparent. Because this information is important clinically, histogram equalization is not advantageous to chest images. It has proven useful in situations where one does not have small, low-contrast detail of clinical interest, such as in CT of the mediastinum.

Offset and Gain Corrections

Often in acquiring a digital image, there are systematic errors introduced that create noticeable and disturbing artifacts. Such artifacts depend on the method of image acquisition; the fact that they are systematic allows one to process and improve the image. For example, in a CT scanner or a slit radiographic unit where one has an array of detector elements, significant differences in leakage (or dark current) and gain occur between elements. In a CT scanner these corrections are done routinely on every scan and are transparent to the user. If, however, the CT unit is out of calibration and the corrections made differ significantly from the ones that should be made, the quality of the resultant image will suffer.

For a slit radiographic device in which a given row is obtained by sampling a given detector element during a scan, the correction takes the form

$$I_{out}(m, n) = G(n) \cdot [I_{in}(m, n) - L(n)] \quad (5)$$

where $L(n)$ is the leakage or dark current offset vector, $G(n)$ is the gain vector and $I_{out}(m, n)$ is the corrected image. Typically, the dark current offset vector is determined just before an x-ray exposure and the gain vector by a calibration scan of air or a known thickness of tissue substitute. Figure 6 shows an example of the chest image in Fig 1, which was acquired with a prototype scanning slit unit and not corrected for detector-to-detector offset and gain variations. The difference in the two images is striking and illustrates the fundamental importance of this type of processing.

POINT OPERATIONS 2: ALGEBRAIC MANIPULATIONS

Addition and Subtraction

Addition and subtraction are important point operations in medical imaging in general and are used routinely in DSA. Two or more images (ie,

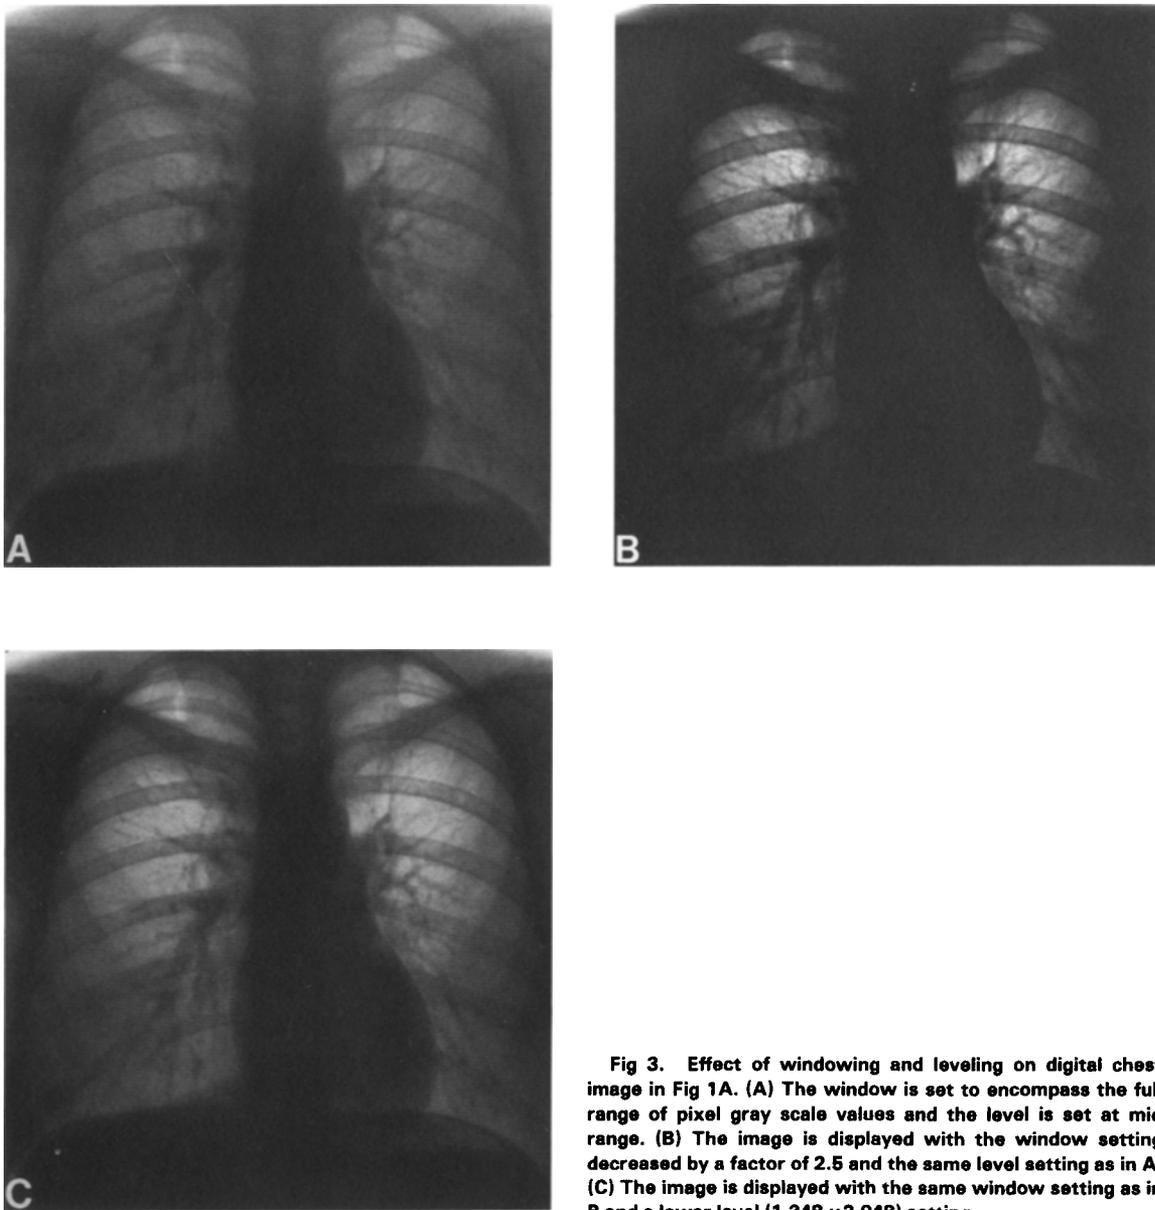


Fig 3. Effect of windowing and leveling on digital chest image in Fig 1A. (A) The window is set to encompass the full range of pixel gray scale values and the level is set at mid range. (B) The image is displayed with the window setting decreased by a factor of 2.5 and the same level setting as in A. (C) The image is displayed with the same window setting as in B and a lower level (1,348 v 2,048) setting.

frame operations) are involved, and the number of pixels in the final image is the sum or difference of the corresponding pixels of the input images.

$$I_{\text{out}}(m, n) = I_1(m, n) + I_2(m, n) \quad (6)$$

$$I_{\text{out}}(m, n) = I_1(m, n) - I_2(m, n) \quad (7)$$

A common example of addition is the averaging of rapidly acquired images in digital subtraction angiography to reduce noise. Provided patient motion is not a problem, the average of N images

$$I_{\text{out}}(m, n) = N^{-1} \sum_{i=1}^N I_i(m, n) \quad (8)$$

results in an image in which the noise level is reduced by a factor of \sqrt{N} compared with one of the initial images. This occurs because the noise fluctuations vary randomly from image to image and add incoherently (ie, tend to cancel out), whereas the anatomical structures of interest are the same in each image of the sequence and add coherently. An example of averaging is shown in Fig 7.

Subtraction is routinely used in digital fluorographic (or fluoroscopic) procedures to obtain an image of vascular structures. Typically, several frames are acquired before and after the injection of contrast media. The images acquired

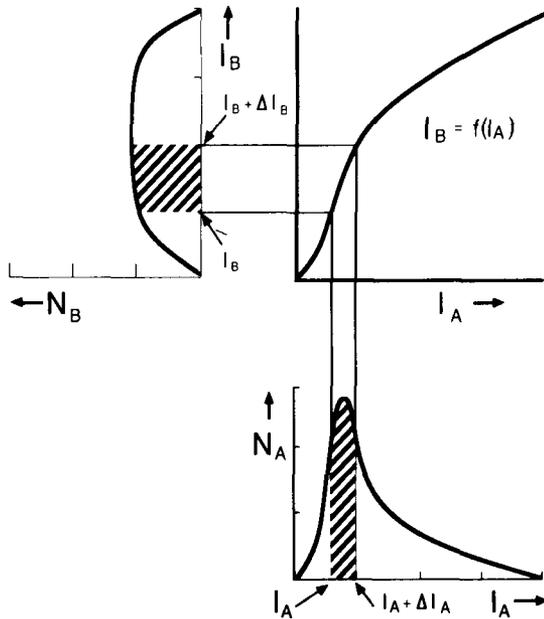


Fig 4. Effect of a nonlinear point operation and its relationship to the input and output image histograms.

before the injection are averaged to obtain a low noise mask, which in turn is subtracted from a postinjection image acquired when the contrast medium is passing through the blood vessels of interest. As shown in Fig 8, subtracting the mask leaves only the radiopaque blood vessels. Usually contrast medium is present in several of the postinjection images, and these can in turn be

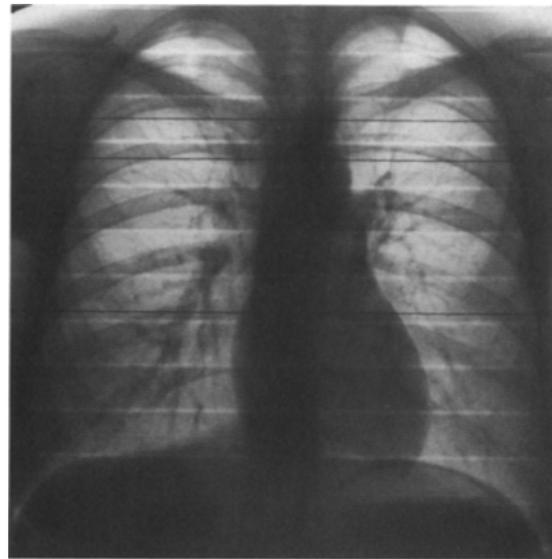
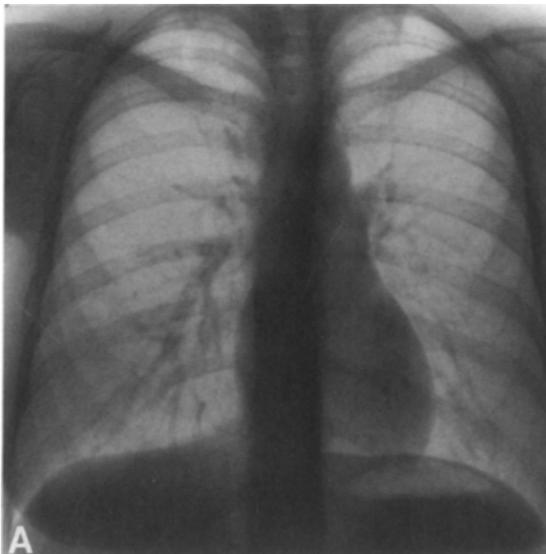


Fig 6. Digital chest image of Fig 1A before applying offset and gain corrections.

averaged or preferably their weighted sum obtained to further reduce the noise level in the final image. In the latter case, each image has a weight based on the amount of contrast medium present in the blood vessel(s). The subtraction is so perfect in modern equipment that virtually all background structures and bony landmarks are absent. This is clinically undesirable, and often the preinjection mask is added to the subtracted

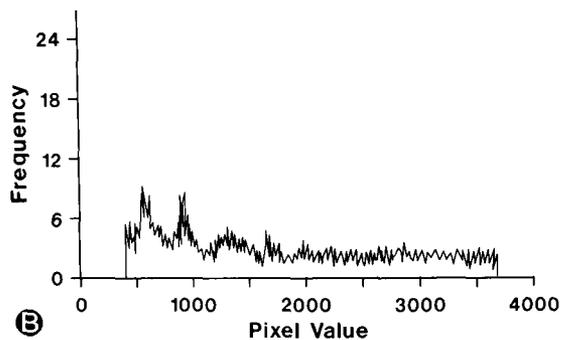


Fig 5. Result of histogram equalization processing on the digital chest image of Fig 1A (A) and on the image's histogram (B).

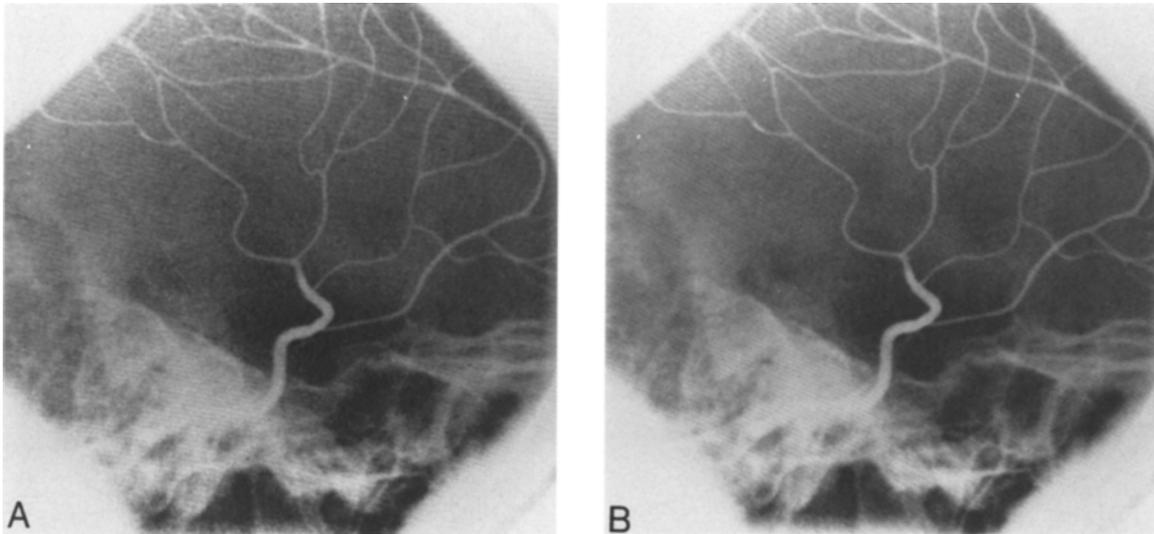


Fig 7. Comparison of a single frame (A) and an average of eight frames (B) of a digital fluorographic image.

image with a fractional weight to obtain an image with high-contrast vasculature and low-contrast bony landmarks. An example of this is shown in Fig 9.

Dual-Energy Techniques

The manipulations performed in dual-energy imaging are, in essence, point operations. The first step is to generate two images of the patient, identical in every respect except that one is obtained with an x-ray spectrum with a low mean energy and the other is obtained with an x-ray

spectrum with a somewhat higher mean energy. These are then decomposed point by point into water and bone images, whose intensity values are proportional to the overlying masses per unit area of soft tissue and bone.¹ The advantage of this type of image processing is that it separates the two types of anatomical information present and provides the radiologist with simpler images to view. Thus, a nodule that may be hidden by the ribs in a conventional chest film is clearly visible on the soft tissue image. Recent results indicate that a dramatic improvement in pulmonary nod-

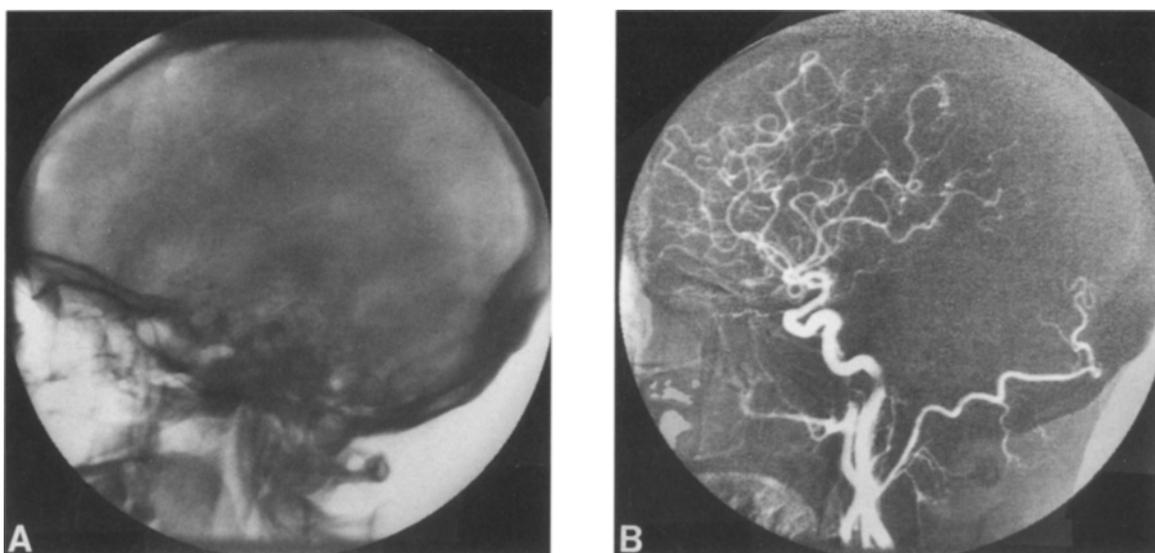


Fig 8. Postinjection DSA image (A) before and (B) after subtraction of mask.

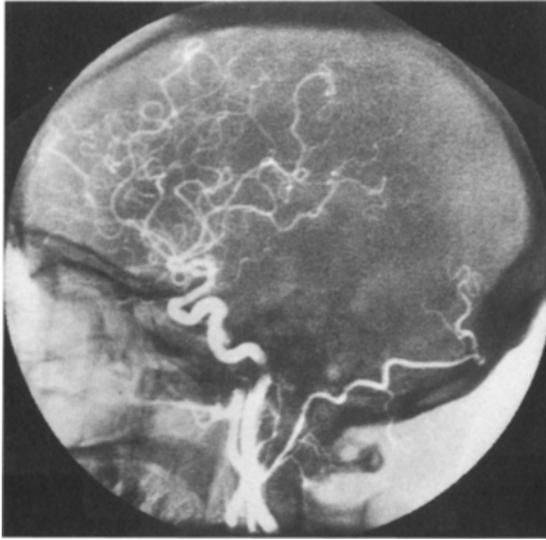


Fig 9. High-contrast fluorographic cerebral arteriogram obtained by adding the postinjection image (Fig 8A) at 0.25 weight and the subtraction image (Fig 8B) at 0.75 weight.

ule detection rates is realized with dual-energy techniques.² In addition, dual-energy chest radiography shows nodule calcifications with greater sensitivity and accuracy than other noninvasive techniques.³ This increased accuracy is an important capability, because the presence of calcium is a primary radiographic criterion for benignancy.

Mathematically, the dual-energy point opera-

tions can be represented as

$$w_b(m, n) = f[S_{lo}(m, n), S_{hi}(m, n)] \quad (9)$$

$$w_s(m, n) = g[S_{lo}(m, n), S_{hi}(m, n)] \quad (10)$$

where $S_{lo}(m, n)$ and $S_{hi}(m, n)$ are the pixel values of the low and high energy images and $f[]$ and $g[]$ are nonlinear transformations empirically determined from calibrations made with different masses per unit area of soft tissue (water) and a bone-equivalent material. A soft-tissue (bone-cancelled) image is then generated by adding a fraction (≈ 0.55) of the bone image to the initial soft-tissue (mass per unit area) image to fill in the voids or "shadows" left by the removal of bone. In addition, a normal or single energy image can be obtained by adding the original low- and high-energy images together. Soft-tissue (bone-cancelled) and bone (soft tissue-cancelled) images (Fig 10) can then be presented to radiologists simultaneously along with the standard single energy image.

LOCAL OPERATIONS 1: GEOMETRIC CORRECTIONS

Spatial Distortion

Often undesirable spatial distortions or changes occur when images are acquired. When it is clinically advantageous, these distortions can be corrected using digital processing techniques.

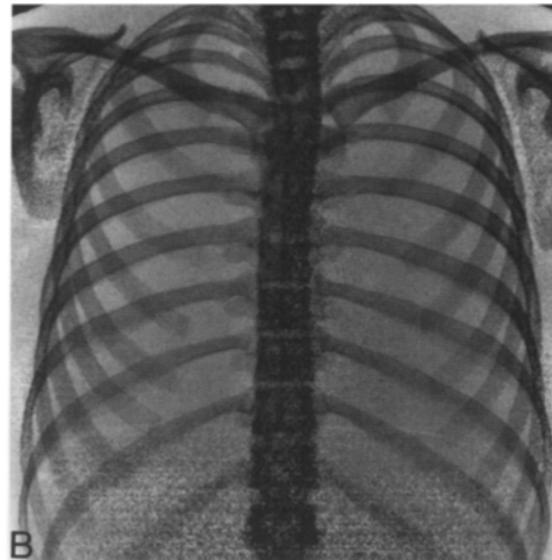
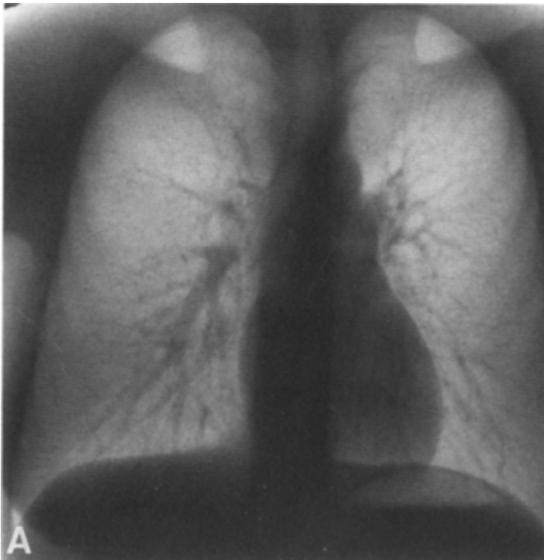


Fig 10. Soft tissue (A) and bone (B) digital chest images of reference patient (see Figure 1A) obtained with Picker-prototype dual-energy chest unit.

Examples of this type of image processing are correcting for motion misregistration in DSA, correcting for misregistration in dual-energy radiography using a sandwich detector, and correcting for image intensifier distortion in digital fluorography. In such cases, the final image is obtained from the original image using a two-stage process. The first step is to map the input image pixels to their location in the output image, and then to use this spatially shifted image to obtain the final image. In general, the first step will cause fractional pixel shifts, and interpolation is required to determine the value of the pixels in the final image. The simplest interpolation scheme is to use the nearest-neighbor value. This, however, is usually not satisfactory in medical imaging, because gray-scale values often change dramatically between adjacent pixel locations. A more desirable approach is to use bilinear interpolation of the four nearest-neighbor pixel locations. More sophisticated interpolation schemes using cubic spline, sinc, or similar functions can also be used. These

require more computational time and may, depending on the response of the image display, offer a clinical advantage.

The improvement that can be realized by using a patient-motion misregistration correction is shown in Fig 11. Misregistration between the front and back elements of a dual-energy sandwich detector results in imperfect bone and soft tissue cancellation and introduces noise. Likewise, when an image intensifier is used, "pincushion distortion" occurs. In the latter two examples, the spatial aberrations or distortions are systematic and can be quantitated by imaging a periodic pattern. Knowing what the calibration image should look like, one can then determine the pixel shifts that are necessary to remove the distortion. Figure 12 shows the improvement that can be realized when pincushion distortion is corrected for. Such corrections are essential in applications requiring precise spatial information from a conventional fluoroscopic imaging chain (ie, electro-optical image amplifier tube and television).

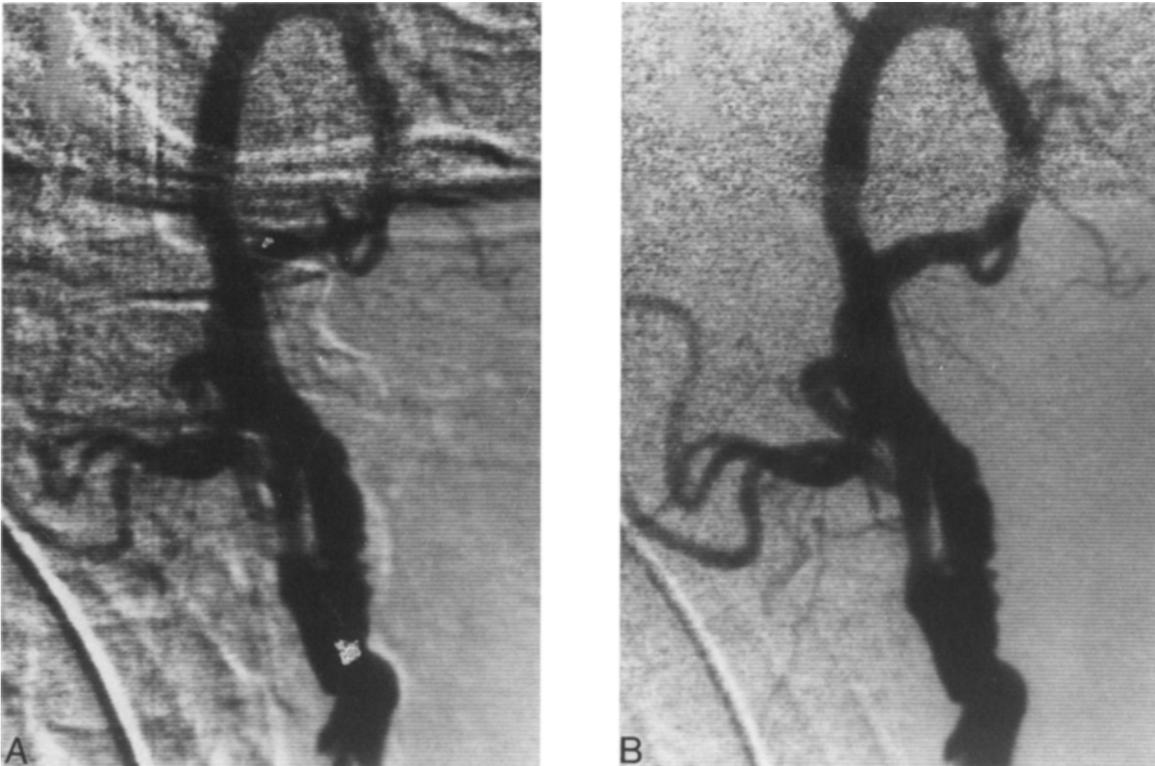


Fig 11. Comparison of DSA images that were obtained (A) without and (B) correcting for patient motion misregistration.

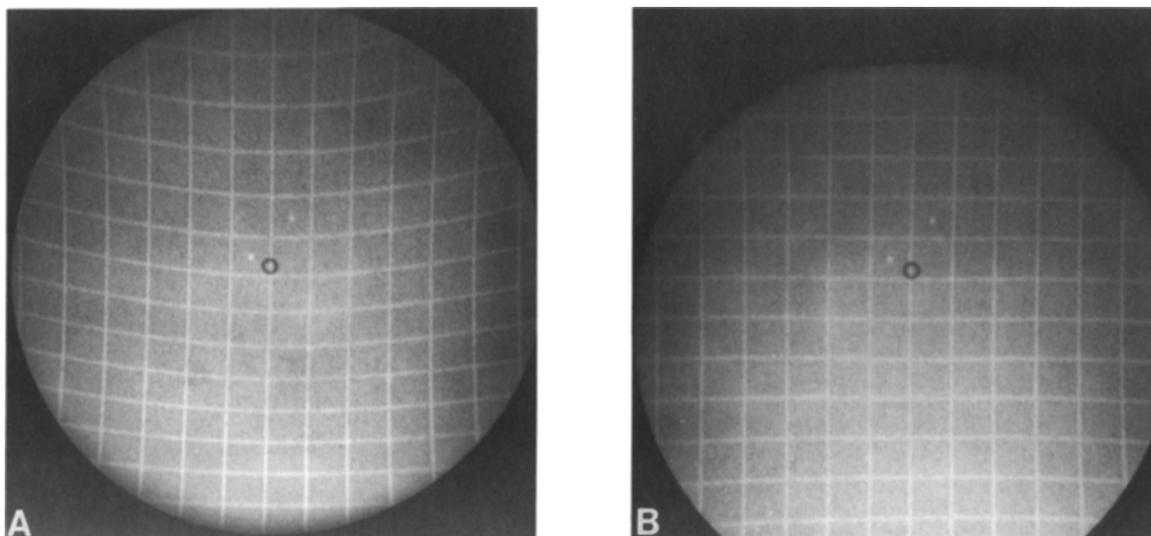


Fig 12. Fluorographic image of grid phantom obtained with the central ray angled at 28.4° to the plane of the image intensifier; ie, the x-ray tube was moved toward the top of figure; Unprocessed (A) and corrected for geometric distortion (B). Reprinted with permission.⁴

LOCAL OPERATIONS 2: FILTERING

Linear Filtering

Edge detail and image appearance is routinely enhanced in computed tomography scout images and in computed radiographic images by using linear and/or nonlinear processing techniques. An image processing operation is linear if the result of processing an image, which is the sum of two images, is equal to the sum of the images processed separately. That is, if

$$H\{I_1(m, n) + I_2(m, n)\} = H\{I_1(m, n)\} + H\{I_2(m, n)\} \quad (11)$$

then the image processing operation $H\{ \}$ is a linear operation. Linear systems analysis has an established theoretical basis and is commonly used in the study of electrical circuits, optical systems, and x-ray imaging systems. With the exception that discrete rather than continuous mathematics is used, the same theoretical framework can be used to study the linear filtering of digital images.

Convolution—Smoothing and Edge Enhancement

The operation of convolving one function with another is an essential concept in image processing and is best understood by considering the one-dimensional case. Mathematically, the con-

volution of the filter function or kernel $g(m)$ with the input image is represented by the expression

$$I_{out}(m) = \sum_{i=-(N-1)/2}^{(N-1)/2} g(i) \cdot I_{in}(m + i) \quad (12)$$

where N , for reasons of symmetry, is an odd integer. If the kernel extends over five pixels, then the following is true.

$$I_{out}(m) = g(-2) \cdot I_{in}(m - 2) + g(-1) \cdot I_{in}(m - 1) + g(0) \cdot I_{in}(m) + g(1) \cdot I_{in}(m + 1) + g(2) \cdot I_{in}(m + 2) \quad (13)$$

For simplicity and notational convenience, $*$ is often used to denote the convolution operation. With this notation, equation 12 reduces to the following.

$$I_{out}(m) = g(m) * I_{in}(m) \quad (14)$$

Extending equations 12 and 14 to two dimensions, one obtains the following relations.

$$I_{out}(m, n) = \sum_{i,j=-(N-1)/2}^{(N-1)/2} g(i,j) \cdot I_{in}(m + i, n + j) \quad (15)$$

$$I_{out}(m, n) = g(m, n) * I_{in}(m, n) \quad (16)$$

The degree and type of filtering, ie, smoothing or edge enhancement, depends on the shape of the kernel. For example, if a rectangular (rect)

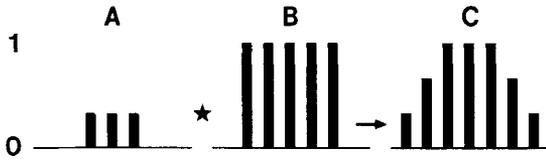


Fig 13. Convolution of three-pixel unity area rect function kernel (A) with five-pixel rect function image (B), yielding image (C).

function kernel (Fig 13A) is convolved with an image (Fig 13B), the result is a smoothed or blurred image (Fig 13C). How one obtains the final blurred image is conceptually illustrated in Fig 14. If one subtracts the blurred image from the original image, only the edges remain (Fig 15). Likewise, if the kernel has a positive central peak (Fig 16A) and negative side lobes, the output image will only contain information about the edges of the input image (Fig 16C). Figure 16C was obtained with one operation, whereas Fig 15D required two operations. This edge image can in turn be added to the original image to obtain an image in which the edges are enhanced (Fig 16D). This image could also be obtained with one operation if the positive central peak of the Figure 16A kernel is increased by one, ie, the kernel has unity area. In addition to illustrating the fundamental concept of convolution and the effects of smoothing and edge enhancement, there are also two salient points in

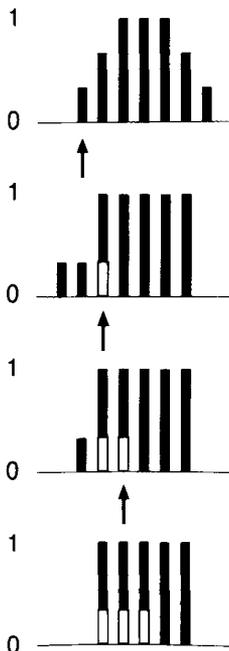


Fig 14. Illustration of the convolution process. The top line represents the resultant function (ie, Fig 13C). At any pixel location, the value of the process is determined by the overlap area of the kernel and the function it is operating on as the kernel is moved from left to right. The overlap area is shown in white in lines 2, 3, and 4 for three adjacent positions of the kernel. The arrow above the lines points to the pixel of the final function that the overlap corresponds to.

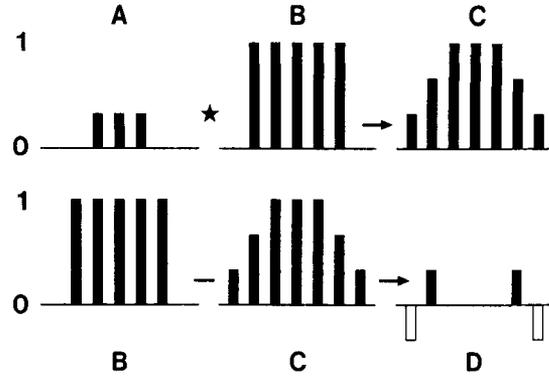


Fig 15. Convolution of unit area rect function kernel (A) with image (B) yielding "smoothed" or "blurred" image (C). An edge image (D) can subsequently be obtained by subtracting image C from image A.

these examples to consider. First, in digital image processing, the kernels are usually symmetric (ie, incorporate an odd number of pixels) and second, if the kernel sum is unity (ie, has unit area), the average pixel values of the input and output images will be the same.

A representative smoothing or blurring kernel is shown for two dimensions.

$$\begin{matrix}
 +1 & +1 & +1 \\
 +1 & +1 & +1 \\
 +1 & +1 & +1
 \end{matrix} \quad (17)$$

Convolution with such a kernel would also usually incorporate division by nine; therefore, as noted above, the average large area pixel values in the output and input images would be the same. The effect of a 3×3 smoothing kernel on a digital chest image is shown in Fig 17A and can

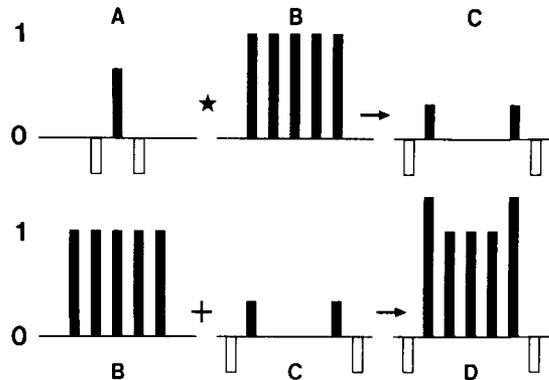


Fig 16. Convolution of zero-area positive central peak, negative-side lobe kernel (A) with image (B), yielding edge image (C). Adding this image to the original image results in edge-enhanced image (D).

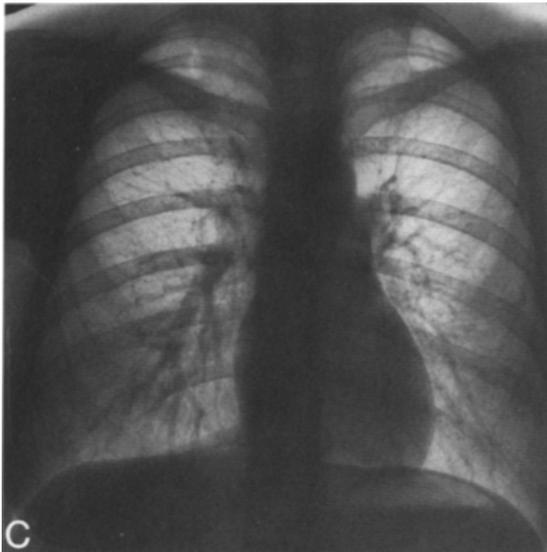
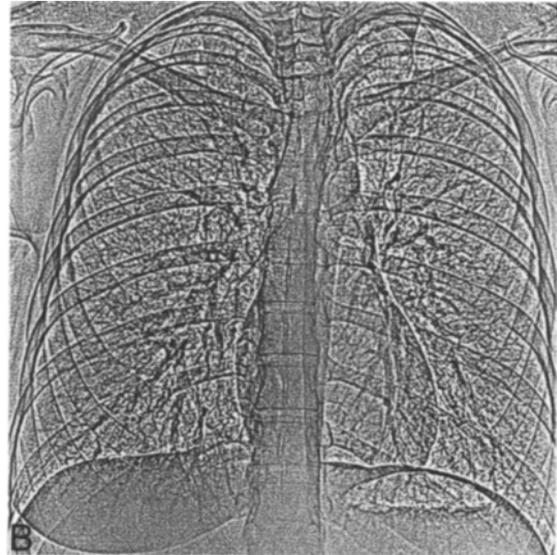
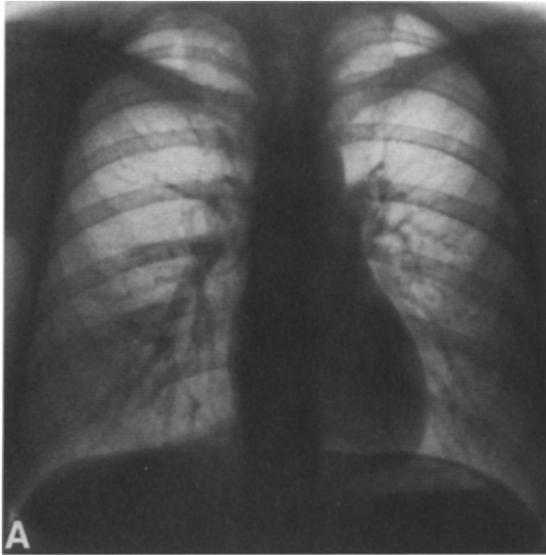


Fig 17. Smoothed chest image of reference patient obtained by convolving Fig 1A image with 3×3 rect kernel (A). Edge chest image obtained by convolving the same image with Laplacian kernel (B). Edge-enhanced image obtained by adding edge image and the original Fig 1A image (C).

be compared with the unprocessed image (Fig 1A). Parenchymal detail and noise are suppressed. A kernel that enhances edges is shown.

$$\begin{array}{ccc} -1 & -1 & -1 \\ -1 & +8 & -1 \\ -1 & -1 & -1 \end{array} \quad (18)$$

This is often referred to as a “Laplacian” kernel or filter, because it approximates an unoriented second-derivative operation. The result of applying it to our reference chest image is shown in Fig 17B. In essence, it leaves only the high spatial frequency edge information and is often used as an “edge detector” in image analysis algorithms.

As can be appreciated by comparing it with the unprocessed image, it has limited clinical utility other than its use as an edge detector. A Laplacian kernel is useful, however, when it is added to the original image (Fig 17C). In practice, this can be combined into a single operation with the following kernel.

$$\begin{array}{ccc} -1 & -1 & -1 \\ -1 & +9 & -1 \\ -1 & -1 & -1 \end{array} \quad (19)$$

As one would expect from comparing these four images, Fig 17C is preferred by radiologists.

An alternative and commonly used linear high-

pass filtering approach is “blurred or unsharp mask subtraction.” As the name states, one convolves the input image with a blurring kernel (ie, 50×50) and then subtracts the result (often fractionally weighted) from the input image. The blurred image contains only low spatial frequencies, and the subtraction of this low-frequency image from the original image suppresses the low frequencies and thereby enhances the final image’s medium and high spatial frequencies. An

example of this two-step approach (ie, low-pass filtering followed by the algebraic operations of subtraction and addition) is shown in Figs 18A-C. As shown in Fig 18D, adding a Laplacian filtered image further enhances the high spatial frequency information. The advantage of this approach over that used to obtain Figure 17C is that only the very low spatial frequencies are suppressed; the medium frequency content is maintained, and the high spatial frequencies are

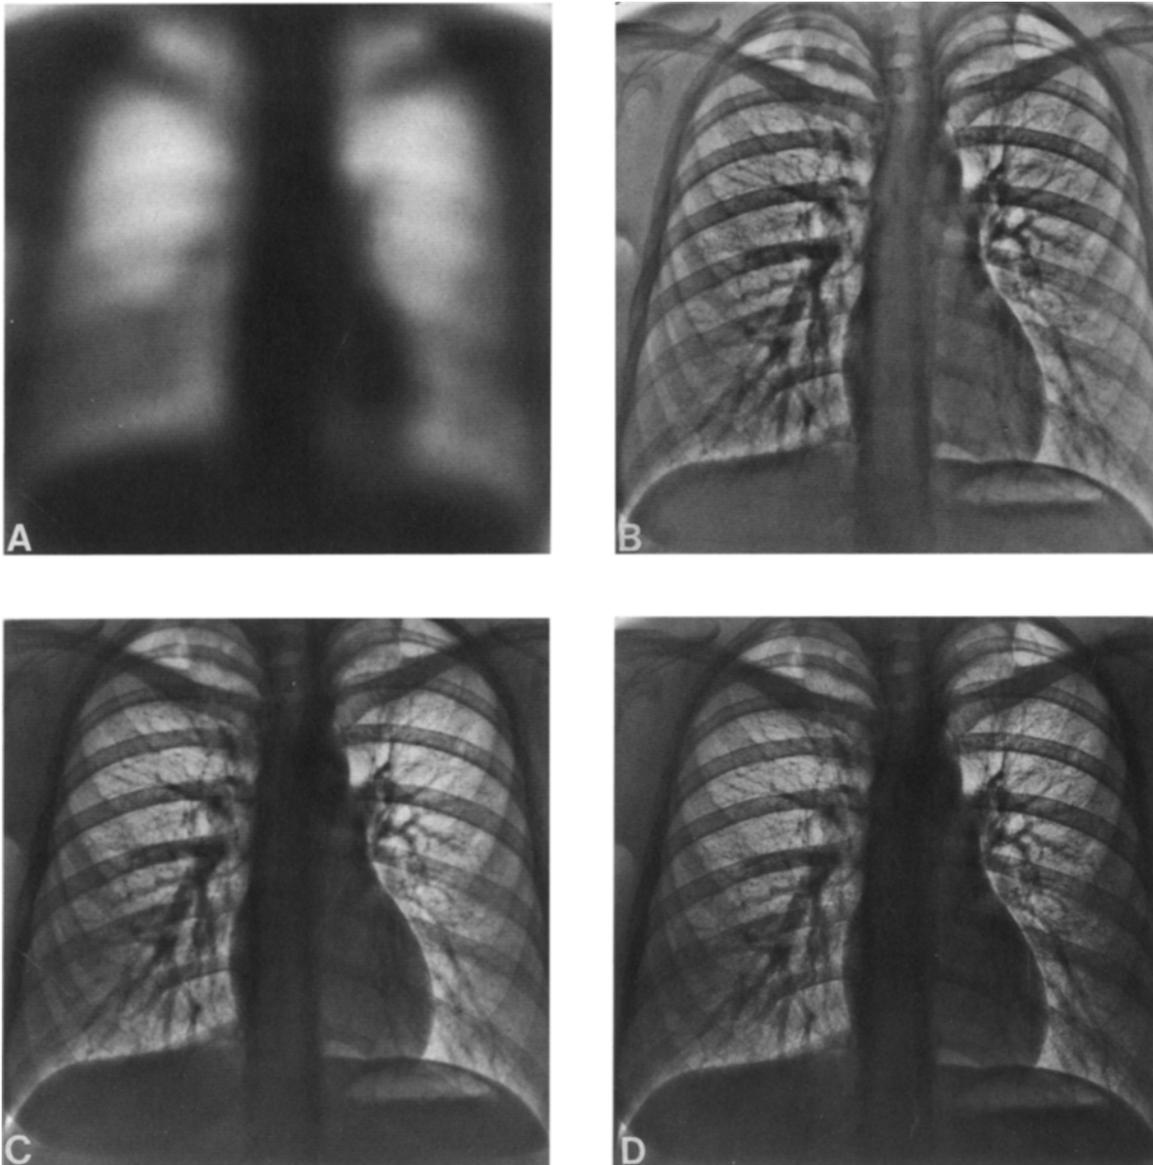


Fig 18. (A) Smoothed (unsharp) chest image obtained by convolving Fig 1A with 50×50 rect kernel. (B) Edge image obtained by subtracting unsharp image from the original image. (C) Blurred or unsharp mask subtraction image obtained by linear combination of B and Fig 1A. (D) Edge-enhanced blurred mask subtraction image obtained by adding C and Laplacian edge image (Fig 17B).

enhanced. As can be appreciated by comparing Figures 17C and 18D, this allows for the display of a wide range of anatomical information as well as emphasis on fine detail in a single image. Of practical interest is that the same result can be achieved by convolving the original image with a single, appropriately designed kernel.

Convolution and the Fourier Transform

The discrete Fourier transform of an image $g(l\Delta p)$ is given by

$$G(f) = \sum_{l=0}^{N-1} g(l\Delta x) e^{-j2\pi f l \Delta p} \quad (20)$$

$$f = k/N\Delta p, k = 0, 1, 2, \dots, N - 1$$

where N is the number of pixels comprising the image and Δp is the size of the pixels. Likewise, the image can in turn be determined from its Fourier transform as follows.

$$g(x) = N^{-1} \sum_{k=0}^{N-1} G(k/N\Delta p) e^{j2\pi kx/N\Delta p} \quad (21)$$

$$x = l\Delta p, \quad l = 0, 1, 2, \dots, N - 1$$

The functions $g(x)$ and $G(f)$ are known as a Fourier transform pair, and for any $g(x)$, $G(f)$ is unique, and vice versa.

A useful property of Fourier transforms is that the transform of the convolution of two functions is the product of the individual Fourier transforms of the two functions in the frequency domain.

$$g(x) * h(x) \rightarrow G(f) \cdot H(f) \quad (22)$$

This implies that one can process an image by first Fourier transforming it to frequency space, multiplying it by the transform of the kernel, and then transforming the result back to the spatial domain. Although this can be done, it usually is not because of the excessive amount of computer time required. An exception to this generalization often occurs in the processing of nuclear medicine images involving small matrices and reasonably short Fourier transform execution times on an array processor. The large matrices used in digital radiography dictate that one use the convolution operation.

Implicit in equations 20 and 22 is that the frequency response of a filter is given by its Fourier transform. Plotted in Fig 19 are the

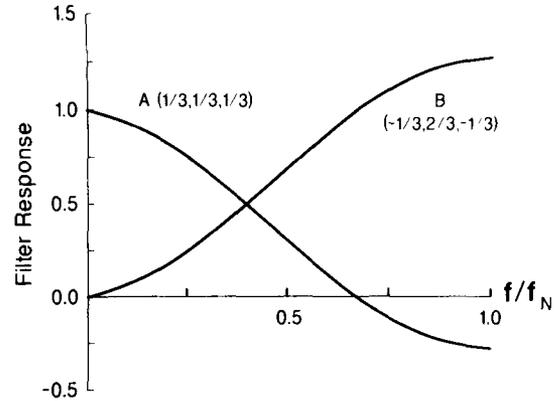


Fig 19. Fourier transform of Figs 13A and 15A rect function kernel (A) and Fig 16A Laplacian kernel (B).

Fourier transforms of the kernels used to obtain the one-dimensional processed images in Figs 13, 14, and 16. As would be expected, the response of the smoothing kernel falls off at higher spatial frequencies; because of this, it is referred to as a low-pass filter. The behavior of the Laplacian filter is quite different and has a response only at higher spatial frequencies; thus, it is a high-pass filter. Adding the original image to the Laplacian filtering operation enhances the high spatial frequencies and also maintains the low spatial frequencies. The ability to analyze the frequency response of a filter is of fundamental interest and also useful in its design. In practice one often uses a high-pass filter that compensates for modulation transfer function fall-off.

Nonlinear Filtering

A high-pass linear filter enhances the high spatial frequency content of the low- as well as the high-pixel-count regions. This often poses a problem in the low-pixel-count regions, because there is little information and considerable noise, which becomes more apparent with a linear high-pass filter. An alternative and often more pleasing result can be obtained by scaling the degree of edge enhancement with pixel count. An example of this type of adaptive filtering on our reference chest image is shown in Fig 20. This image was obtained by a linear combination of the original unprocessed input image and the edge-enhanced blurred mask-subtraction image (Fig 18D). At the high-pixel-count regions (ie, in the lung), the edge-enhanced blurred mask-subtraction image is given a weight of one and

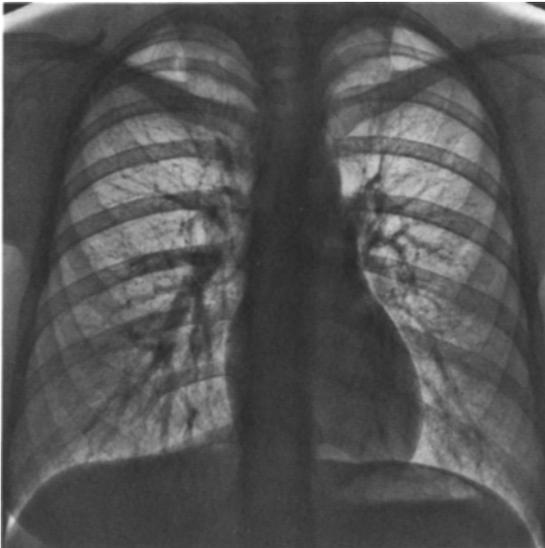


Fig 20. Adaptively filtered reference chest image obtained by linearly combining Fig 1A image with Fig 18D image. In the lung field (ie, the high-pixel-count region) the Fig 18D image had a weight of 1.0 and the unprocessed image a weight of zero. In the subdiaphragmatic region (ie, the low-pixel-count region) the Fig 18D image had a weight of zero and the unprocessed image a weight of 1.0. In between these two regions, the weight of the processed image increased and the weight of the unprocessed image decreased linearly with pixel count.

the original unprocessed image a weight of zero. In the low pixel count regions (ie, below the diaphragm), the edge-enhanced blurred mask-subtraction image is given a weight of zero and the unprocessed image a weight of one. In between the two regions, the weighting of the edge-enhanced blurred mask-subtraction image increases linearly, and the weighting of the unprocessed image decreases linearly with pixel count of the original image. In comparison with our previous best (Figure 18D), the degree of parenchymal detail is similar and the noise in the subdiaphragmatic region and mediastinum is less. Needless to say, this image is preferred by radiologists over the other types of filtered images of this patient.

DISCUSSION

There are several important practical applications of point- and local-image-processing operations in digital radiography. Windowing and leveling, histogram shaping, and correcting for equipment nonlinearities are commonly used point operations. DSA is based on the frame point operations of addition and subtraction. Soft-

tissue (bone-cancelled) and bone (soft tissue-cancelled) dual-energy images are obtained by using nonlinear frame-point operations. Local geometric operations are used in DSA to correct for patient motion misregistration and can also correct for imaging-system distortions.

Local-filtering operations are used to enhance either the low, medium, or high spatial frequencies of an image. Generally, the objective is to suppress much of the low spatial frequency (quasi-DC) and to enhance the high spatial frequency (small detail and edge) information. With blurred- (unsharp) mask subtraction, this is accomplished by convolving the image with a low-pass kernel and then subtracting the resultant unsharp mask from the original image. Alternatively, one can convolve the image with an edge-enhancing kernel. The optimal degree of filtering is unknown and is likely to vary from examination to examination. Probably more important than what filtering is used is the fact that it is not varied from day to day or week to week. Radiologists' pattern recognition depends on image consistency, and if the images appear different from day to day, radiologists lose their bases of reference and risk misreading them.

In addition to enhancing the diagnostic information content of an image, one or more of the digital image-processing techniques described above are used when quantitative information has been obtained. Examples are nuclear medicine functional studies, ie, cardiac ejection fraction, segmented wall motion, and effective renal plasma flow. Examples in other areas are quantification of blood vessel percent stenosis (DSA) and pulmonary nodule calcium content (dual-energy chest radiography).

Feature extraction and the development of algorithms for computer-aided diagnosis are research areas that hold considerable promise. The general strategy is to analyze the physical characteristics of the features of interest and then to use image-processing techniques (ie, edge detection, gray-level thresholding, etc) to isolate these features in the image of interest. The characteristics of the isolated features are in turn compared with reference data. Areas of current interest are lung nodule detection, microcalcification detection and analysis in mammography, and interstitial disease. The results to date are provocative, and techniques being developed could have a significant effect on the practice of medicine.

BIBLIOGRAPHY

Brigham EO: The Fast Fourier Transform. Englewood Cliffs, NJ, Prentice-Hall, 1974

Castleman KR: Digital Image Processing. Englewood Cliffs, NJ, Prentice-Hall, New Jersey, 1979

Gonzalez RC, Wintz P: Digital Image Processing (ed 2). Reading, MA, Addison-Wesley, 1987

Hall EL: Computer Image Processing and Recognition. San Diego, CA, Academic, 1979

Pratt WK: Digital Image Processing. New York, NY, Wiley, 1978

Rosenfield A, Kak AC: Digital Picture Processing. San Diego, CA, Academic, 1976

REFERENCES

1. Lehmann LA, Alvarez RE, Macovski A, et al: Generalized image combinations in dual kVp digital radiography. Med Phys 5:659-667, 1981

2. Niklason LT, Hickey NM, Chakraborty DP, et al: Simulated pulmonary nodules: Detection with dual-energy digital versus conventional radiography. Radiology 160:589-593, 1986

3. Fraser RG, Hickey NM, Niklason LT, et al: Calcification in pulmonary nodules: Detection with dual-energy digital radiography. Radiology 160:595-601, 1986

4. Chakraborty DP: Image intensifier distortion correction. Med Phys 14:249-252, 1987