

Computer Processing Methods for Nuclear Medicine Images

Tracy L. Faber and Russell D. Folks

Emory University School of Medicine, Atlanta, Georgia

Objective: This review is a discussion of computer processing methods for modern nuclear medicine images.

Methods: Basic image processing and quantification techniques are described, and their application to specific nuclear medicine studies is presented. Generation and use of regions of interest, time-activity curves, tomographic reconstruction and oblique reformatting are all addressed, as well as common nuclear medicine protocols for processing multiple gated blood pool studies, first-pass cardiac studies, renograms and tomographic perfusion studies.

Results: Since processing cannot be separated from acquisition and display methods, short summaries of necessary information on these topics are included with primary emphasis placed not only on how and what processing and quantification steps are taken, but also on *why* they are taken.

Conclusion: Although some nuclear medicine images require very little processing, others are quite complex and are not always fully understood. The technologist's comprehension of these procedures will ensure nuclear medicine's successful foray into an ever-changing future.

Key Words: Computer processing; quantification techniques; region of interest; planar images; digital images; gated acquisition; dual-isotope acquisition; smoothing; filtering; profile; time-activity curve; background; filtered backprojection

J Nucl Med Technol 1994;22:145-162

This is the third article in a four-part series on computers in nuclear medicine. Upon completion, the technologist should be able to list computer processing methods, describe processing methods for common nuclear medicine studies and understand why these processing steps are taken.

Computer processing is a prerequisite for display and quantification of all modern nuclear medicine images. The outcome of clinical reading and diagnosis directly depends on the technique and type of processing performed. The appearance of images can be drastically changed, which may either enhance or degrade the presentation of important information. Compli-

cated mathematical techniques can be applied to image data to provide either accurate or inaccurate quantitative measures of physiological function. To ensure that image quality is enhanced and quantified values are correct, the user must have a competent understanding of the processing techniques.

As computer hardware and software capabilities have expanded, the complexity of the processing techniques applied to nuclear medicine images has increased. Surprisingly, though, most specific methods can be described in terms of some fairly simple general processes. If the basic ideas are understood, then the more complicated specific applications can be more easily comprehended.

General understanding of a processing technique includes the knowledge of what an operation really does, how it works and why it is important. The grasp of this information allows the user to adjust strategies to particular situations. Fully understanding the basis of processing procedures helps in solving problems and overcoming processing failures. Although push-button macros and protocols that can process and quantify images in a standard way are efficient and helpful, they almost always require user interaction, and understanding, at some point. The user has a responsibility not only to solve problems but also to guide and interact with a semiautomatic processing routine.

In short, computer image processing is a very important part of clinical nuclear medicine. The following sections should explain how and why some of the most common procedures are performed. These sections should help the technologist apply them correctly in day-to-day situations, and even judge which methods might be appropriate for processing new types of studies or extracting new information from standard studies. More information on general image processing may be obtained from Gonzales et al. or Jain (1,2). More detailed descriptions of various nuclear medicine analysis techniques may be found in Gottschalk or Harrison (3,4).

PLANAR NUCLEAR MEDICINE IMAGE PROCESSING

Digital Images

A digital image is a two-dimensional set, or matrix, of numbers associated with some measurement. In nuclear

For reprints contact: Tracy L. Faber, Division of Nuclear Medicine, Emory University Hospital, 1364 Clifton Road, N.E., Atlanta, GA 30322.

medicine, the numbers in the image represent the number of counts detected as they emanate from the patient's body. This set of numbers can be viewed as a picture, if the numbers are somehow "mapped" to colors or shades of gray. Usually, high counts or numbers are mapped to bright colors while low counts are mapped to dark colors. Each number or element of the two-dimensional matrix is called a pixel. Each pixel has a size in millimeters or centimeters. In general, the area of a pixel defines the area of the body "covered" or "seen" by that pixel. That is, a 5×5 -mm pixel contains counts that arise from a 5×5 -mm area of the patient. The size of the image matrix, expressed in pixels, is generally a factor of 2 (64×64 , 128×128 , 256×256 , etc.) and defines the total amount of body area from which counts are collected. A 128×128 -pixel image that has pixel size 5×5 mm, extends over a (128×0.5 cm/pixel $\times 128$ pixel $\times 0.5$ cm/pixel =) 64×64 -cm portion of the patient's body.

While the display of the digital nuclear medicine image is extremely important, the actual numbers or counts in the image matrix create the display. The appearance of the image can be changed by mathematically modifying its underlying numbers. These values can also be used to give quantitative or numerical measurements of radioactivity within organs. These measurements can then be combined or otherwise operated on to create physiologically important parameters, such as ejection fraction, glomerular filtration rate or tumor size.

Planar Acquisitions

Static Acquisitions. The basic nuclear medicine acquisition is the static planar view. A single image is collected from the gamma camera placed near the patient's body. If additional views are desired, either the camera or the patient is moved to acquire another image. A liver scan is an example of a static study. Usually, in static acquisitions the same acquisition parameters (pixel size, matrix size and acquisition time) are the same for all views.

Dynamic Acquisitions. A dynamic study is one in which a continuous series of images is acquired over time. These are used to watch the time course of the radiopharmaceutical; such information contains clues to an organ's functioning. Often, the rate at which dynamic frames are acquired varies. Many frames may be acquired during the first few minutes in order to see details of how the pharmaceutical is absorbed by the organ. Later, frames may be acquired at a slower rate, as concentration changes more slowly during washout, redistribution or clearance. Other image parameters, such as matrix or pixel size, may also be changed during dynamic acquisitions. A renogram is an example of a commonly acquired dynamic study.

Gated Acquisitions. A gated acquisition is one in which the start of acquisition is triggered by a physiological signal, primarily the electrocardiogram (ECG). In these gated cardiac studies, the key point is that a user-selected number of views in time, or frames, are collected for each cardiac cycle. Prior to the acquisition, the R-R interval of the ECG is examined by the computer to determine the patient's heart

rate, which is divided by the number of gated frames needed to give the acquisition duration of each frame. During acquisition, the computer "directs" counts to the first frame when the R-wave from a QRS complex is detected. This first frame is collected for its allotted duration, and then counts are "redirected" to the second frame and so on.

This continues until the next QRS complex. Each collected frame is created using hundreds of cardiac cycles, so the resulting images are representative of an "average" heartbeat. This kind of gated acquisition is called "frame mode." Additionally, in cardiac perfusion studies, the patient's breathing pattern can induce motion of the chest and heart, particularly when the scan is performed soon after an exercise stress test. For this reason, some laboratories perform perfusion scans using respiratory gating, although, this technique is less common than use of the ECG as a trigger.

One drawback of the frame mode is that irregular heartbeats (or respirations, in the case of respiratory gating) are included in the acquisition. This causes the average heartbeat image to contain some abnormal contractions, and this can compromise image quality. In addition, if the heart rate changes, each gated frame will start to hold larger or smaller portions of the cardiac cycle than originally intended. For example, consider the case of a heart rate that speeds up to twice its original rate with a 16-frame gated acquisition. Originally, the end-systolic frame occurred in about the middle of the 16-frame study, near frame 8.

However, as the heart rate doubles, the end-systolic frame will occur nearer to frame 4. This means that the average heartbeat seen in the gated set is actually varying radically. Also, frames 8–16 will have fewer counts than frames 1–8, since the entire cardiac cycle actually fits into only 8 frames instead of the original 16 with a doubled heart rate. Most modern nuclear medicine systems help reduce this problem by continuously monitoring heart rate and adjusting frame acquisition times to account for changes. This adjustment is rarely perfect, however, and often the final frames of the study will have reduced counts because of heart rate changes during the acquisition.

A different gated acquisition mode is called the "list mode," in which every detected count is stored and listed, along with its location. Time markers are also placed in this list, to give some indication of *when* each count was acquired. In order to acquire gated information, markers denoting the occurrence of the R-wave trigger of the QRS complex can be stored in the list. After the acquisition, the list of counts, their positions and the timing information can be used to create a set of gated frames. However, because the time and R-wave markers are also stored, counts acquired during irregular heartbeats can be discarded. Likewise, the duration of each frame can be easily modified if the heart rate changes, since the timing information indicates when each count occurred in the cardiac cycle. The drawback of list mode acquisitions is that computer storage requirements are very high. In the frame mode, only the running summation of counts is stored, but the list mode stores every count and its location separately. For this reason,

while planar gated studies may be feasibly acquired in the list mode, gated tomographic acquisitions are generally acquired in the frame mode.

Dual-Isotope Acquisitions. Most new hardware and software allow for photons from two different radioisotopes to be acquired simultaneously into separate images. This is possible when the energies of the emitted photons from the two nuclides are sufficiently different so that the hardware can determine which photons emanate from which isotope. Acquired counts from one nuclide get directed into one image matrix, while counts from the second nuclide get directed into a second image matrix. This allows for more than one physiological process, as measured by the two radioisotopes, to be imaged at one time.

Dual-isotope acquisitions can usually be obtained in either static, gated or dynamic modes, although this depends on a number of factors related to imaging hardware and software. Dual-isotope acquisitions require twice as much computer memory and disk space storage as regular single-isotope acquisitions. Gated acquisitions, therefore, that already require large amounts of memory and storage space, may not always be performed in a dual-isotope mode.

Processing dual-isotope acquisitions is generally no different than processing two separately acquired images. Sometimes, however, the image of the lower energy radionuclide must be corrected in light of the scattered counts of the higher energy nuclide. When the higher energy counts bounce off other nuclei within the body, they lose some of their energy and may then be mistakenly detected as photons from the lower-energy nuclide. This "down-scatter" from the higher-energy photons degrades the image of the lower-energy photons, unless some correction is performed.

General Planar Postprocessing

Display. In many types of planar studies, little actual processing is done except for displaying the images in a way such that they can be easily interpreted by the physician. There are two common operations performed on all images that do not affect the raw data within the digital matrix, but simply change the way the images appear on the computer screen.

Image Interpolation. Interpolation allows an image to be enlarged so that it is easier to see and analyze. Interpolation creates new pixels in between the old ones, so the final result is more pixels in the matrix, each with a smaller pixel size than in the original image. Interpolation is almost always bilinear; that is, the new pixel values in the zoomed image are created by computing the average count value between the old numbers. Bilinear interpolation keeps the resulting zoomed image smooth no matter how large it gets. In contrast, another type of interpolation, "nearest-neighbor" interpolation, creates new pixels by repeating the count values of the old numbers. This kind of interpolation results in a "blocky" looking image that becomes more distracting as the image is zoomed larger.

Image Scaling. Scaling is another function that affects how the image appears on the screen. Some factor, e.g., 4 or

1/256, is multiplied with every pixel in the image. Most commonly, this factor is automatically chosen so that the highest count becomes 255 and the lowest count becomes 0. These numbers are important since most nuclear medicine computers can still only display values between 0 and 255. However, it is usually possible for the operator to choose his or her own scale factor, e.g., one that maps count levels from 0 to 10 to new values ranging from 0 to 255. With this particular mapping, all counts originally greater than 10 are now mapped to 255. This kind of scaling allows very low-count regions of the image to be better visualized.

Scaling two different images to the same maximum and/or minimum value is called normalization. It allows easier comparisons between images by making their maximal counts equivalent.

Image Parameters. More commonly, the same effect of the scale factor is achieved by altering the color map. Every color map, including the monochromatic gray scale, has two parameters which define how it interacts with the image. These are termed either the window and the level or the "upper bound" and "lower bound." The lower bound and lower level are equivalent measures; they determine the image count level that gets mapped to the bottom of the color map. Any image values below the lower level will also appear as this bottom color. The upper level of the color map is the image value that gets mapped to the top of the color scale; likewise, any value greater than the upper level gets mapped to the highest color. Window refers to the length in image counts of the color map so that the upper level is equivalent to the level + window.

These parameters are displayed in different ways on various computer systems. Since many systems can only display images in the range of 0–255, some systems will scale the images to this range, and then display the upper and lower bounds in terms of these 0–255 values. Other systems, even if they must scale the images for display, still define the bounds in terms of the original count values. Finally, the bounds can be displayed in terms of percentages of the maximum image count.

It should be noted that scaling and color map adjustment can drastically change the appearance of an image. While this can be used to great advantage—for example, to enhance certain important image features—it may also be confusing or even misleading. Many physicians expect clinical images to be displayed in a standard way; if one deviates from this standard, some notation of this fact should be made on the image. In addition, some quantitative programs attach very specific meanings to specific colors; color maps should not be adjusted when this is the case.

Smoothing or Filtering. The term filtering is used to indicate operations that either smooth or sharpen images; that is, increase or decrease image blurring. In planar nuclear medicine applications, however, only smoothing is frequently employed. While filtering can be described mathematically in complex terms, the actual operation is usually simple.




FIGURE 1. Smoothing operation. (A) The original image matrix. (B) The 3×3 smoothing kernel. (C) Centering the kernel over one pixel of the image. (D) Weighted average of the center pixel and its neighbors: each pixel under the kernel is multiplied by the kernel weight over it, and the results are summed together. (E) Pixel underneath the kernel's center is replaced by new weighted average. (F) Steps in C–E repeated for each image pixel; this is the final, smoothed image.

Smoothing generally requires definition of a smoothing kernel or small matrix; often a few smoothing kernels are predefined in nuclear medicine computer systems. The smoothing kernel is generally square with a matrix size of 3×3 pixels, 5×5 pixels or 7×7 pixels. The numbers contained in the matrix are called weights. In a sense, the kernel is placed so that its middle pixel is on top of one pixel in the image. The kernel will then cover a 3×3 (or 5×5 , etc.) pixel region of the image.

The counts in the image pixels underneath each element of the 3×3 kernel are multiplied by the kernel weights and the results are all summed together. The image pixel beneath the center of the kernel is replaced by this summation. This operation is repeated for every image pixel. Figure 1 shows the process graphically. In this process, each image pixel is replaced by a weighted average of itself and some of its neighbors in its immediate neighborhood. This averaging operation tends to blur the image. As the kernel gets larger,




FIGURE 2. Original stress and delayed ^{201}Tl planar cardiac perfusion scintigrams shown in top row. Smoothed versions shown underneath in the second row.

and/or the weight of the center pixel of the kernel gets lower, the smoothing operation blurs the image more and more. However, along with the blurring comes noise removal; very high or low “spiky” noise pixels get averaged out. This is visible in the images displayed in Figure 2. Smoother images may be easier to analyze, but it is possible to over-smooth images to the point where the blur degrades analysis accuracy.

Gated or dynamic images may be filtered or smoothed in the time dimension as well. For example, the counts at pixel (30,30) in frame 3 of a gated study might be replaced with a weighted average of the pixels at location (30,30) in frames 2, 3 and 4. Time averaging gives the appearance of smoother motion in gated images, and once again, removes noisy pixels.

Profiles. A graph of image counts along any line drawn in an image is called a profile. Figure 3 displays a planar cardiac perfusion image and a profile that travels horizontally through it across the left ventricle. Profiles show how count




FIGURE 3. A profile taken through the left ventricle in a planar cardiac perfusion study. Note the two main profile peaks corresponding to the lateral and septal ventricular walls.

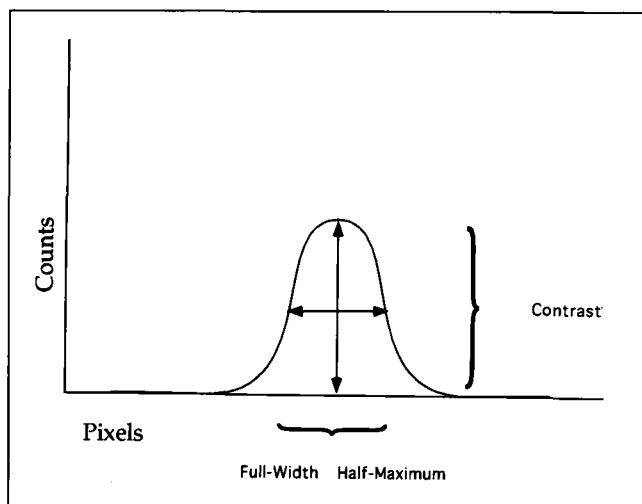


FIGURE 4. How a profile taken through a point source is used to calculate contrast and resolution (FWHM).

levels rise and fall in a structure more accurately since, as noted before, color maps can mask count changes.

Profiles are often used as part of quality control to analyze contrast and resolution. Contrast is the difference between count rates in an object and in the background, and it can be measured by computing count values at the peak and valley of a profile. Resolution is a measure of how close two points can be and still be discerned as two objects in an image. The physics of collimation cause a point source to appear quite blurred in a nuclear medicine image, and as blurring increases, resolution is said to decrease. Resolution is measured by creating a profile through an image of a point source, and measuring the width of the profile at one-half its maximum value. The width is called the full-width at half-maximum (FWHM). The more the peak profile is spread out, the larger the FWHM and the lower the resolution. Measurements of both contrast and resolution are displayed in the profile of a point source shown in Figure 4.

Image Registration. Image registration is the shifting of the image, i.e., the matrix of values which represents the count distribution within the two-dimensional plane in which it resides. As an example, consider an image composed of a 64×64 matrix of pixels. The image may be placed on a field of view whose coordinate system is left-handed, i.e., position: $x = 0, y = 0$ is at the upper left, so that the x -coordinate values increase from left to right, and y -coordinate values increase from top to bottom.

Consider an area of interest whose left upper corner is at pixel position where $x = 32$ and $y = 32$. Assume that we wish to shift the area of interest 2 pixels downward in the field of view (in the positive y direction) and 2 pixels rightward in the field of view (in the positive x direction). After such a shift, the same area of interest now has its left upper corner at a pixel position where $x = 34$ and $y = 34$.

The need to shift the image may arise because of patient motion during specific frames of a dynamic or tomographic imaging sequence or out of a need to superimpose images acquired at different times. Image registration has applica-

tion in cardiac perfusion imaging, which is particularly susceptible to motion artifacts owing to the quantification which may be performed. Additional applications include gastric emptying studies and parathyroid imaging.

REGIONS OF INTEREST

The region of interest (ROI) is probably the most often used processing tool in nuclear medicine. An ROI is a closed boundary surrounding some feature or area, for which some statistics are desired.

Creating an ROI

Most systems allow creation of simple geometric ROIs, e.g., squares, rectangles, ellipses, etc. More complicated ROIs can be drawn freehand, using a trackball, light pen or mouse. The key to drawing an appropriate ROI is simple in theory, but very complicated in practice. First, however, the technologist needs to know why that region is being drawn. ROIs are used in many different situations, some of which will be discussed in further sections. It is important to understand not only the processing techniques that will use the ROI but also basic anatomy in order to create meaningful ROIs.

Automatic techniques for creating ROIs are available on most clinical systems. The most common type of automatically generated ROI is one based on a "threshold" or "isocontour." Some count value is chosen, i.e., the threshold value, and every pixel above that value is included in some ROI. The threshold value may be automatically set by a processing program; more frequently, the user decides what that threshold should be. Note that this threshold technique can give more than one ROI in an image; for example, a single threshold in a renogram can give three ROIs—one surrounding each of the two kidneys and a third surrounding the bladder. If only one ROI is desired, generally the user must designate the important one.

Isocontours are similar to thresholds, but this term is *usually* used to indicate more than one threshold applied to an image at one time. Multiple count values separated by some constant difference (e.g., 10, 20, 30, 40 . . . or 25, 50, 75, 100 . . .) are used as thresholds to generate multiple ROIs. These isocontour lines are similar to those used on geographical maps containing contour lines of similar altitudes. In nuclear medicine, isocontours follow curves of similar count values. Isocontours can be helpful in deciding an appropriate single threshold to define the desired ROI. Figure 5 shows circular, freehand-drawn, threshold-generated and isocontour ROIs applied to the same image.

Some processing programs use sophisticated automatic or semiautomatic techniques to find ROI boundaries. While such programs run the gamut from simple image processing techniques to artificial intelligence approaches, and are often referred to as boundary or edge-finding methods, they all have as their goal the generation of a meaningful ROI. One approach is so common that it bears explanation. Often a first or second derivative calculation is used to help define




FIGURE 5. Different types of ROIs. (A) A circular ROI placed around the left ventricle in a planar blood pool study. (B) Hand-traced ROI drawn around the left ventricle. (C) ROIs generated using a threshold value of 70% of the image count maximum. Note that three ROIs are created: one around the left ventricle, one around the right ventricle and one around the spleen. (D) Isocontours generated using thresholds at 40%, 50%, 60%, 70% and 80% of the image count maximum. Note that some of the contours (at 70%) correspond to the ones shown in (C).

edges. The first derivative is also called the gradient; the second derivative is also called the Laplacian.

These operations have stringent mathematical definitions, but they are very simple concepts in practice. Both calculations analyze how count values change from pixel to pixel. Large count changes indicate the boundary of a structure, and both first and second derivatives are looking for these large changes. In truth, however, nuclear medicine images are not of such high quality that these derivative techniques often work perfectly. Usually, additional knowledge about anatomy is somehow encoded into the method to help find appropriate edges, and frequently the user is asked to intervene or edit the final ROI.

Operations on Planar ROIs

In one sense, all standard nuclear medicine processing techniques start with an ROI. Since the image almost always includes background structures and objects that are not specifically being studied, the first step is usually to isolate the important area using an ROI. Further quantitation generally involves operations on the ROI boundary itself or on the portion of the image that it encloses.

ROI Statistics

The simplest thing to do with an ROI is to sum the counts in the pixels it encloses. Additionally, mean counts, maximum and minimum counts, and the standard deviation of the counts can be computed. Sometimes the shape of an ROI is important. The area it encloses can be calculated by summing the number of pixels surrounded by the region and multiplying that by the pixel sizes:

$$\text{ROI area} = \text{No. of pixels in ROI} \times \text{pixel size} \times \text{pixel size.} \quad \text{Eq. 1}$$

Similarly, the width or height of an ROI can be measured by computing the width or height of the ROI in pixels and multiplying by the pixel size:

$$\text{ROI width} = \text{No. of pixels along ROI width} \times \text{pixel size.} \quad \text{Eq. 2}$$

Circumferential analysis is a technique used when the ROI is nearly convex or oval, without notches or irregular portions in its border. The center of the ROI is computed, and radii are drawn from this center point out to the ROI boundaries at various angles around the ROI. The length of the radii at each angle can be used to analyze the shape of the ROI. In addition, the mean or maximal count levels encountered along each radius can be used as a regional measure of function. Circumferential analysis is most often used in cardiac imaging, and specific examples will be discussed in more detail there.

Background Subtraction

ROIs are often used as a first step for background subtraction in planar images. Recall that planar images are summations of counts collected through the entire thickness of the body. Structures behind and in front of the organ of interest contribute to the number of counts in an ROI placed around that organ. These additional structures which are not of interest in the study are part of the background. If they extend beyond the organ of interest, then a second ROI can be placed near the original ROI to cover only background structures.

The second, or background, ROI will contain counts emanating *only* from background regions. Then the background-counts-per-pixel is calculated by computing the average count value within the background ROI. This average background count is subtracted from every pixel in the organ ROI to give the true number of counts in the organ of interest at every point; this operation is called background subtraction. The total number of counts within the organ is computed by summing up these background-subtracted organ ROI counts.

Sometimes the background ROI is defined by enlarging the organ ROI by a small amount and considering the doughnut-like region between that larger ROI and the organ ROI as the background. Other programs use simple, rectangular ROIs drawn by the user for background subtraction. More complicated methods use mathematical functions to predict how the background might change from one portion of the image to another. The primary concern is that any background ROI be truly representative of the background structures that contribute counts to organ ROI. Often, that task requires an expert, the technologist.

Time-Activity Curves

Dynamic and gated studies are used to investigate how the amount of radiopharmaceutical, or how the shape of the organ, changes with time. Statistics generated from ROIs placed on the same organ at various frames of a dynamic




FIGURE 6. A time-activity curve with an exponential function (dotted line) fit to its tail. The rate constant of the exponential function is 0.20.

study can be charted to display such changes graphically. The time-activity curve (TAC) is the most common of these graphs. In the time-activity curve, total or maximum counts in an organ ROI are plotted against time or frame number.

Often, what should happen to activity within a normal organ is known. For example, pharmaceuticals are cleared from the kidney at a rate related to the kidney's physiological state, as well as to the concentration of the pharmaceutical in that organ. This knowledge of what a time-activity curve should look like allows us to fit a mathematical function to the curve; if a common function such as a line or an exponential is used, analysis of the curve's shape is fairly simple. Different kinds of physiological parameters can then be extracted from this curve-shape analysis.

The most common types of functions fitted to time-activity curves are lines, exponentials and cosines. Washout or clearance often happens exponentially; that is,

$$\text{Counts in ROI at time } t = \text{Maximum counts} \times e^{(-k \times t)} \quad \text{Eq. 3}$$

The ability to fit this kind of function to a time-activity curve should be provided by nuclear medicine software; the fitting process involves finding an appropriate value for "k" in Equation 3, so that the resulting function closely matches the time-activity curve. This k is generally called a rate constant, and it describes how fast the time-activity curve falls off toward zero in the later frames of a dynamic study. Thus, it describes how quickly a pharmaceutical is cleared or washes out of an organ. A low value of k usually indicates that an organ is not functioning properly. Figure 6 shows a typical time-activity curve with a final exponentially decaying washout portion. An exponential in the shape of Equation 3 has been fitted to this portion of the time-activity curve, and the rate constant is reported as an exponent.

Cosine functions are generally fitted to time-activity curves created from ROIs placed over the heart chambers or

portions of the chambers in gated cardiac studies. The cosine is a cyclic wave function, so it approximates the cyclic heartbeat. The cosine takes the form:

$$\begin{aligned} \text{Counts in ROI at time } t &= (\text{Maximum Counts in ROI}) \\ &\times \text{cosine}(t/\text{heart rate} + \text{phase}). \quad \text{Eq. 4} \end{aligned}$$

Here again, the fitting of the function is not something that should concern the technologist; rather, the nuclear medicine computer should provide this capability. The result of the fit is the value of the phase parameter in Equation 4. This is called the phase angle, and it indicates when in the cycle the frame with the maximum ROI counts occurs. Note that in a normal left ventricle, the phase angle should be 0, indicating that the first frame of the study has the maximum counts. This is because the gating starts with the QRS complex, at end diastole, when the left ventricle is at its largest. Thus, an ROI placed over the left ventricle should contain more counts at end diastole than at any other frame. Phase angles are generally used in cardiac studies to determine if a chamber of the heart, or a region within a chamber, contracts abnormally, or out of phase, with the rest of the organ.

A few additional parameters are frequently extracted from the time-activity curve. The time-to-peak measures the time it takes for the time-activity curve to reach its maximum value after injection of the radiopharmaceutical. "Upslope" describes the angle that the initial portion of the time-activity curve makes with the horizontal axis—that is, the slope of the curve during uptake—and is expressed in counts/min (or counts/sec). The half-time, or $T_{1/2}$, is the time required for the time-activity curve to fall from its maximum to half of its maximum, that is, the time required for half of the tracer to washout or clear from the organ.

CLINICAL IMAGE PROCESSING TECHNIQUES

Multiple Gated Blood Pool Studies

The multiple gated blood pool study (MUGA) is a gated planar view of the blood within the heart chambers and is frequently used to compute ejection fraction (EF) (5). ROIs are drawn around a background area and then around the left ventricle (LV) at end diastole (ED) and end systole (ES). The counts in the LV ROI minus the average background counts are proportional to the LV chamber volume. Therefore, ejection fraction can be calculated as

$$\begin{aligned} \text{Ejection Fraction} &= \\ &\frac{(\text{Background-subtracted LV counts at ED} \\ &- \text{Background-subtracted LV counts at ES})}{\text{Background-subtracted LV counts at ED}} \quad \text{Eq. 5} \end{aligned}$$

The actual implementation of this processing technique can vary greatly from system to system. It can all be done manually, with the technologist deciding which of the frames are ED and ES, drawing the three ROIs, and calculating EF




FIGURE 7. Ejection fraction calculations from a MUGA study. Upper left image shows a user-drawn ROI surrounding the left ventricle (outer boundary) in the end-diastolic frame. This ROI is used as an aid in gradient-based automatic left ventricular edge detection. This automatically generated boundary is the inner ROI. Upper right image shows the automatically created left ventricular edge in the end-systolic frame. A user-drawn background ROI is shown to the lower right of the left ventricle. In the lower left portion of the figure, the left ventricular chamber counts are plotted using a time-activity curve. From this information, the ejection fraction is computed as 45%.

based on the total ROI counts. There are entirely automatic, or nearly entirely automatic, programs that find LV boundaries at all frames and create a time-activity curve that describes LV ROI count changes over the entire cardiac cycle. ED and ES frames can then be determined as the frames with the highest and lowest LV counts, respectively. With a background ROI, then, EF can be directly calculated. Figure 7 shows some of the steps of an ejection fraction calculation from a typical MUGA study.

Segmental wall motion can also be evaluated from MUGA studies (6). Since the blood in the LV chamber is contiguous with the myocardium, the way the chamber blood moves is equivalent to how the endocardium moves. Once again, wall motion analysis requires finding the boundaries of the LV, in at least the ED and ES frames. A circumferential analysis technique is then used to measure wall motion. The center of the LV ROI is automatically determined, and radii are drawn from this center to the boundary at a number of different angles. The change in radius length from ED to ES at each angle is a measure of how the wall is moving at that angle. Often, these motion values are expressed relatively, in percentages:

% segmental motion at an angle

$$= \frac{\text{ED radius at that angle} - \text{ES radius at that angle}}{\text{ED radius at that angle}} \times 100\%. \quad \text{Eq. 6}$$

Figure 8 shows one implementation of segmental wall motion computations from a typical MUGA study. Note in particular




FIGURE 8. Segmental wall motion calculations from the same MUGA study shown in Figure 7. In the top row, from left to right, are shown (A) the end-diastolic frame with the left ventricular ROI and its center point; (B) the end-systolic frame with the left ventricular ROI; and (C) the end-diastolic frame with the left ventricular ROI and six radial segments drawn from the ROI center to its edges. The center portion of the figure shows the relative lengths of the radial segments at end diastole (outer boundary) and end systole (inner dotted boundary.) This patient has somewhat reduced septal radial shortening (segments 4 and 5).

the calculation of segmental wall motion by measuring LV radii at ED and ES.

Numerous other techniques have been proposed for procuring additional data from MUGA studies. Absolute left ventricular volumes have been calculated by measuring the size of the ROI placed around that chamber (7). A simple version of this method can be employed by assuming that the left ventricle is a perfect ellipsoid. The LV volume can be calculated from a lateral planar MUGA view as:

$$\text{LV volume} = \frac{4}{3} \times \pi \times \text{LV ROI height} \times \text{LV ROI width} \times \text{LV ROI width}, \quad \text{Eq. 7}$$

where the width and height have been converted from pixels to centimeters.

As mentioned in the discussion about time-activity curves, the left ventricle can be analyzed using phase analysis (8). The LV ROI can be divided into sectors, and each sector can be analyzed with a time-activity curve fitted to a cosine function. Any sector with an abnormally large phase angle is probably contracting abnormally. These sectors may also be used as the ROIs in Equation 5, so that an ejection fraction is associated with each. This regional ejection fraction should, like segmental wall motion analysis, correspond with left ventricular functioning.

Cardiac First-Pass Studies

In cardiac first-pass studies, a bolus of high-activity is rapidly injected intravenously, and images are obtained as that bolus passes through the cardiac chambers. First-pass studies allow images to be collected while the bolus is contained entirely in the LV or right ventricle (RV), thus

eliminating problems with background counts or overlapping chambers, such as those which occur in MUGA studies.

First-pass studies must be acquired using a high-rate dynamic frame mode, or using list mode in order for the pictures to show the bolus as it moves through each chamber separately. Following acquisition, the technologist generally chooses a frame in which the LV is clearly seen, and draws an ROI around it. There are, of course, some available automatic programs to create this ROI. A time-activity curve is generated from this ROI as it is applied to a sum of 3–6 contiguous cardiac cycles in the dynamic study. In a manner similar to MUGA analysis, both EF and segmental wall motion can be calculated from the ROIs at ED and ES (9). First-pass studies also allow the RV to be easily separated from the LV, since the bolus is effective only in the RV during some frames. Therefore, right ventricular EF can also be calculated (10). Unfortunately, the irregular, nonconvex shape of the RV makes wall motion analysis using circumferential profiles difficult.

First-pass studies are also frequently used to analyze shunts, where some of the blood from the LV is pumped back into the RV instead of into the systemic circulation at each heartbeat (11). The amount of blood shunted back into the RV can be evaluated by calculating the amount of activity seen in the lungs. In normal studies, a lung time-activity curve from a first-pass study has one large peak after the bolus is initially pumped out of the RV. Then there is a smaller peak in about 8 sec after the bolus has been pumped back through LV, the systemic circulation and the RV. (This second peak is smaller because as time goes on, the bolus starts mixing with nonradioactive blood, and becomes less concentrated.)

In patients with left-to-right shunts, however, the initial peak of the lung time-activity curve is followed immediately by a second peak, caused by blood shunting from the LV to the RV and being pumped directly back into the lungs. Shunts are evaluated by fitting a mathematical function called a gamma-variate curve to the first peak of the lung time-activity curve, and then subtracting this fitted function from the original time-activity curve. Another gamma variate function is fitted to the second peak seen in this subtracted time-activity curve. The total area under the first gamma variate curve is related to the volume of flow through the lungs; the total area under the second gamma variate curve is related to the volume of shunted blood. The standard parameter used for evaluation of shunts via first-pass studies is the ratio of systemic-to-pulmonary flow, where systemic flow is equal to the initial pulmonary blood flow minus what was shunted. Therefore,

$$\frac{\text{pulmonary flow}}{\text{systemic flow}} = \frac{\text{Area under first fitted curve}}{\text{Area under first fitted curve} - \text{Area under second fitted curve.}}$$

Eq. 8

Calculation of this pulmonary-to-systemic ratio is known generically as the Qp/Qs method of shunt detection. Another method, which has been used in various forms since the early 1960s, and which is still offered in clinical nuclear software, is the C2/C1 method, in which the time-to-peak activity is determined from a lung time-activity curve. A ratio consisting of the counts at twice the time-to-peak (in the numerator) and the peak counts (in the denominator) is calculated. This result is expressed as a percentage, and values greater than 35% are considered abnormal.

Cardiac Perfusion Studies

Planar perfusion images show pictures of blood flow to the myocardium. These images document the existence of ischemic heart disease or prior myocardial infarctions. Thallium-201 was the primary agent used for perfusion studies up until a few years ago; currently there are two ^{99m}Tc agents also being used: ^{99m}Tc-sestamibi and ^{99m}Tc-teboroxime. All of these radiopharmaceuticals are delivered to the myocardium in proportion to perfusion, so poor or no uptake is consistent with stenosed coronary arteries or dead muscle. The main technique used to process these studies is a circumferential profile analysis that evaluates variations in perfusion around the circumference of the LV (12,13). This analysis is generally done for images acquired under a resting state, as well as for those acquired after a patient was injected with the pharmaceutical at peak exercise or pharmacological stress.

Perfusion analysis again starts with an ROI automatically or interactively drawn around the epicardial boundary of the LV, often a simple circle, since absolute position is not so important as the fact that the ROI extends beyond the epicardium. From the center of the ROI, radii are drawn at different angles out to the ROI boundary. Counts encountered along each radius are evaluated; the highest count value is assumed to be the peak perfusion value at that angle. The peak counts are often graphed versus the angle at which they were located, thus creating a circumferential profile curve. Usually, this curve is compared against an average circumferential profile, generated by performing the analysis on many normal volunteers.

When the patient's curve (perfusion) at a certain angle is significantly lower than the normal curve (perfusion), then the patient is considered to have a perfusion defect at that location. Rest and stress perfusion are often compared; simply put, a perfusion defect seen at a specific area in a stress study that is also present in a rest study generally indicates dead myocardium, or myocardial infarction at that location. A defect at stress that disappears at rest is generally indicative of living tissue that simply does not receive enough blood flow during stress because of coronary artery stenosis. Such reversible defects may be treatable using angioplasty or coronary artery bypass grafts. Figure 9 shows the results of processing a planar ²⁰¹Tl study using circumferential analysis.

There are many different variations on this general technique. The details of perfusion acquisition and analysis techniques used in a specific institution should be understood,




FIGURE 9. Circumferential analysis of a planar ^{201}Tl cardiac perfusion study. The top row shows 45° LAO views of the stress and delayed studies; on their right are displays of the highlighted maximal count pixels detected around the left ventricular circumference during analysis. The bottom left graph is a circumferential profile of stress perfusion (top curve) and the expected normal stress perfusion values (bottom curve.) The bottom right graph is the circumferential profile (top curve) and the expected normal washout values (bottom curve.) This patient's stress perfusion and washout values are greater than normal at all angles, so this study is considered normal.

since they affect the appearance of the image and the shape of the circumferential profiles. The basic ideas are very similar, however. A patient's myocardial perfusion can be evaluated by comparing the relative counts seen in the LV at rest and stress to that of a normal population.

Renal Studies

Nuclear renal studies today typically involve assessment of both structure and function. Careful visual inspection of the renogram and the static images is used to glean anatomic information such as differential function between right and left kidneys, decreased arterial blood flow, presence of mass lesions and obstruction of urinary emptying. While numerous renal studies involve only blood sampling and well-counting, here we will focus on those which involve image processing.

The basic tools of image processing—ROIs, their associated time-activity curves, and image subtraction—are used in the calculation of quantitative results from dynamic renal images. A number of corrections can be applied to the data in the course of processing in order to account for scan table attenuation, camera dead time, renal depth, scatter and attenuation within the patient and other factors. These calculations can often be quite complex. Typical image processing steps include the following:

1. Raw dynamic image data are summed or reframed so that the arrival time of activity in the kidneys can be determined.
2. ROIs are defined for the whole kidneys and, optionally, the cortical areas.

3. ROIs representing background are user- or computer-defined. A number of techniques have been used to improve and automate background selection. These include count thresholding to optimize the choice of background area and using a standard ROI shape, such as a box or ellipse. It is important to exclude kidney counts from the background, to sample all areas that contribute counts to the kidney ROI and to be consistent.
4. Background-corrected time-activity curves are plotted from the dynamic image data and compared to the expected curve shape and clearance rate.

Glomerular filtration rate (GFR) can be estimated using $^{99\text{m}}\text{Tc}$ -labeled diethylenetriaminepentaacetic acid (DTPA) and dynamic imaging. The GFR indicates how well substances are filtered out from the blood into the kidney. Iodine-131-labeled orthiodohippurate (OIH) has been used for many years in studies of kidney function to obtain a measure of effective renal plasma flow (ERPF). ERPF indicates how well OIH is delivered to the kidneys via the blood.

Relative ERPF rates are determined by calculating the early uptake of OIH in each kidney. ROIs must be created to compute the total counts taken up in both kidneys during the 1–2- or 2–3-min dynamic acquisition. The necessary ROIs include those two surrounding each of the kidneys, as well as a right and left background ROI. Relative ERPF for the right kidney is then calculated as:

Relative right ERPF

$$= \frac{\text{Background-subtracted counts in right kidney ROI}}{\text{Total background-subtracted counts in both right and left kidney ROIs}}$$

Eq. 9

and relative ERPF for the left kidney is calculated as:

Relative left ERPF

$$= \frac{\text{Background-subtracted counts in left kidney ROI}}{\text{Total background-subtracted counts in both right and left kidney ROIs}}$$

Eq. 10

Total ERPF (as opposed to *relative* ERPF) can be determined using empirical equations that attempt to quantify the amount of radioactivity in the kidneys (14). The amount of radioactivity detected in the kidney depends on how much activity was injected, the pharmaceutical uptake and the amount the photons are attenuated as they emanate from the body. The injected activity can be computed by counting the emissions from the syringe before and after injection. Uptake can be measured using background-subtracted ROIs

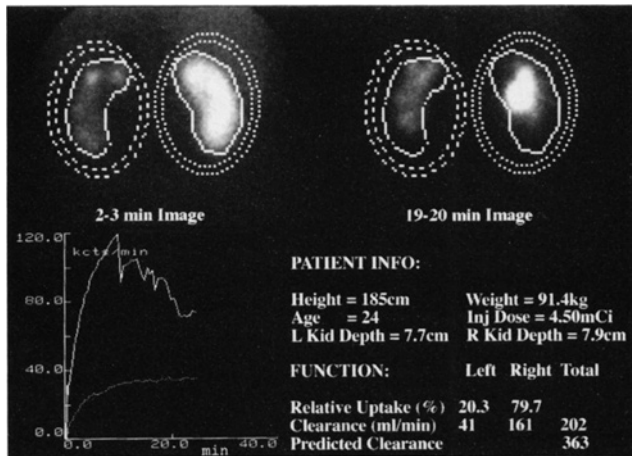


FIGURE 10. Renogram results using MAG3. Upper left image shows the summed dynamic frames from 2–3 min. The total kidney ROI has been drawn, and a doughnut-shaped background ROI has been placed outside of the kidney ROI. The upper right image shows the summed dynamic frames from 19–20 min with similar ROIs. Note that both uptake and clearance of MAG3 is badly reduced in the left kidney. Bottom left are the time-activity curves of the background-subtracted left and right kidney ROIs. The upper curve corresponds to the right kidney, and the lower curve corresponds to the left kidney. Once again, the shapes of the curves indicate abnormal left kidney function. Patient information required for calculating the ERPF, and those results, are shown in the lower right.

over the kidneys. Attenuation can be estimated by multiplying an approximate attenuation factor with each kidney's depth from the body surface. Finally, kidney depth can be estimated according to the patient's height and weight.

$$\text{TOTAL ERPF} = 5.029 \times 0.36987 \times \text{OIH uptake} - (0.000231476 \times \text{OIH uptake} \times \text{OIH uptake}), \quad \text{Eq. 11}$$

where

OIH uptake

$$= \frac{\text{Background-subtracted kidney counts} \times Y \times Y \times 100}{1\text{-min count of injected radionuclide}},$$

and

$$Y = 13.3X + 0.7.$$

Here, Y is equivalent to kidney depth, and X is the patient's weight in kilograms divided by his height in centimeters.

Figure 10 shows one renogram processing display screen. Left and right kidney and background ROIs are shown. A calculation similar to Equation 11 has been performed to determine ERPF in this patient study.

Relative GFR can be determined using DTPA with methods similar to those employed to calculate relative ERPF. The accumulation of DTPA during the first 1–3 min of scan-

ning is proportional to GFR. Therefore, equations analogous to relative right ERPF and relative left ERPF can be employed to compute relative GFR. Total GFR can be measured using Gates' equation (15):

$$\text{GFR} = (\% \text{ renal uptake of DTPA}) \times 9.81270 - 6.82519, \quad \text{Eq. 12}$$

where % renal uptake of DTPA =

Right kidney background-corrected counts

$$\frac{e^{-0.153 \times Y}}{\text{preinjection counts}}$$

$$+ \frac{\text{Left kidney background-corrected counts}}{e^{-0.153 \times Y}},$$

preinjection counts

and Y is once again kidney depth.

By the end of 1993, approximately 50% of the renal scans performed in the U.S. were being done with MAG3 (mercaptoacetyl triglycine), a pharmaceutical which can be labeled with ^{99m}Tc . Because of the superior imaging properties of ^{99m}Tc as compared to ^{131}I , and because MAG3 is filtered and excreted similarly to OIH, this new tracer is an effective replacement for the OIH/DTPA combination. Methods are available for calculating MAG3 clearance using blood sampling techniques, as well as a purely camera-based technique.

Whether the function study is done using OIH or MAG3, (or one of several other radiopharmaceuticals currently available), renogram curves are an important qualitative way to assess blood flow and function. In addition, analysis of the curves after administration of the diuretic captopril is used in the study of patients suspected of having renovascular hypertension.

Gastric Emptying

Gastric emptying studies are concerned with finding the emptying half-time of the stomach contents, since certain disease processes can result in prolonged emptying times and other symptoms. Solid-phase emptying can be measured using ^{99m}Tc introduced into a solid meal, such as chicken liver, eggs or oatmeal. Liquid-phase emptying, which occurs much faster, is measured using a liquid such as orange juice. The two studies can be done simultaneously using the dual-isotope technique. In this case, ^{111}In might be the tracer used for the liquid phase.

Anterior and posterior static images are acquired at regular intervals for up to 2 hr. A single ROI is drawn for the anterior images, and another is drawn for the posterior images. The geometric mean counts are determined for each pair of images, and a time-activity curve is constructed, from which the half-time is calculated from the curve slope.

Parathyroid Imaging

Parathyroid imaging is used in the detection of parathyroid adenoma and parathyroid hyperplasia. The basic method involves a dual-isotope acquisition. This may be any combination of tracers, one of which is taken up by the thyroid and the other by the thyroid and parathyroid. The tracer combination is typically ^{99m}Tc (as pertechnetate) and ^{201}Tl , or a combination of ^{99m}Tc -sestamibi and ^{111}In . The images from the lower energy tracer should be corrected for downscatter from the higher energy tracer.

An ROI is generated for the thyroid, and smaller ROIs are placed within the area of normal thyroid activity to determine mean counts in each pole of the gland. The ratio of these counts for each of the two tracers is used to rescale, or normalize, the two images so that their new count maxima are identical. Following count normalization, the thyroid image is subtracted from the thyroid and parathyroid image. The resulting image should delineate areas of abnormal parathyroid uptake. There have been a number of modifications to the basic procedure described above, including changing the order in which the tracers are given, the relative doses, and the time interval between doses, but the basic image processing steps remain the same.

TOMOGRAPHIC IMAGE PROCESSING TECHNIQUES

Acquisition

If enough planar views are acquired at various angles around the patient, then these views, or projections, can be mathematically combined or reconstructed to create a three-dimensional stack of slices showing the distribution of radioactivity at every point in the body. This stack is referred to as a tomogram. One slice of the tomographic stack can be created from each row of the planar projection image set. Tomography eliminates some of the problems of planar imaging, such as occlusion by objects in front of or behind the organ of interest.

Compared to a two-dimensional view of the organ, a three-dimensional slice-by-slice tomogram allows every portion of the object to be completely seen and analyzed. Much of tomographic processing mirrors that of planar processing, since each tomographic slice is just a digital image similar to each planar view. Like the planar views, each tomographic slice has an associated matrix size and pixel size. In addition, tomographic slices have a thickness, which is usually equal to the pixel size in the vertical direction of the original projection. Pixel sizes in the tomogram depend on the pixel size in the horizontal direction of the projection.

Like planar views, tomographic acquisitions can be static, dynamic, gated and/or dual-isotope. Because a tomogram requires obtaining about 64 projection views, tomographic acquisition takes much more time than a planar acquisition. For this reason, dynamic acquisitions are not frequently acquired. However, with new multiheaded rotating gamma cameras, dynamic tomography is becoming feasible and both brain and cardiac perfusion have been evaluated using dy-

amic tomographic acquisitions. Gated tomography is possible with ^{99m}Tc cardiac agents, including the applications of blood pool and perfusion studies. Research in the processing of gated cardiac perfusion studies in particular indicates that gated tomography may well become a commonplace clinical tool because of its superiority over static perfusion imaging.

TOMOGRAPHIC RECONSTRUCTION

Filtered Backprojection

Reconstruction is the process of creating the three-dimensional stack of tomographic slices by combining a set of multiple two-dimensional planar views. The most widely available technique for performing this task is called filtered backprojection. Backprojection can be most easily understood by considering how a point source is reconstructed.

Figure 11 shows how projections of the point source are acquired and used in backprojection to create a reconstruction of a slice through the original point. In each planar view, the point source appears not as a point, but rather as a blurred circle. A horizontal profile drawn across one row of the projection, through this blurred circle, represents the counts from the projection that will contribute to one slice of the tomogram. In backprojection, the counts in each projection's profile are spread out over the slice from which they were acquired perpendicularly to the original planar view. Notice how the final backprojection, after all planar views have been processed, is very blurry, and has a so-called "star artifact" associated with it. The star artifact is characterized by the rays seen emanating from the reconstructed point.

The blurring and the star artifact can be eliminated with filtering; this is the filtering portion of filtered backprojection. Two filters are actually used. First, a smoothing filter, similar to the ones described in the general processing section, is applied to each planar projection, smoothing it in both the vertical and horizontal directions. This filter helps reduce noise, which would otherwise be greatly amplified by the second filter and degrade the image quality too much. The second filter is called a ramp filter which performs a "sharpening" operation that gets rid of the blurring and star artifact caused by backprojection. More detailed information on filters as they are used in reconstruction is given by Galt et al. (16).

Some systems refer to the first filter as a prefilter; the second filter is often called the reconstruction filter. The prefilter is usually a Butterworth or Hanning filter, although many others are available. Different types of filters smooth the images in slightly different ways, but there are some common characteristics. Both Butterworth and Hanning filters are described by cutoff or critical frequency values, which determine *how* smooth the image becomes, in a manner similar to the size of the smoothing kernel. While cutoff and critical frequency are two different parameters that describe the shape of the filter, they are similar in their most important aspects. The lower the cutoff or critical frequency value, the smoother the resulting image.

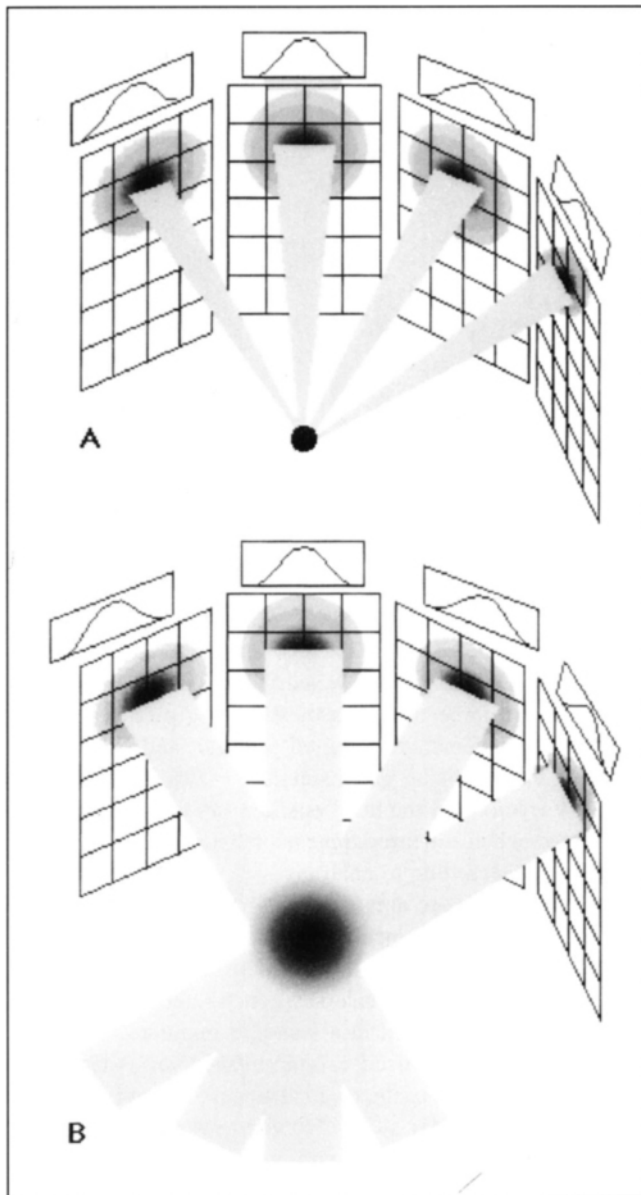


FIGURE 11. (A) Four projections created from a point source. Profiles taken through one row of the projection plane shows that the point has been blurred. (B) Backprojection from one row of the projections into the reconstruction slice. Blurring persists into the reconstruction and a star-shaped artifact is created.

Butterworth filters also have another parameter, the order, which also affects the final image, but not nearly so much as the critical frequency value. In general, the lower the order, the smoother the resulting image. In contrast to the prefilter, there is only one commonly used reconstruction filter, the ramp filter. Ramp filters have no variable parameters. Figure 12 shows the effects of using only a ramp filter during reconstruction of one transverse slice, and then the effects of using both a ramp filter and a Butterworth smoothing filter of various cutoff values.

Some systems actually combine the smoothing and sharpening filters into one operation; this is less common than it used to be. The documentation for such combined filter




FIGURE 12. (A) One reconstructed slice of a cardiac perfusion study using only a ramp filter. (B) Same slice reconstructed using a Butterworth filter of critical frequency = 0.6/cm and order = 10 prior to the ramp filter. (C) Same as (B), but with critical frequency = 0.4/cm; order = 10. (D) Same as (B), but with critical frequency = 0.2/cm; order = 10.

systems may never mention a ramp filter, but if a system has only one filtering operation, it no doubt incorporates a ramp filter. In contrast, some systems will call combined filters “ramp-Butterworth” or “ramp-Hanning.” Combined filters have associated cutoff values, but in fact the cutoffs refer only to the smoothing Hanning and Butterworth filters. The cutoffs in combined filters have the same effect as in separate prefilters: the lower the cutoff, the smoother the resulting image.

Attenuation Correction

Emitted gamma rays are often absorbed by nuclei of atoms within the body. Gamma rays coming from near the surface of the body have less chance to encounter other atoms since they almost immediately enter into the air. Gamma rays coming from the center of the body have a much larger chance of encountering bones and muscle and of being absorbed in those tissues. This absorption is called attenuation. Because of attenuation, most tomograms tend to show fewer counts coming from the center of each transverse slice than from the edges of the slice.

Attenuation Coefficients. The amount by which each type of tissue may attenuate a certain energy photon is known. The factor that measures this attenuation amount is called the attenuation coefficient. Its magnitude depends upon the energy of the photon and the density (specifically, the electron density) of the tissue. Higher-energy photons are less likely to be absorbed than lower-energy photons. Also, as tissue density increases, photon attenuation within that tissue increases. For example, bones have the highest attenuation coefficient of any natural material in the body.

Attenuation Map. The only way to accurately correct for the loss of attenuated photons is to use an attenuation map of all the structures in the body, indicating the position of bones, muscle and all other attenuating tissues. Attenuation

maps can be generated using transmission scans such as x-ray transmission computed tomography, which by definition creates pictures whose intensity levels are related to tissue densities. In practice, this approach is rarely used. Instead, ad hoc methods are employed, most of which make very simplistic assumptions about the attenuation map. Often it is assumed that the attenuation coefficients within the body are all the same constant value. This assumption is fairly good for abdominal tomography, just fair for head tomography, and, unfortunately, particularly poor for the thorax.

The primary processing step required for generating an attenuation map is to create an ROI that surrounds the outer boundary of the body. This can be performed with threshold techniques or with interactive ROI tracing. The area within the ROI is filled with the appropriate attenuation coefficients; usually, a simple constant-attenuation coefficient map is created. This map is then used as the primary input to an attenuation correction algorithm.

While attenuation correction methods are often complex, the ultimate goal of each of them is to create an image whose count levels more closely represent the actual distribution of the radionuclide in each tissue. To do this, they add counts to the original tomographic slice based on the attenuation coefficient and the distance of each pixel from the body surface. Both of these pieces of information are obtained from the attenuation map, so its accuracy is of paramount importance. The body outline should follow the true skin contours as closely as possible.

GENERAL TOMOGRAPHIC IMAGE PROCESSING

Oblique Reorientation

Frequently, the three-dimensional blocks of tomographic data are reoriented or resliced at angles to create new slice-by-slice views. The reslicing can also be termed “resampling” or “reformatting.” Originally, each tomographic slice is perpendicular to the projection planes, and since the projections are parallel to the axis of the body, the slices are transaxial or transverse. Transverse slices cut across the length or height of the person being imaged. When these slices are stacked, a three-dimensional set or block of data is created that shows radionuclide uptake at every point in the area of the body being imaged. Figure 13 demonstrates how the slices stack up to form a three-dimensional set. This three-dimensional stack can be sliced vertically across its width or depth to create coronal or sagittal slices; a demonstration of this is also shown in Figure 13. However, the block can also be resliced at any angle the user desires, and the new sections are termed “oblique” slices.

Often, the three-dimensional stack of transverse brain tomographic slices is recut so the new transaxial slices are actually parallel to the sylvian fissure of the brain. The user is asked to draw a line on the original sagittal slices to indicate the angle of the sylvian fissure, and slices parallel to this line are created. These new slices are then once again stacked to create a new three-dimensional volume which can




FIGURE 13. A brain perfusion tomogram is a three-dimensional stack of slices. Some transverse sections from this stack are shown in the upper right. One sagittal and one coronal slice from the stack are shown in the bottom image; note that these two slices are both perpendicular to the transverse sections shown in the upper left.

be cut vertically to create new near-sagittal and near-coronal views. These three new sets of slices are often called re-oriented transverse, re-oriented sagittal and re-oriented coronal. This reslicing gives standard views of the imaged organ. Variations in the head position between patients are eliminated when all three-dimensional tomographic stacks are resliced according to anatomy.

Cardiac images are almost always resliced in a very specific way. Because the position of the heart within the body can vary greatly from person to person, reslicing the three-dimensional stack of transverse sections allows every cardiac image to be viewed in a standard manner. Therefore, the appearance of resliced cardiac images varies less from patient to patient than the original transverse images do.

While the particular steps involved in creating these specific slices depend on the manufacturer of the nuclear medicine system used, most follow a general pattern. The original transverse slices are displayed, and the user is asked to identify the left ventricular long-axis (the line extending from the apex of the left ventricle to the aortic valve plane). The block of data is resliced in planes perpendicular to the original transverse slices and parallel to the denoted long-axis to create near-sagittal sections. These new sections are called vertical long-axis (VLA) slices. The user is then asked to once again identify the long-axis on these VLA slices. The VLA stack is resliced along planes perpendicular to VLA sections and parallel to the denoted long-axis.

These new sections, which cut horizontally through the left and right ventricles, are termed horizontal long-axis (HLA) slices. A second reslicing of the VLA stack is performed, this time using a plane perpendicular to both the stack and to the denoted long-axis. These sections cut through the left ventricle perpendicularly to the long-axis and are called short-axis slices. Figure 14 shows these various standard cardiac slices, along with the steps required to create them.




FIGURE 14. Creation of oblique cardiac sections. (Row 1) Original transverse slices with user-drawn long-axis. (Row 2) Vertical long-axis slices generated from Row 1 with a user-drawn long-axis. (Row 3) Horizontal long-axis slices generated from Row 2. (Row 4) Short-axis slices generated from Row 2.

The various angles of the user-drawn lines and the lines used to create oblique slices are often called by various terms by different manufacturers. Some systems use astronomical or aeronautical terms to denote the angles, such as "azimuth," "pitch" and "roll." Others may use no specific term at all. In some implementations, each line must be drawn starting at the apex and ending at the base, while other implementations have no such limitations. Only the manufacturer can give explicit details as to how oblique reslicing must be performed on a given system. While there may be standard defined reslicing planes, there is definitely no standard defined method for creating them.

Display

Tomographic images are generally displayed as a set of slices, similar to how a set of multiple planar views would be displayed. A scale factor can be applied or the color map can be changed so the appearance of a slice changes. Usually, multiple slices are all scaled using the same multiplicative factor, so colors or intensities in different slices can be easily compared. Most systems allow each slice to be scaled separately, however, so the brightness or color of the same original count value can change drastically from slice to slice. While both scaling modes are useful in various cases, it is important to know how the images on the current display are scaled in order to analyze them properly.

Regions of Interest

ROIs are created on tomographic slices in the same manner as they are created on planar views. It should be noted, however, that many structures extend through more than one tomographic slice. Therefore, a complete analysis of an organ requires boundary finding or ROI drawing in every slice in which that organ appears. Thresholding is the only frequently used automatic ROI creation technique for simultaneously processing multiple slices. Generally, a user can

vary the threshold to separate the organ from background. The same threshold is applied to all slices. This will result in objects having closed ROIs surrounding their boundaries in every slice that they appear. Those ROIs which surround the same organ in multiple slices will be connected in the sense that the computer will know that they all refer to one object. When thresholding methods are not sufficient, the technologist is usually asked to draw the ROIs in every slice where an analysis is desired.

Operations on Tomographic ROIs

ROI Statistics. The total number of counts in an organ in a tomographic image can be calculated by summing the total counts in each ROI surrounding that organ in every slice. Organ volume can be calculated by summing the ROI area of the object in each slice and multiplying the sum by the slice thickness. Circumferential analysis is used, primarily on a slice-by-slice basis, with the results from consecutive slices presented together in the same report. This allows analysis of the perfusion of an entire organ, if the organ is fairly regular, like the LV or cortex of the brain.

Background Subtraction. Background subtraction is used less frequently with tomography than with planar studies. Tomograms solve the problem of occluding objects that contribute to the background of an organ of interest. In some cases, however, the radiopharmaceutical can be absorbed by structures other than, but very near to, the organ of interest. In this case, scatter and some remains of the reconstruction star artifact can tend to overlap onto the organ to be studied. These cases do demand background subtraction, which is generally carried out similarly in tomography as in planar imaging.

Time-Activity Curves. Time-activity curves are used in tomographic slices as they are in planar studies. Time-activity curves can be generated either from a two-dimensional ROI placed in a single tomographic slice, or from a three-dimensional collection of ROIs surrounding an organ that extends through several slices. With dynamic or gated studies, organ function can be analyzed by investigating how ROI statistics change over time.

SPECIFIC TOMOGRAPHIC IMAGE PROCESSING TECHNIQUES

Cardiac Perfusion Tomogram Processing

Tomographic perfusion studies created using ^{201}Tl or $^{99\text{m}}\text{Tc}$ -sestamibi are processed using a three-dimensional extension of planar processing techniques. Each short-axis slice of the rest and stress tomogram is analyzed circumferentially to find the maximum count (perfusion) value at numerous angles about the circumference of the left ventricular myocardium. The perfusion rates are compared to known normal values to determine if blood flow has been reduced at rest, stress or both.

There are numerous approaches for achieving this analysis, and specific implementations vary. Generally, the following steps are employed, as shown in Garcia et al. (17).

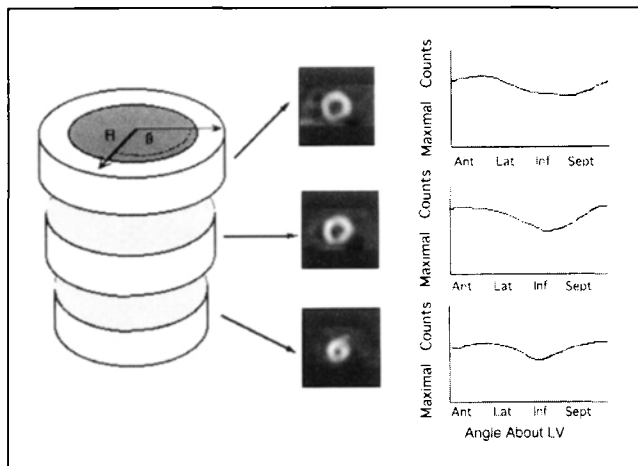


FIGURE 15. Diagram of cardiac perfusion analysis in tomography. Each left ventricular short-axis section is analyzed using a circumferential analysis technique.

The user denotes the base and apex of the left ventricle on the vertical or horizontal long-axis slices, using a mouse, light pen or trackball. This indicates the extent of short-axis slices that will be processed. On short-axis sections, the user designates the center of the left ventricle and, using a circle, indicates the outer boundary of the epicardium. This boundary is used to limit the search for maximal myocardial counts. In each short-axis slice, a search for maximum counts is performed at a number of angles, usually about 40, around the left ventricular center. By graphing the maximal counts versus the angle at which each was found, a circumferential profile is generated for each short-axis slice. This type of analysis is demonstrated in Figure 15.

Alternatively, the maximal count values can be color-coded and charted as a picture in ever-widening circles from left ventricular apex to base. This gives a "bull's-eye" or "polar map" display. In bull's-eye displays, colors near the center indicate apical counts; colors near the outer boundary indicate basal counts. As one travels from the top of the polar map around it in a clockwise direction (through the top, right, bottom and left sides), counts determined from the anterior, lateral, inferior and septal walls of the LV are displayed.

Various other calculations can be made from the raw count values that were determined from the short-axis slices. Stress and rest perfusion are compared to the known normal values; abnormal sectors can be denoted on the polar maps using a special color. Black is often used to denote abnormal sectors; thus, such polar maps are called "blackout" maps. Reversibility is a measure of the relative perfusion increase from the stress study to the rest study.

When stress perfusion is abnormal in a section of the left ventricle, but the rest perfusion in that section is normal, the area of abnormal stress perfusion is said to be reversible. The amount of reversibility of every part of the left ventricle can also be compared to normal values, and any LV portion that demonstrates abnormal reversibility can be displayed in a reversibility blackout map. In contrast, portions of the left

ventricle that demonstrate abnormal stress perfusion, but reverse to show normal perfusion at rest, may also be denoted using special colors in the polar map. Often, white is used to indicate abnormal stress perfusion that reverses at rest, and such maps are then called "whiteout" maps. Figure 16 shows the short-axis slices from rest and stress, of both a normal and an abnormal ^{99m}Tc -sestamibi study. The rest and stress polar maps, as well as the rest and stress blackout maps, are also shown.

As with almost all nuclear medicine processing methods, tomographic perfusion analysis has many different approaches. For instance, some techniques do not use circumferential analysis on those short-axis slices near the apex but sample this section of the left ventricle in a three-dimensional, rather than slice-by-slice, manner (18). Three-dimensional displays of left ventricular perfusion are becoming more common, to complement the bull's-eye displays. This may be helpful to eliminate the distortion of bull's-eye displays. Bull's-eyes tend to distort the location and shape of perfusion abnormalities, since apical sections are mapped to a small center portion, but basal sections are mapped to a very wide outer portion. However, the philosophy of almost all of the tomographic perfusion analysis methods are remarkably similar. Maximal counts in many left ventricular portions are compared to known normal values to analyze myocardial perfusion, with the ultimate goal of diagnosing the existence of coronary artery disease or myocardial scarring.

Tomographic Cerebral Blood Flow and Metabolism Processing

While cerebral blood flow and metabolism are not quantitated using a standard method across the many nuclear medicine systems, there are a number of general approaches that appear frequently. Technetium-99m-HMPAO is a commonly used brain perfusion agent for SPECT. It is delivered to the brain in proportion to blood flow and can indicate numerous conditions such as stroke, Alzheimer's disease and brain tumors. One frequently used analysis technique in these studies is to compare the counts seen in ROIs drawn over similar regions on both the right and left sides of the brain (Fig. 17). Sometimes, a circumferential analysis technique can be used to describe ROIs toward the outer boundary of the cortical gray matter; maximal counts in the ROI sectors spanning the circumference of the cortex are considered to be measures of perfusion (19).

While normal brain perfusion can vary greatly, most regions should be fairly symmetrically perfused in both left and right hemispheres. A large difference between maximal counts in ROIs drawn over the same structure in both the left and right hemispheres may indicate an abnormality in the lower-count ROI. Comparisons may also be made between regions drawn in different areas of the same hemispheres, when the normal relative perfusion rates are known. Cortical perfusion in a frontal lobe ROI can be compared with that in a temporal lobe ROI, for example. If normal studies have indicated that the frontal lobe generally has greater perfusion




FIGURE 16. (A) Stress (top) and rest (bottom) short-axis slices from a normal ^{99m}Tc -sestamibi cardiac perfusion study. Polar maps shown on the bottom of the figure are (top) stress and rest perfusion maps and (bottom) stress and rest perfusion blackout maps. In this case, no regions are blackened out, since none are abnormal. (B) Stress (top) and rest (bottom) short-axis slices from a ^{99m}Tc -sestamibi cardiac perfusion study of a patient with an irreversible inferior perfusion defect and a reversible anterior perfusion defect. The polar maps shown on the bottom of the figure are (top) stress and rest perfusion maps and (bottom) stress and rest perfusion blackout maps. In the stress blackout map on the left, note the defects in both the anterior and inferior walls. The anterior defect has disappeared in the rest blackout map on the right, indicating that it is a reversible defect. The inferior defect persists in the rest blackout map, so it is an irreversible defect.

than the temporal lobe, then this should be reflected in the maximal count levels seen in the two ROIs. These basic, relative perfusion comparisons are often the most sophisticated quantitative technique performed in brain perfusion tomography.

One perfusion study that is not so frequently encountered now as it was a decade or so ago is the ^{133}Xe cerebral blood flow study. This perfusion study is a dynamic tomographic imaging technique which measures the washout of radioactive xenon from the brain (20). The patient breathes the xenon gas until its concentration reaches equilibrium within the blood. Then, a number of tomographic acquisitions are done after the patient starts breathing room air once again. The manner in which the radioactive gas washes out of each region in the brain indicates the amount of blood flow within that region. Fairly complicated techniques are used to calculate the brain blood flow in ml/g at every pixel in the tomographic slices.




FIGURE 17. ROIs placed over various regions on superior (left) and mid-axial (right) transaxial slices of a cerebral perfusion tomogram. Maximal counts within in each ROI are compared from left to right hemispheres, to determine relative perfusion rates.

Xenon-133 brain blood flow imaging has never been widespread because of the difficulties with rapid and repeated dynamic tomographic acquisitions and because of the introduction of other perfusion agents that do not require such complicated techniques. However, dynamic ^{133}Xe tomograms give quantitative brain blood flow data that allow each patient's image to be compared more easily to normal values. The newer radiopharmaceuticals imaged with static tomography can give only relative blood flow information. In addition, the ^{133}Xe brain blood flow study is one of the only examples of dynamic acquisition and time-activity curves applied to tomography.

CONCLUSION

In some respects, the amount of processing performed on nuclear medicine images is surprisingly low. Many planar studies are still read from a screen or film after only minimal filtering and zooming. Tomographic images always require computer reconstruction, but again, many are read with little further standardized processing.

In contrast, a few processing techniques are extremely complicated. Some use highly complex mathematics that are not fully understood by most users. Quantitative values appear as if by magic, and sometimes they are accepted as truth without question.

From these two extremes, we attempt to find some middle ground. Day-to-day procedures that affect image quality and quantified measurements should be fully understood. These include any process that requires the user to make decisions, designate image areas or otherwise interact with the operation. Accurate clinical readings and diagnosis from nuclear medicine images depend on suitable processing. In addition, new applications and protocols are more easily developed when the computation methods are well-understood. Part of nuclear medicine's future depends on appropriate processing and accurate quantification of images in order to standardize and objectify image interpretation.

REFERENCES

1. Gonzalez RC, Wintz PA. *Digital image processing*, second edition. London: Addison-Wesley; 1977:1-503.
2. Jain A. *Fundamentals of digital image processing*. Englewood Cliffs NJ: Prentice-Hall; 1989:1-569.

3. Gottschalk A, ed. *Diagnostic nuclear medicine*, second edition. Baltimore, MD: Williams and Wilkins; 1988:1-1154.
4. Erickson JJ, Rollo FD. *Digital nuclear medicine*. Philadelphia: J.B. Lippincott; 1983:1-240.
5. Strauss HW, Zaret BL, Hurley PJ, et al. A scintigraphic method for measuring left ventricular ejection fraction in man without cardiac catheterization. *Am J Cardiol* 1971;28:575-580.
6. Schad N. Nontraumatic assessment of left ventricular wall motion and regional stroke volume after myocardial infarction. *J Nucl Med* 1977;18:333-341.
7. Sandler H, Dodge HT. The use of single plane angiocardigrams for the calculation of left ventricular volume in man. *Am Heart J* 1968;75:325-334.
8. Links JM, Douglass KH, Wagner HN. Patterns of ventricular emptying by Fourier analysis of gated blood-pool studies. *J Nucl Med* 1980;21:978-982.
9. Jengo JA, Oren V, Conant R, et al. Effects of maximal exercise stress on left ventricular function in patients with coronary artery disease using first pass radionuclide angiocardigraphy: a rapid, noninvasive technique for determining ejection fraction and segmental wall motion. *Circulation* 1979;59:60-65.
10. Berger HJ, Matthay RA, Loke J, et al. Assessment of cardiac performance with quantitative radionuclide angiocardigraphy: right ventricular ejection fraction with reference to findings in chronic obstructive pulmonary disease. *Am J Cardiol* 1978;41:897-905.
11. Maltz DL, Treves S. Quantitative radionuclide angiocardigraphy: determination of Qp:Qs in children. *Circulation* 1973;47:1049-1056.
12. Garcia E, Maddahi J, Berman D, Waxman A. Space-time quantitation of thallium-201 myocardial scintigraphy. *J Nucl Med* 1981;22:309-317.
13. Van Train KF, Berman D, Garcia E, et al. Quantitative analysis of stress thallium-201 myocardial scintigrams: a multicenter trial. *J Nucl Med* 1986;27:17-25.
14. Schlegel JU, Halikiopoulos HL, Prima R. Determination of filtration fraction using the gamma scintillation camera. *J Urol* 1979;122:447-450.
15. Gates GF. Split renal function testing using Tc-99m-DTPA: rapid technique for determining differential glomerular filtration. *Clin Nucl Med* 1983;8:400-407.
16. Galt JR, Hise H, Garcia EV, et al. Filtering in frequency space. *J Nucl Med Technol* 1986;14:152-162.
17. Garcia EV, van Train K, Maddahi J, et al. Quantification of rotational thallium-201 myocardial tomography. *J Nucl Med* 1985;26:17-26.
18. Garcia EV, Cooke CD, Van Train KF, et al. Technical aspects of myocardial SPECT imaging with ^{99m}Tc -sestamibi. *Am J Cardiol* 1990;66(suppl):23E-31E.
19. Buchsbaum MS, Ingvar DH, Kessler R, et al. Cerebral glucography with positron tomography. Use in normal subjects and in patients with schizophrenia. *Arch Gen Psychiatry* 1982;39:251-259.
20. Devous MD Sr, Stokely EM, Chehabi HH, et al. Normal distribution of regional cerebral blood flow measured by dynamic single-photon emission tomography. *J Cereb Blood Flow Metab* 1986;6:95-104.