Fundamentals of Image Processing in Nuclear Medicine

12

C. David Cooke, Tracy L. Faber, and James R. Galt

Contents

C.D. Cooke (✉), T.L. Faber, and J.R. Galt
Emory University School of Medicine, Atlanta, GA, USA

12.1 Learning Objectives

The purpose of this chapter is to introduce the reader to the fundamentals of image processing in Nuclear Medicine. It is not meant as a comprehensive guide, but more as an overview and introduction to those topics important to understanding the various forms of image processing. At the conclusion of this chapter, it is hoped that the reader will be able to briefly describe the following:

1. The history of imaging in Nuclear Medicine.
2. The components and processes necessary for generating an image.
3. Spatial resolution and how to calculate it.
4. Image interpolation.
5. How images are converted to color and displayed.
6. Image filtering and its effect on images, as well as have a working knowledge of the two most common filters in Nuclear Medicine.
7. The Fourier Transform and its use in nuclear medicine.
8. The process of analyzing regions of interest.
9. Different ways of segmenting images to extract useful information.
10. Different ways of generating and displaying three-dimensional images.
11. The principles of image registration.
12. The importance of image normalization.

12.2 Introduction

Most Nuclear Medicine procedures are, by their very nature, image-oriented. As early as 1950, printed images of the Thyroid, using I-131, could be obtained

M.M. Khalil (ed.), *Basic Sciences of Nuclear Medicine*, DOI: 10.1007/978-3-540-85962-8_12,

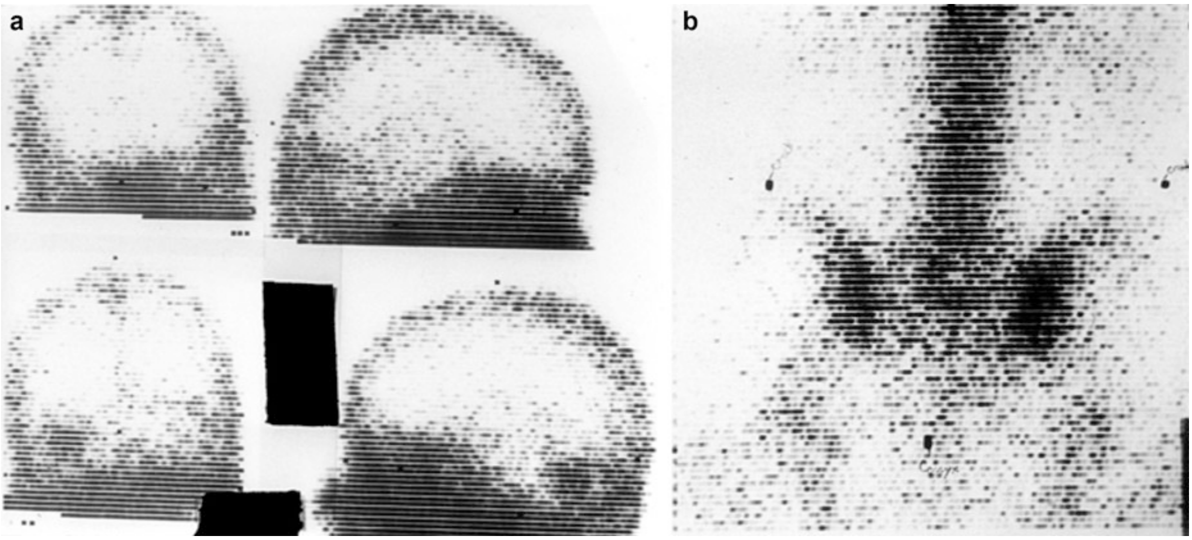

Fig. 12.1 Early rectilinear scans. (**a**) Brain perfusion scan (note defect in bottom left and bottom right images). (**b**) Bone scan (three markers can be seen delineating the left and right iliac crests and the coccyx) (Images courtesy of Naomi Alazraki, MD, Emory University and VA Medical Center, Atlanta, GA)

using the rectilinear scanner developed by Benedict Cassen [1] (see Fig. 12.1). In 1958, the Anger scintillation camera [2], developed by Hal Anger, could be used to image an entire organ at one time; and by 1963, Kuhl and Edwards were using the Anger camera to produce tomographic images [3]. It has been over 50 years since Hal Anger's initial developments, and the traditional gamma camera remains largely unchanged, though now, most have better electronics resulting in increased sensitivity, spatial resolution, energy resolution, and stability. Space will not permit for a detailed history of the gamma camera; however, those who are interested might enjoy the article by Ronald Jaszczak [4].

Today, Nuclear Medicine images are acquired and stored on computers, making the data easily accessible and readily available for further processing and quantification. The goal of this chapter is to familiarize the reader with the basic fundamental techniques of image processing, and specifically, those techniques most useful in Nuclear Medicine.

12.3 Image Presentation

12.3.1 Analog and Digital Images

The scintillation camera was first developed by Hal Anger of the University of California at Berkley in the late 1950s and early 1960s [5], long before computers were in common use and even longer before computer displays that could adequately display a diagnostic image were invented. A major component of these early cameras was circuitry that could convert the output of the photomultiplier array into three voltages representing the x and y locations and the energy (z, the brightness of a scintillation is proportional to the incident photon's energy) of the scintillation photon. If the z voltage matched the voltages calibrated to define the energy window, the x and y locations were sent to an oscilloscope and a momentary flash would appear at a corresponding point on the screen. By placing a focusing lens between the oscilloscope screen and a piece of x-ray film the film was exposed one flash of light (corresponding to one scintillation event) at a time. This type of image is called an analog image, because the x and y locations of the oscilloscope flashes that exposed the film are continuous; they can occur anywhere on the screen. Very nice images could be created with these systems, but they also had some drawbacks. Like all film-based systems, they were subject to under- and overexposure if the technologist cut the exposure time too short or let it go too long. If the clinician wanted to monitor the activity over time, several pieces of film had to be exposed or a system for making several images on one piece of film had to be used.

The progress of computers over the years brought about the development of digital scintillation cameras where a computer is an integral part of the system that processes a scintillation event. Digital cameras provide images in the form of a computer matrix or a digital image. Each cell of this computer matrix is called a picture element or pixel, and represents the counts collected at a corresponding point on the camera face. The values of the pixels are often called pixel counts, because they represent the count of the number of scintillation events detected. The great benefit of a digital image is that it is available for computer processing and display. The exposure time is not a problem because the brightness of the image can be controlled when it is displayed. The pixel counts are also available for computerized analysis instead of being limited to the strictly visual interpretation common with analog images.

The image, or scintigram, is built over time by calculating the x and y location of each scintillation event and, if the energy is determined to be within the desired energy window, incrementing the corresponding pixel. If this is done in real time, the acquisition is said to be in frame mode. If, instead, a list of events and locations is kept it is said to be in list mode. An image is produced from a list mode acquisition after the acquisition is completed. List mode acquisition requires a great deal of more memory and disk space than frame mode, because all scintillation events (regardless of energy) are recorded and more information is kept per event. Regardless of how the image is formed it still needs to be displayed. Scintigram formation is illustrated in Fig. 12.2 and scintigram display is illustrated in Fig. 12.3.

Images are displayed by coding the pixel values to the brightness, or color, of the pixels in the image displays. Usually, scintigrams are displayed in a grayscale format, where black is equal to zero counts and white is equal to the brightest pixel in the image. At times, this is inadequate to capture the dynamic range of an image. It is possible to rescale the image so that a pixel count well below the image maximum but better matching the area of clinical interest is scaled to white. An example of this would be a bone scan that includes the bladder as shown in Fig. 12.4.

Another benefit of digital imaging in nuclear medicine is the ability to acquire images in different formats (Fig. 12.5). Static scintigrams are formed when the scintillation camera acquires a single image from a single location. Dynamic scintigrams are a sequence of static images acquired by defining phases, where each image is acquired for a specific period of time. Gated scintigrams of the heart (often called multigated or MUGA scans, but officially known as equilibrium radionuclide angiography or ERNA) are formed as a sequence of static images that each represent a segment of the heart's EKG R-to-R interval. These images are built up over many heart beats. Whole body images are made by using the scintillation camera to scan fromthe patient from head to foot (or vice versa) and building an image that appears to be one long image made with a very large camera.

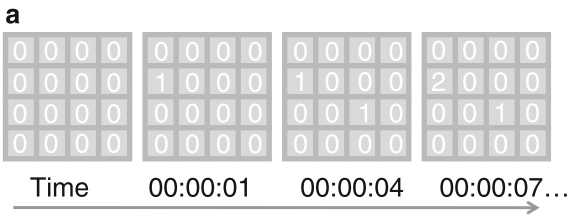

b

Time	X	Y	Energy
00:00:01	1.1	3.4	136
00:00:02	3.8	3.4	100
00:00:04	3.5	2.2	145
00:00:07	1.4	3.1	142
…			

Fig. 12.2 (**a**) In frame mode acquisition, each pixel in the image is incremented as scintillations are counted at a particular point on the camera face. This panel shows how pixels are incremented as photons with the desired energy are detected and localized by the camera. (**b**) In list mode, the same scintillation events are placed into a list and the image is formed from the list after acquisition. Note that if this is a ^{99m}Tc image the second entry may be excluded from the image because of its energy

Tomograms represent slices through the body and are reconstructed from projection data. SPECT images are traditionally made by acquiring static images in an arc around the patient of at least 180° and reconstructing tomograms from these projections. While it is possible to construct projections from a PET acquisition, such a display has little value and is not needed for reconstruction.

12.3.2 Image Math

At times, it may be desirable to perform basic mathematical functions on scintigrams and tomograms. Since these images are made up of pixels and those

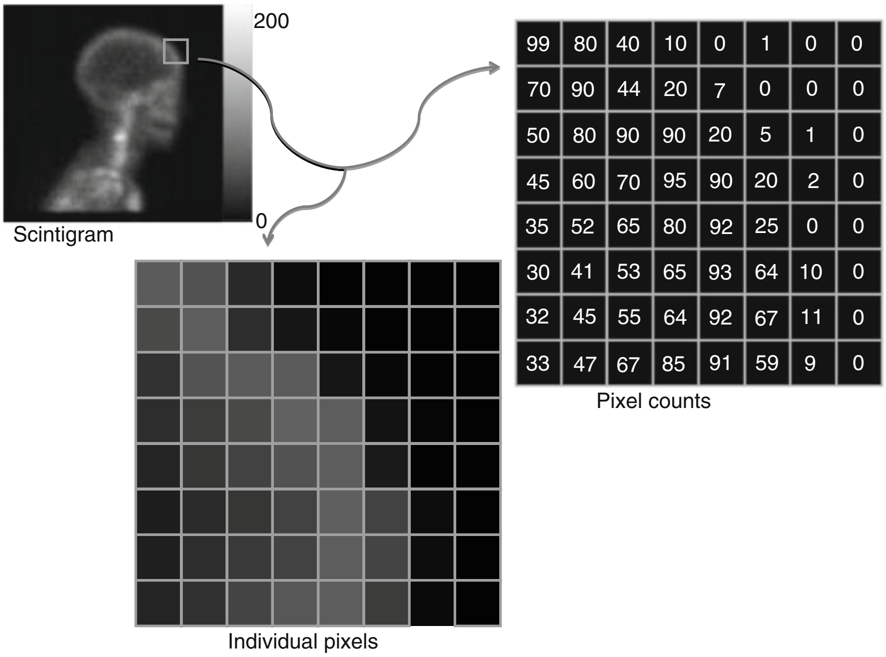

99	80	40	10	0	1	0	0
70	90	44	20	7	0	0	0
50	80	90	90	20	5	1	0
45	60	70	95	90	20	2	0
35	52	65	80	92	25	0	0
30	41	53	65	93	64	10	0
32	45	55	64	92	67	11	0
33	47	67	85	91	59	9	0

Fig. 12.3 A scintigram is a matrix or array of individual pixels, where each pixel has a numeric value related to the number of scintillation events detected in that location. The image is displayed by coding the brightness, or color, of each pixel of the display to the pixel counts

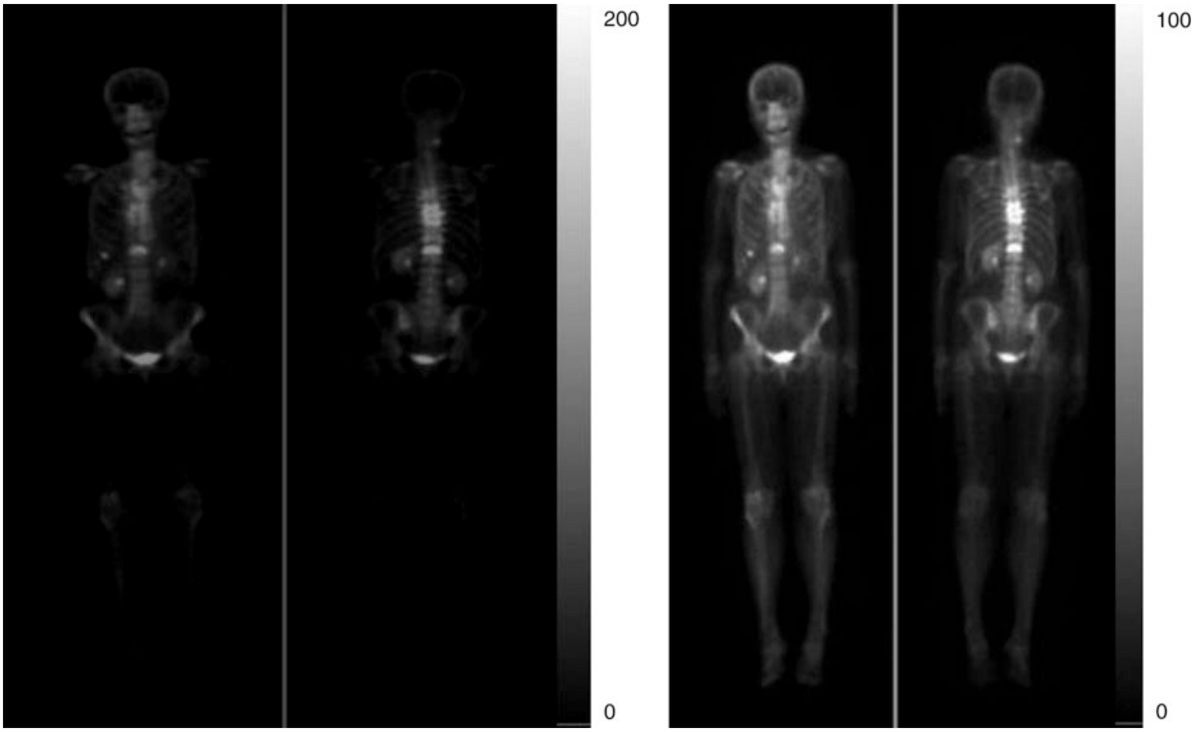

Fig. 12.4 Some scintigrams, such as the bone scan shown here may contain areas of dynamic range well outside of the range of pixel counts in the area of clinical interest. This may be due to high organ uptake in a different organ or, as is the case here, the extraction and excretion of the tracer into the bladder. If the image is displayed scaled from zero to the image maximum, the organ of clinical interest (for example, the bones in the left hand panel) may have too little contrast to be read. By reducing the maximum pixel of the display, it is possible to increase the displayed contrast in the area of clinical interest as shown by the right hand panel. The process of adjusting the displayed brightness and range to suit the scintigram is called windowing

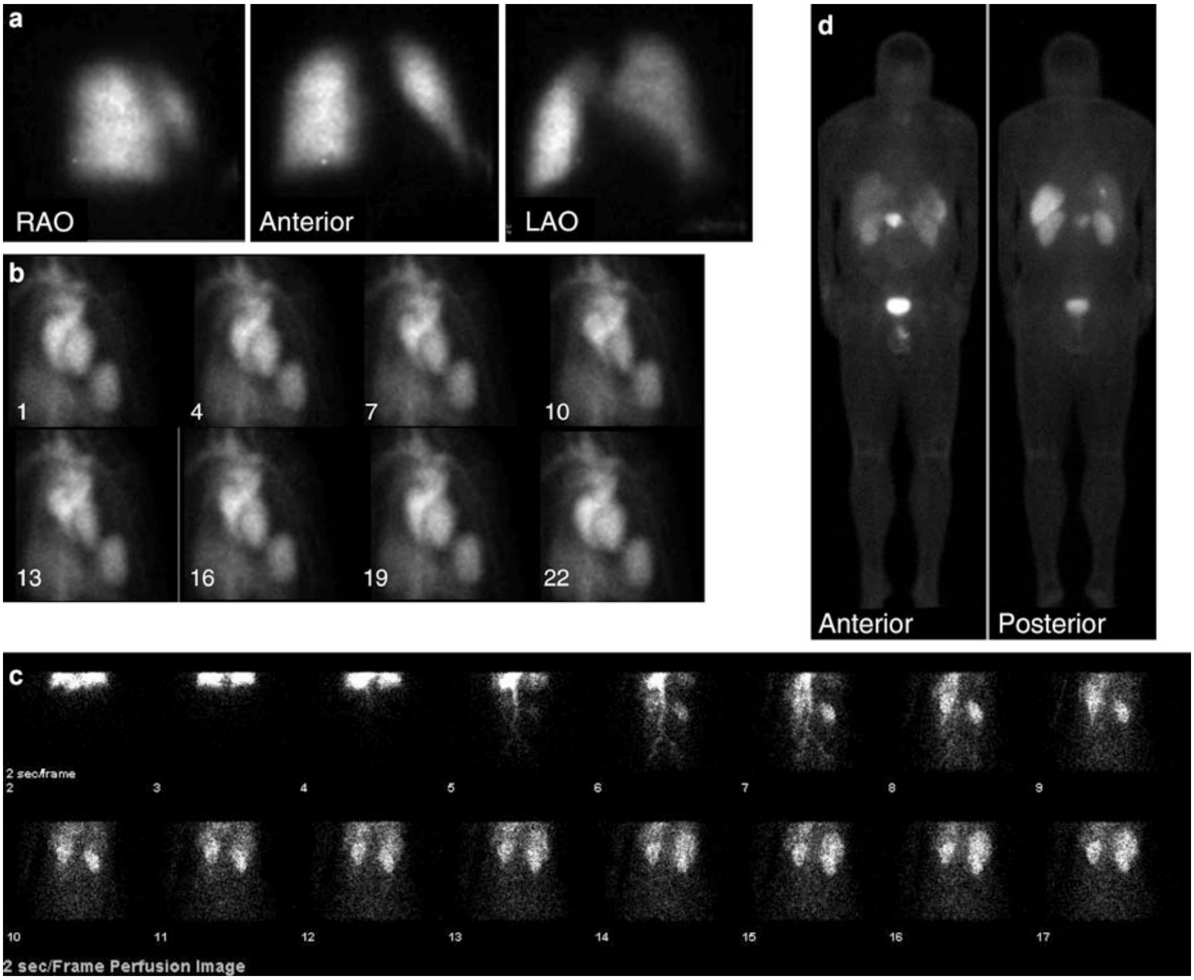

Fig. 12.5 Planar scintigrams may be acquired in different formats. (**a**) Static scintigrams are formed when the scintillation camera acquires a single image from a single location. The technologist manually moves and positions the camera at each desired position. Three views taken in the ^{99m}Tc MAA perfusion lung scan are shown here. (**b**) Gated scintigrams of the heart are acquired over several minutes and show the beating of the heart over an "average" cycle. Each frame of the image represents one segment of the heart's R to R interval. In this example, 8 of the 24 frames that were acquired are shown. The sequence shows the cardiac blood pool starting at diastole, proceeding to systole, and ending back at diastole. (**c**) Dynamic scintigrams allow the clinician to follow a process over time. The example shows the uptake phase of a ^{99m}Tc renogram. Beginning immediately after injection, each of the images in the sequence represents 2 s. The progress of the tracer can be followed as it appears in the blood stream and begins to appear in the kidneys. (**d**) Whole-body scans are acquired by scanning the length of a patient's body with a large field of view scintillation camera. An image is formed that appears to have been made with a single very large scintillation camera. This example shows the uptake of ^{111}In OctreoScan in a patient with an islet cell tumor in the pancreas and metastasis in the liver. (The authors would like to acknowledge Raghuveer K. Halkar, MD, Emory University, Atlanta, GA, for his assistance with this figure.)

pixels define a pixel count at a specific location in space, it is possible to add, subtract, multiply, or divide images by a constant or to perform the operation between two images. Addition of images may be used to reduce the number of images in a dynamic image set by summing frames within a phase, a process known as reframing. This is used in renal imaging, for example, to reduce the number and improve the displayed quality of the frames of the dynamic image, where the original images appear to be very noisy on display but are preferable for extraction of the renogram curves. Subtraction of two images may be used to highlight the difference between two images as in ^{99m}Tc sestamibi – ^{123}I subtraction scintigraphy for hyperparathyroidism where subtraction of the ^{123}I scintigram removes the thyroid gland from the images, making it easier to locate abnormal parathyroid tissue [6] (Fig. 12.6). Prior to subtraction, the two images must be carefully normalized, so that the thyroid glands have the same pixel count values in both scintigrams. This is done by determining a constant that is multiplied by each pixel in the ^{123}I scintigram to give the same pixel counts (on average) as in the ^{99m}Tc sestamibi scintigram. Thus when the images are subtracted, only the areas that represent different activity levels are left.

12.3.3 Matrix Size and Spatial Resolution

Digital images are, as the name implies, discrete representations of the objects that were imaged. This discrete nature can be quantified in terms of Matrix Size and Spatial Resolution. The spatial resolution can actually be broken down into two separate categories: (1) the inherent resolution of the actual imaging device (both with and without a collimator) and (2) the resolution or size of each individual pixel in the image.

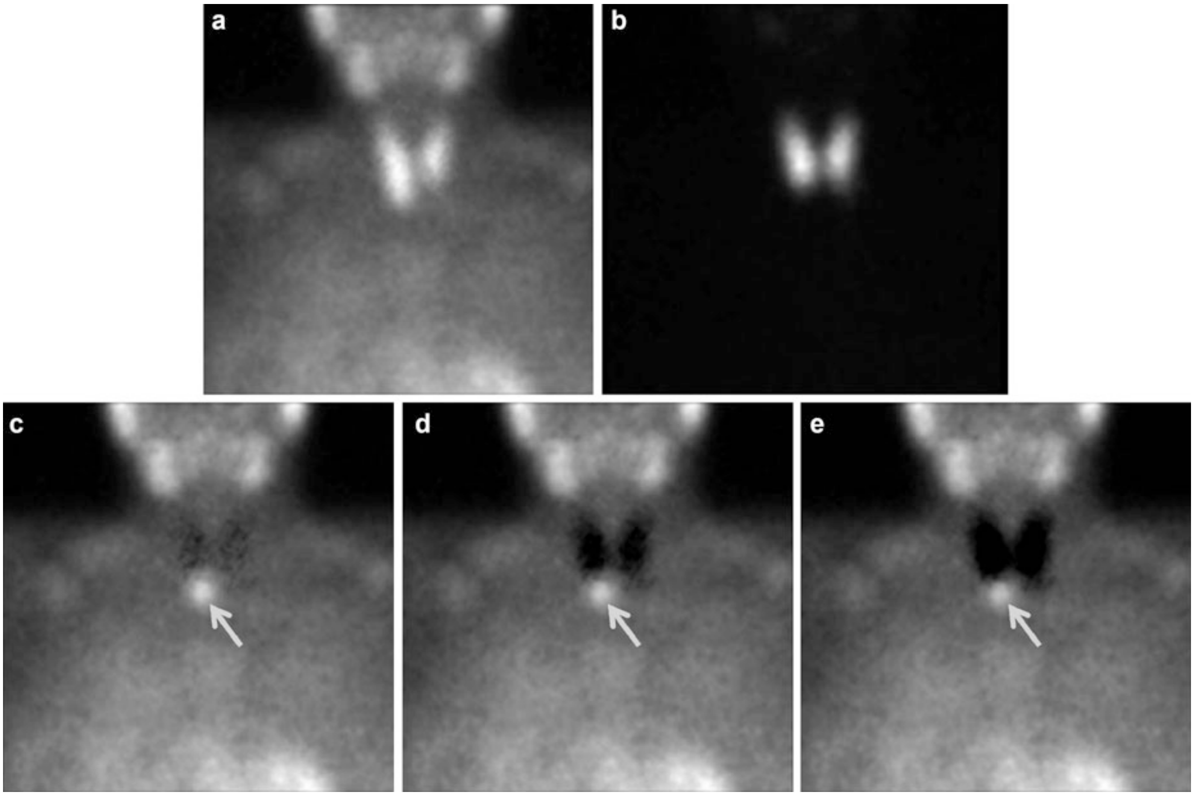

Fig. 12.6 ^{99m}Tc sestamibi – ^{123}I scintigraphy illustrates how one image may be subtracted from the other. While ^{99m}Tc sestamibi (**a**) localizes in both thyroid and parathyroid tissue, ^{123}I (**b**) only localizes in the thyroid tissue. The ^{123}I scintigram is normalized to have the same counts over the thyroid gland as the ^{99m}Tc sestamibi scintigram and subtracted from that image resulting in subtraction scintigram (**c**, **d**, and **e**), making abnormal parathyroid tissue more conspicuous (*arrow*). In the subtraction scintigram labeled C, the ^{123}I scintigram, which has many times fewer counts in the thyroid than the ^{99m}Tc sestamibi scintigram, has been normalized to match the pixel count level of the ^{99m}Tc sestamibi image. In (**d**) and (**e**) it has been normalized to have progressively more counts. (The authors would like to acknowledge Raghuveer K. Halkar, MD, Emory University, Atlanta, GA, for his assistance with this figure.)

12.3.3.1 Matrix Size

Before we delve into calculating Spatial Resolution, we need to briefly cover image matrices. As mentioned above, an image is made up of a collection of discrete picture elements or pixels. These images (collections of pixels) are most often arranged in a square grid or pattern whose size is given as the number of pixels across the face of the image (columns) by the number of pixels down the face of the image (rows); this is also referred to as the image matrix size. For instance, an image that has 64 rows by 64 columns (4,096 total pixels) would be said to have a matrix size of 64 × 64 pixels. Knowing the matrix size of an image is important when calculating the spatial resolution of an imaging system.

12.3.3.2 Spatial Resolution

The spatial resolution of an imaging system can be measured with (extrinsic) and/or without (intrinsic) a collimator, by placing a point source on the face of the camera and acquiring an image. The spatial resolution is then calculated by extracting the counts and location of a row of pixels (called a profile) in this image that go through the point source and measuring the width of the point source profile at half the height of the point source profile (see Fig. 12.7). This number is called the Full Width at Half Maximum (FWHM) and is a direct measure of the spatial resolution of the imaging system. As an example, Fig. 12.7 shows a profile drawn through a point source whose maximum count is 2,908. At half this height (1,454) the width of the

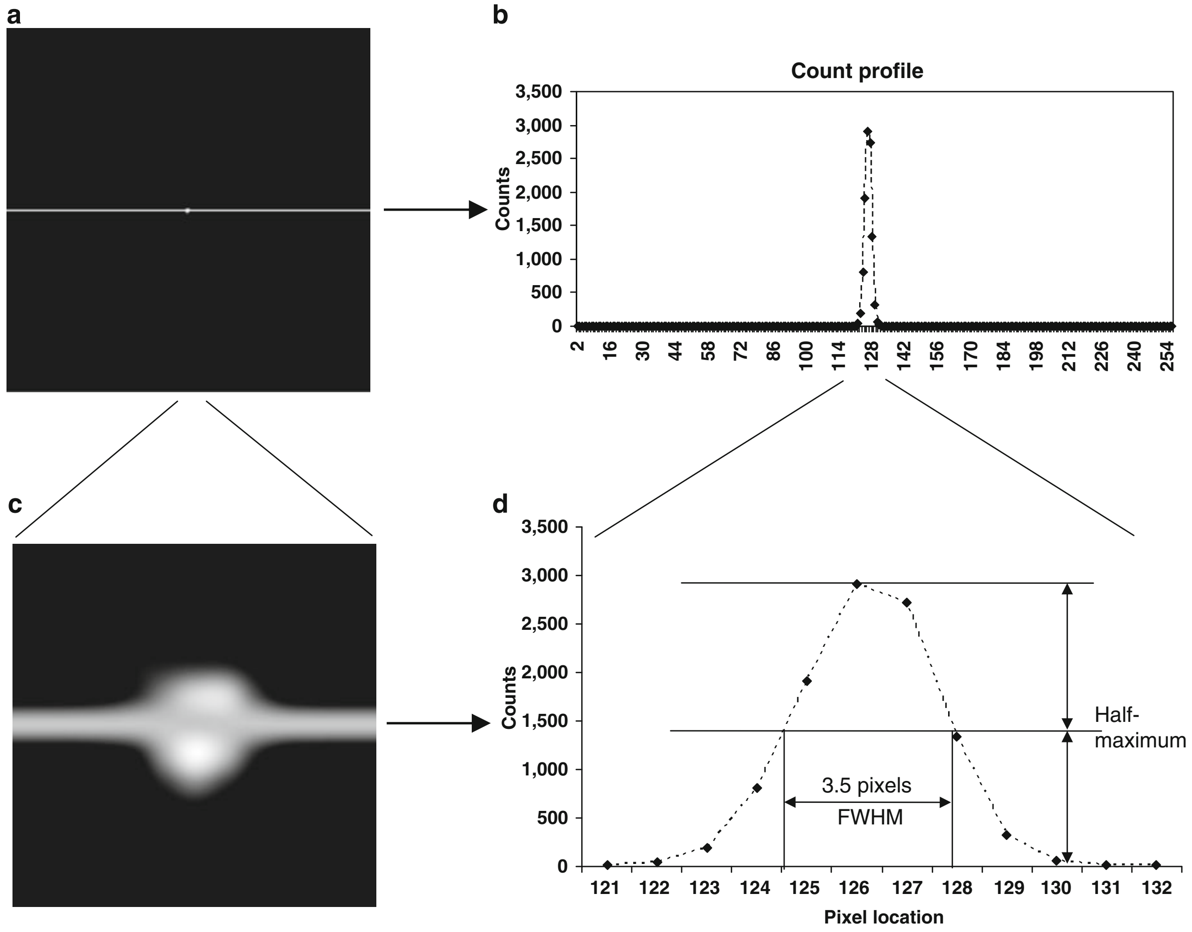

Fig. 12.7 Measuring the FWHM. (**a**) Original point source acquired in a 256 × 256 matrix, (**b**) Plot of counts vs. pixel location for the original image, (**c**) 12 × 12 pixel zoom of original point source image, and (**d**) Plot of counts vs. pixel location for the 12 center pixels showing how the FWHM is calculated

profile is 3.5 pixels, so the spatial resolution of this imaging device is 3.5 pixels. To convert this to mm, we need to know the matrix size of the image and the field-of-view of the camera. For this example, the point source was acquired on a camera with a field-of-view of 532 mm, and the acquired image has a matrix size of 256 × 256 pixels. Putting all of this together, the resultant spatial resolution is:

$$\begin{aligned}\text{Spatial resolution} &= 3.5\ \text{pixels} * (532\ \text{mm}/256\ \text{pixels})\\ &= 7.27\ \text{mm}\end{aligned}$$

The second part of the above equation, (532 mm/256 pixels) is the answer to the second category of the resolution question, the resolution or size of each individual pixel in an image. In the above case, each pixel in this image is 2.08 mm, and is specific for this camera size and this image matrix size. If we had acquired the image in a 128 × 128 pixel matrix, instead of a 256 × 256 pixel matrix then the pixel size would have been 4.16 mm (532 mm/128 pixels). Likewise, if we had used a different camera, with a different field-of-view, this would have also affected the pixel size.

Before we leave this topic, let's explore, briefly, why the FWHM is a measure of the resolution of the imaging system. Another way to refer to the resolution of the imaging system would be to indicate its "resolving" power, or its ability to resolve two different point sources. In other words, how close can two point sources be and the system still be able to distinguish the two point sources. If you look at Fig. 12.8, you will see an example of two point sources spaced generously apart (10 pixels, panel A). As the point sources are moved closer together, the counts begin to overlap (5 pixels, panel B). At the point where half of each profile overlaps each other, we can no longer resolve the two different point sources; for all practical purposes, it looks like one big point source (3 pixels, panel C). You will notice that this happens when the two point sources are spaced exactly one FWHM apart (FWHM = 3.5 pixels for this particular camera), which is why, this number is often used as the resolution of the imaging

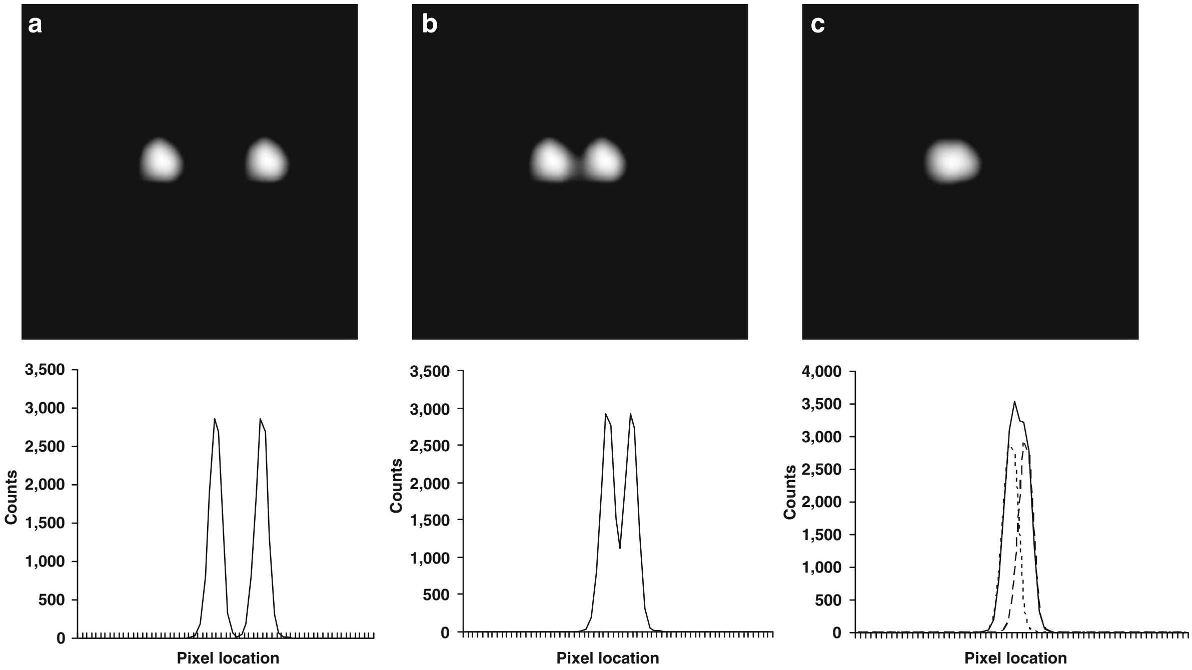

Fig. 12.8 Explanation of FWHM as resolution or resolving power. (**a**) Two-point sources spaced 10 pixels apart. (**b**) Same two-point sources spaced 5 pixels apart. (**c**) Same two-point sources spaced 3 pixels apart; the original profiles of the two separate point sources are seen as dashed lines

system (the point at which two point sources can no longer be resolved) [7, 8].

12.4 Image Interpolation

Another aspect of images and matrix size is changing the size of the image; for instance, you might want to make an image smaller or larger for display purposes. This is usually accomplished by interpolating the original image pixels and making them smaller or larger; in most instances, we are generally interested in making the images larger. As an example, if you were to display a 64 × 64 pixel image on a typical 19″ computer screen whose resolution is 1,280 × 1,024, the image would only take up 1/320th of the screen and would measure approximately inch square. To make the image easier to see, we could zoom or interpolate the image to a larger matrix size, say 256 × 256, which would result in an image size 1/20th of the screen or approximately 3 square in. There are many ways to interpolate an image; we will take a look at three of the more popular methods: Nearest Neighbor (or Pixel Replication), Bilinear Interpolation, and Bicubic Interpolation. An important point to remember is that regardless of the interpolation method used, you can never add more detail or information to an image than what was there to begin with.

12.4.1 Nearest Neighbor

Nearest neighbor interpolation is the process of copying pixels to make an image larger, and is only useful when your final image size is an integer multiple of the original image size. In our previous example, we could have used the nearest neighbor method to interpolate our 64 × 64 image into a 256 × 256 image by replicating each original image pixel 4 times in the new image. Let's illustrate this with a much simpler example. Suppose we have a 2 × 2 image that we want to interpolate to a 4 × 4 image (effectively doubling its size). Figure 12.9a shows our original 2 × 2 image with 4 pixels whose values range from 3–13, and Fig. 12.9b shows our resultant 4 × 4 image with the nearest neighbor interpolation. The advantage of this technique is that it can be done very quickly on any computer and all of the original image data is preserved; however, the results are not always visually pleasing (typically, the resultant images will look blocky). The main disadvantage of this technique is that you can only interpolate images by integer multiples of the original image size. Figure 12.9d shows a central

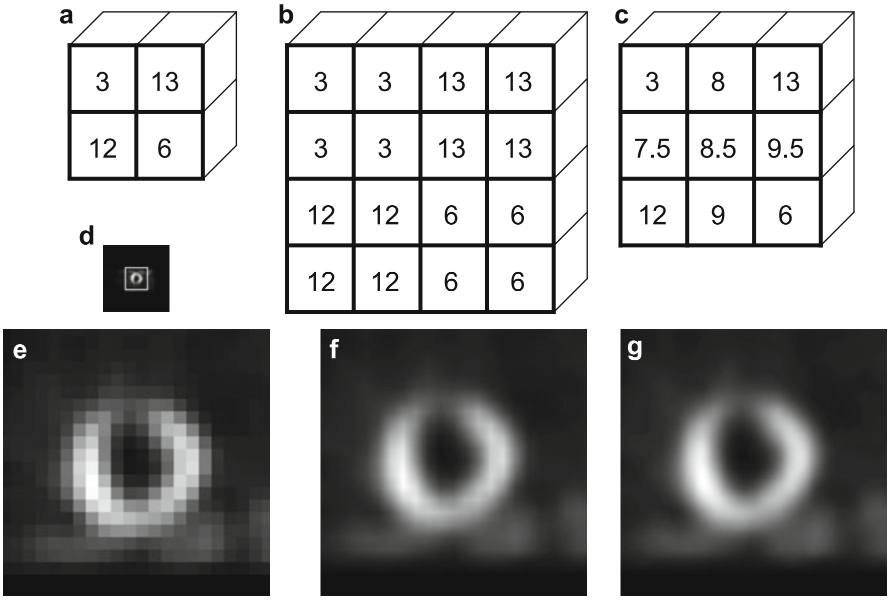

3	13
12	6

3	3	13	13
3	3	13	13
12	12	6	6
12	12	6	6

3	8	13
7.5	8.5	9.5
12	9	6

Fig. 12.9 Example of image interpolation. (**a**) Original 2 × 2 image with 4 pixels, (**b**) Nearest neighbor interpolation to 4 × 4 image, (**c**) Bilinear interpolation to 3 × 3 image, (**d**) original 64 × 64 central short-axis slice from a SPECT myocardial perfusion scan, (**e**) nearest neighbor interpolation of central 21 × 21 pixels of original image (shown as white box in (**d**)) to 256 × 256, (**f**) bilinear interpolation of same area as (**e**), and (**g**) bicubic interpolation of same area as (**e**)

64 × 64 pixel short-axis slice from a myocardial perfusion scan, with a white box around the central 21 × 21 pixels. Figure 12.9e shows a nearest neighbor interpolation of this central 21 × 21 pixel region zoomed to a 252 × 252 pixel region (12 × the original size); notice how blocky the resultant image is.

12.4.2 *Bilinear Interpolation*

Bilinear interpolation fills in the missing pixels by linearly interpolating between the four neighboring pixels (the closest 2 × 2 neighborhood). As an example, we can use linear interpolation to zoom the original 2 × 2 pixel image shown in Fig. 12.9a to a 3 × 3 pixel image. Using bilinear interpolation (linear interpolation applied separately in the x and then the y directions), the resultant image is shown in Fig. 12.9c. Notice how each edge pixel value is halfway between the pixel values on either side (both in the x and y directions – this is a special case for the pixels on the edge of the image), and the very center pixel (8.5) is the average of the four corner or neighbor pixels ($(3 + 13 + 12 + 6)/4 = 8.5$). The main advantage of this method over the nearest neighbor method is that the resultant image does not need to be an integer multiple of the original image and it produces an image that is more visually pleasing (not as blocky); however, it does take more computing power. Figure 12.9f shows a bilinear interpolation of the same central 21 × 21 pixel region zoomed to a 252 × 252 pixel region. Note that this image still contains some slight blockiness.

12.4.3 *Bicubic Interpolation*

Bicubic interpolation has shown improvement over bilinear interpolation by considering the nearest 4 × 4 neighborhood of pixels (instead of the nearest 2 × 2 neighborhood) and using a weighted average of these 16 pixels to calculate the value of the new pixel. In this technique, closer pixels are weighted higher than further pixels and the result is a sharper image than can be achieved with bilinear interpolation. The disadvantage is that it takes even more computing power to calculate. Figure 12.9g shows a bicubic interpolation of the same central 21 × 21 pixel region zoomed to a 252 × 252 pixel region. Note how smooth this interpolation is (there is no apparent blockiness).

12.5 Image Display and Lookup Tables

12.5.1 *Pseudocolor Displays*

In medical imaging, images are usually displayed by assigning a color to each pixel value consisting of red, green, and blue (R, G, and B) intensities. This type of display is called a pseudocolor display. The mapping from pixel value to color is called a lookup table (LUT) or video lookup table (VLT). The number of entries in the lookup table; that is, its length, determines the number of colors that can be displayed, and while historically this has been limited to 256, in theory, it may be much larger. In fact, the number of entries in the lookup table should be equal to the maximum value in the image, so that entries exist for every possible pixel value. However, it is always possible to scale the image prior to display, so that the lookup table size is more a matter of hardware considerations.

The depth of the lookup table is split into three portions, one each for the red, green, and blue values of the color to be assigned. In most cases, each of these portions is one byte, so the total depth of the lookup table is three bytes. Again, this is more a consideration of the hardware and software used for display and may not be user selectable. Note, however, that the size of the R, G, and B entries determines the range of possible values for those colors. That is, if each entry is one byte long, then only the values of 0–255 can be used. Note that this allows 2^{24} possible colors to be placed in the lookup table (256^3); however, as noted above, the number that can be placed there at one time is a function of the length of the lookup table rather than its width. Figure 12.10 shows how the LUT works to define a color that will be used to display each pixel.

Pseudocolor displays are very flexible, because the color mapping associated with the LUT can be interactively changed. A completely new LUT can be loaded almost instantaneously so that the image may be viewed in color, or using a gray scale, without

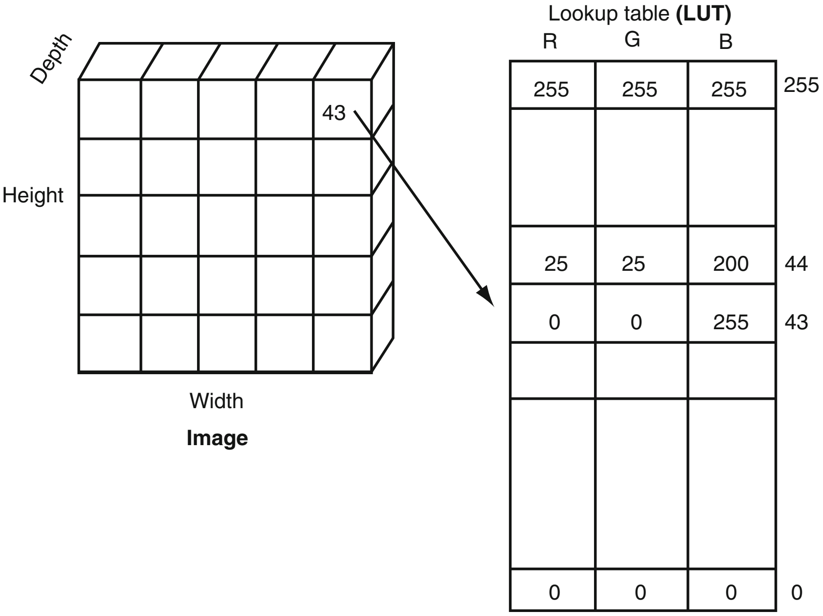

Fig. 12.10 Graphic explanation of pseudocolor displays. Each pixel value in the image is used as an index into the video lookup table (LUT). *Red* (R), *green* (G), and *blue* (B) color values are obtained from the LUT at that index entry. These values are then used to display that pixel

modifying the image data itself. The most common operation for changing a LUT is that of modifying the lower bound or upper bound. The lower and upper bounds are the pixel values at which the LUT changes from its first to its second color or from its next-to-maximum to maximum color. Generally, if the LUT has 256 entries, then the default lower bound is 0, and the upper bound is 255. (The lower bound is also called the level, and the difference between the upper and lower bounds may be called the window.) However, if the lower bound is changed to 10, then the LUT will be compressed, so that the value that was originally at 0 will be placed at all values between 0 and 10, the value at 255 will remain the same, and the other colors that originally ranged from 0 to 255 will be compressed, using interpolation, between those values. Raising the lower bound acts as a type of background subtraction, because low values are now displayed using the same color as “0,” so that these low intensity regions seem to disappear. Lowering the upper bound tends to enhance lower values, because smaller pixel values are now displayed using the portions of the color map that used to be assigned to the higher intensity pixel values. Figure 12.11 shows how modifying the color tables changes the display of a bone scan.

Some particular lookup tables are important to understand. The gray scale contains the same R, G, and B values for each entry, so the “colors” are all gray. They range from black (R,G,B = 0) at the first entry, to white (R,G,B = 255) at the top entry, so that background colors are dim and high pixel values are near white. The reverse gray scale ranges from 255 at the first entry to 0 at the top entry, so that background colors are white, while high pixel values are black. Discrete lookup tables change to different colors in the space of a single entry, and are often scaled so that these color changes occur every 10%. In this way, a quantitative assessment of the pixel values can be obtained – a green value may be associated with pixel values equal to 50–60% of the maximum in the image, for example. However, discrete lookup tables, and indeed, any color table that changes hue in the space of a few pixel values can result in what are called false contours. Two pixel values that are very similar in value may end up being assigned very different colors. This leads the eye to believe that there are hard boundaries between the two pixel values when in fact, there is very little difference between them.

12.5.2 True Color Displays

A second type of display requires that each pixel value in the image contain values for red, green, and blue colors. This is typically not the case for any medical

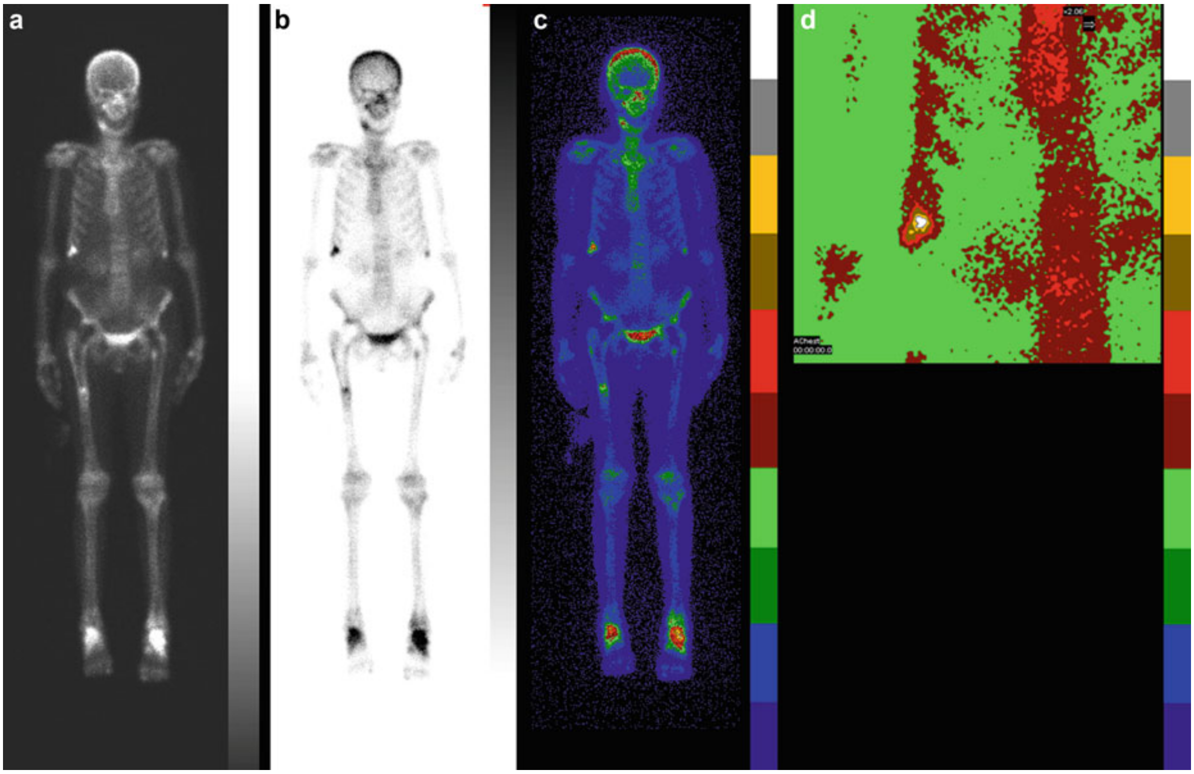

Fig. 12.11 Bone scan displayed using different lookup tables (LUT). (**a**) Grayscale mapping, with bright shades of gray (*white*) indicating large counts. (**b**) Reverse gray scale, where dark shades of gray indicate large counts. This LUT has been modified so that its lower bound is greater than 0, giving the image the appearance of having fewer counts in the background. (**c**) Ten-step color scale, so that at every 10% of maximum counts, the color changes dramatically. When a portion of the image is zoomed (**d**), the rapid changes from about 50–100% of maximum counts in the hot region of the right rib cage can be appreciated

image; however, such displays can be generated and are frequently used for 3D graphics. In these images, the data has three "planes," each associated with the red, green, and blue values used for display. Again, these planes are usually one byte in width, so that the R, G, and B values can range from 0 to 255. No lookup table needs to be used to translate the image into a display; rather, the three planes already encode the colors. Note that in true color displays with byte-sized planes, 2^{24} colors can be displayed at one time all in the same image. This color resolution is necessary in three-dimensional graphics, since many more colors are necessary to give the display realism. However, given that nuclear medicine images really have only one value per pixel, true color displays are not generally needed for displaying slices.

12.6 Image Filtering

Image filtering can be thought of as a mathematical process applied to an image for various reasons, such as removing noise or enhancing features. This process can occur before, during, or after the image reconstruction process, and in most Nuclear Medicine procedures, it is employed to remove noise. It is a very necessary process for SPECT reconstruction, as shown in Fig. 12.12.

Because filtering is a mathematical process, it is much easier to convert images into the math "world" (better known as frequency space) for filtering, rather than converting the filters into the image world. This process of converting images into the math world, or frequency space, is accomplished using the Fourier

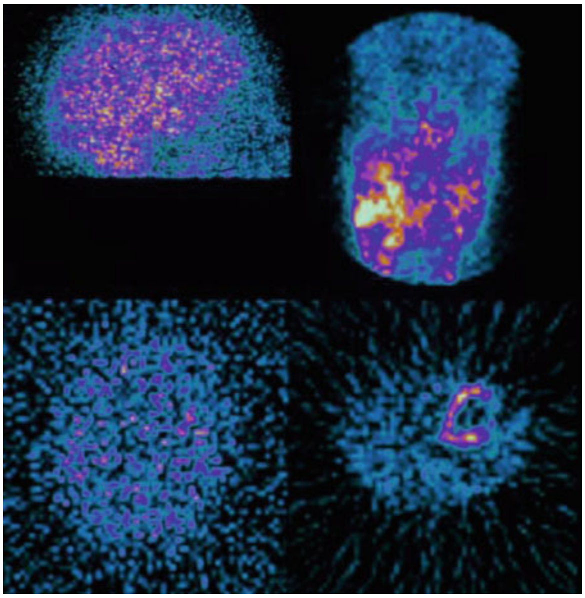

Fig. 12.12 *Top left panel* shows one frame from a Brain Perfusion SPECT acquisition, bottom left shows the same SPECT acquisition reconstructed with no filtering applied. *Top right* shows one frame from a Myocardial Perfusion SPECT acquisition, *bottom right* shows the same SPECT acquisition reconstructed with no filtering applied

Transform, and is covered in the next section. However, it would be good to take a step back and examine how image objects and properties relate to this new math world called frequency space.

12.6.1 Frequency Space

From Collins Essential English Dictionary [9], frequency is defined as the number of times that an event occurs within a given period. For instance, a clock pendulum that swings back and forth every second would have a frequency of one cycle (one back and forth swing) per second, or 1 Hz. Similarly (though overly simplified), since a digital image is made up of pixels, there are a finite number of these elements that can be turned on and off. Let's take the example of an image that has 64 pixels across the image. If we turn half of the pixels (32 pixels) on and then the remaining half of the pixels off, this would translate into a frequency of 1 cycle (pixels on, then off) per 64 pixels or 0.015625 cycles/pixel (Fig. 12.13a). If this process were halved such that 16 pixels were on, then off, then on, then off, this would translate into a frequency of 2 cycles/64 pixels or 1 cycle/32 pixels or 0.03125 cycles/pixel (Fig. 12.13b). If we continue this process of halving the number of pixels that are on or off (Fig. 12.13c–e), we eventually get to the point where every other pixel is on and off, which translates into a frequency of 1 cycle/2 pixels or 0.5 cycles/pixel (Fig. 12.13f). This special case represents the maximum frequency (on/off cycles) that can be represented in a digital image, and is called the Nyquist Frequency (named after the Swedish-American engineer Harry Nyquist).

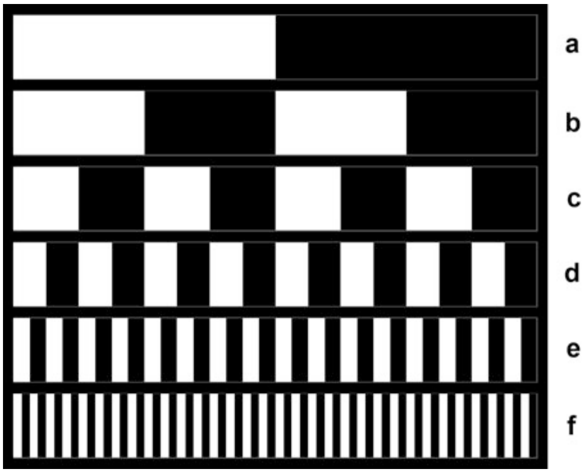

Fig. 12.13 Panels **a–f** show increasing numbers of cycles/pixel for a 64 pixel image (image was zoomed for display purposes). (**a**) 0.015625 cycles/pixel (1 cycle/64 pixels), (**b**) 0.03125 cycles/pixel (2 cycles/64 pixels), (**c**) 0.0625 cycles/pixel (4 cycles/64 pixels), (**d**) 0.125 cycles/pixel (8 cycles/64 pixels), (**e**) 0.25 cycles/pixel (16 cycles/64 pixels), and (**f**) 0.5 cycles/pixel (32 cycles/64 pixels)

Up to this point, we have been using cycles/pixel as the unit of measure for frequency, but other units may also be used. As just noted, the maximum frequency that can be displayed in a digital image is 0.5 cycles/pixel, which can also be represented as a percentage or fraction of the Nyquist frequency. Thus 0.5 cycles/pixel $= 1 \times$ Nyquist frequency $=$ 100% of the Nyquist frequency (note that when using units of percent or fraction of the Nyquist frequency, the maximum frequency is now greater than 0.5). In addition, if the physical size of the detector is taken into consideration, then the units of frequency can be converted into cycles/cm using the following formulas:

$$\text{freq(cycles/cm)} = \text{freq(cycles/pixel)} \\ *(\#\ \text{pixels})/(\text{size of detector in cm})$$

or

$$\text{freq(cycles/cm)} = \text{freq}\ (\%\ \text{of Nyquist}) \\ *0.5\ \text{cycles/pixel} \\ *(\#\ \text{pixels})/(\text{size of detector in cm})$$

For instance, in our above example, if the field of view was 40 cm, then the Nyquist frequency would be 0.5 cycles/pixel × 64 pixels/40 cm = 0.8 cycles/cm (note that when using units of cycles/cm, the maximum frequency can again be greater than 0.5).

Now that we know what a frequency is, the next obvious question is how do frequencies relate to image objects?

12.6.2 Spatial Domain Versus Frequency Domain

If we look again at Fig. 12.13, we notice that as more pixels are turned on and off (more cycles), the corresponding frequency also increases (from 0.015625 to 0.5 cycles/pixel). Following this line of reasoning, low frequencies correspond to the slowly varying portions of images, while high frequencies correspond to quickly varying portions of images. This means that "edges" (points in an image with an abrupt change in intensity such as going from black to white, or white to black) contain numerous high frequency components. In addition, image noise (spurious bright and dark pixels scattered randomly throughout the image) also contains numerous high frequencies; this makes sense if you stop and think about it. Randomly distributed bright and dark pixels will have very abrupt edges when compared with their neighbors, and as mentioned before, it's these edges that contribute to the high frequency components. However, real edges in images (such as boundaries between myocardium and background) also contribute high-frequency components. So, filtering is always a balance between removing as much noise as possible while preserving as much resolution and detail as possible.

12.6.3 Fourier Transform

12.6.3.1 Types of Filters

Generally speaking, all filters fall into one of three classes: (1) Low-Pass filters (filters that pass low frequencies), (2) High-Pass filters (filters that pass high frequencies), and (3) Band-Pass filters (filters that pass a narrow range of frequencies). From the previous section, we learned that noise in Nuclear Medicine images is almost always a high-frequency component; therefore, the goal of most of the filtering we do in Nuclear Medicine is to reduce this noise, which means we mostly employ low-pass filters. The one exception to this is the ramp back-projection filter, a High-Pass filter used in the Filtered Back-Projection process (this topic is covered in a later chapter in this book). Band-Pass filters are a special class of filters, and since they are not generally used in Nuclear Medicine, we will not discuss them here.

There are many Low-Pass filters that we could discuss, including Metz, Weiner, Hanning (or Hann), and Butterworth. Because of space constraints, we shall focus on two of the more popular filters, Hanning and Butterworth.

Hanning Filters

In years past, the most common filter in Nuclear Medicine was the Hanning filter. It is a relatively simple filter with one parameter defining its characteristics, the cutoff frequency. The mathematical definition of the Hanning Filter is as follows:

$$w(f) = 0.5 + 0.5\cos\left(\frac{\pi f}{f_m}\right) \quad \text{if } |f| < f_m$$

where f = spatial frequencies of the image; f_m = cutoff frequency.

By definition, this filter is defined to be 0 at all frequencies greater than the cutoff frequency, completely removing them from the image. Figure 12.14 shows three different Hanning filters with cutoff frequencies of 0.4 cycles/cm, 0.6 cycles/cm, and 0.822 cycles/cm. Note how the filter rolls quickly toward the cutoff frequency and is 0 for all frequencies above the cutoff frequency. This figure also shows

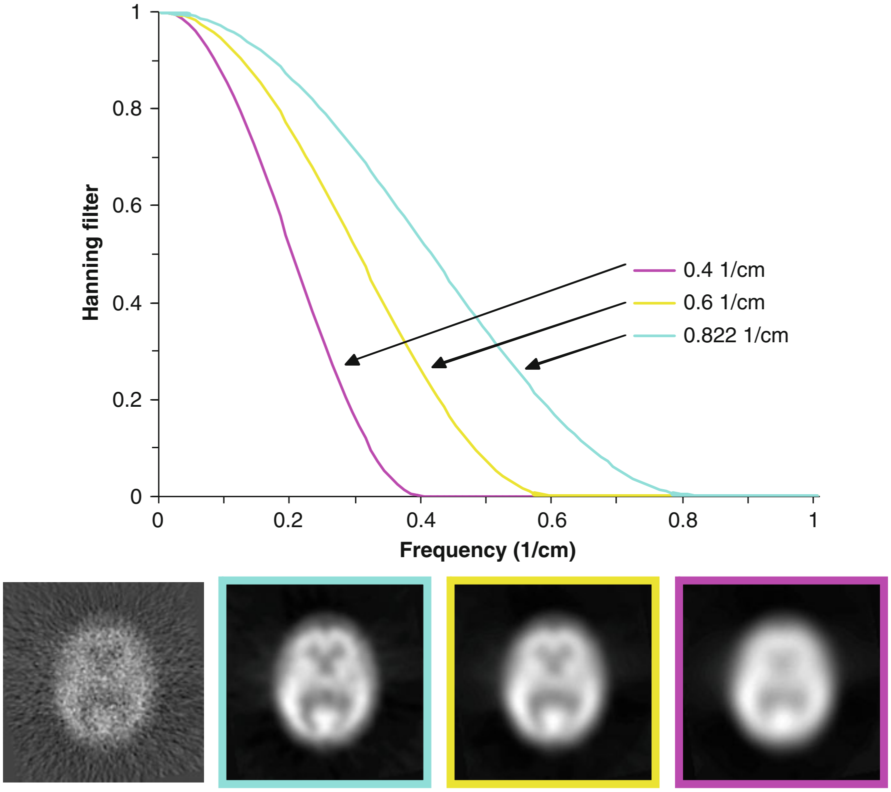

No filter 0.822 cycles/cm 0.6 cycles/cm 0.4 cycles/cm

Fig. 12.14 Hanning filter with three different cutoff frequencies shown, 0.4 cycles/cm, 0.6 cycles/cm, and 0.822 cycles/cm. Note how changing the critical frequency effectively shifts the Hanning filter up and down the Frequency axis of the plot

a typical brain SPECT study filtered with these three filters. Because the Hanning filter rolls-off to 0 so quickly, it is a very good smoothing filter (it very effectively removes high-frequency components, including noise and edge detail); however, it tends not to preserve image resolution.

Butterworth Filters

Sometime in the 1990s, the popularity of the Butterworth filter started gaining momentum. Today, it seems to be the filter of choice in Nuclear Medicine. Perhaps, it is because the Butterworth Filter is so versatile, it can be constructed in such a way that mimics the properties of a Hanning filter, but it can also do a lot more.

There are two parameters that define the Butterworth filter; they are (1) Critical Frequency and (2) Power Factor (sometimes called Order). Notice that the first parameter is called the "critical frequency" instead of the "cutoff frequency" as it was for the Hanning filter. This is because the critical frequency is the point at which the Butterworth filter starts its roll-off toward 0, rather than being defined as 0, as in the Hanning filter. In fact, mathematically, the Butterworth filter never reaches 0, it merely approaches 0. Unfortunately, there is some confusion as to the correct name of the second parameter (what I have called power factor). Mathematically, the power factor $=$ $2 \times$ order, though there are some manufacturers who do not follow this rule. The only real way to determine if the Butterworth filter, as implemented in your particular equipment, is using the power factor or order is to examine the equation that was actually implemented. Here is the equation for the Butterworth filter using power factor:

$$w(f) = \frac{1}{\sqrt{1 + \left(\frac{f}{f_c}\right)^p}}$$

where f = spatial frequencies of the image; f_c = critical frequency; p = power factor.

In the case where a manufacturer has implemented the Butterworth filter using order instead of power factor, then the equation would look like this (note that p has been replaced with $2n$, where n is the order):

$$w(f) = \frac{1}{\sqrt{1 + \left(\frac{f}{f_c}\right)^{2n}}}$$

In practice, the critical frequency of the Butterworth filter behaves much like the cutoff frequency of the Hanning filter, in that, as the critical frequency is changed, the effect is to shift the resultant Butterworth filter up and down the Frequency-axis, as shown in Fig. 12.15. The power factor effectually changes the steepness of the Butterworth filter's roll-off toward 0, as shown in Fig. 12.16. Because of this versatility (the ability to change not only the frequency of the roll-off, but also the steepness of the roll-off), the Butterworth filter can be used to more effectively remove noise while still preserving resolution. Note how the noise in the Brain images at the bottom of Fig. 12.16 is affected by the change in power factor (increasing the power factor effectively reduces the higher frequency components, such as noise), but the detail in the interior structures is preserved.

12.6.3.2 Fourier Transform in Filter Application

In this section, we shall examine the process of actually applying a filter to an image. As mentioned in the previous section, filtering usually takes place in the mathematical world of frequency space, because filter implementation and application is easier. The complete filtering process is shown in Fig. 12.17 and goes something like this: an image is converted to frequency space, the filter is applied, and the resulting image is converted back to image space. The process of converting an image into frequency space is accomplished using the Fourier transform [10–14]. This mathematical operator is nondestructive, completely reversible, and can be computed fairly quickly for images whose dimensions are a power of 2 (using the

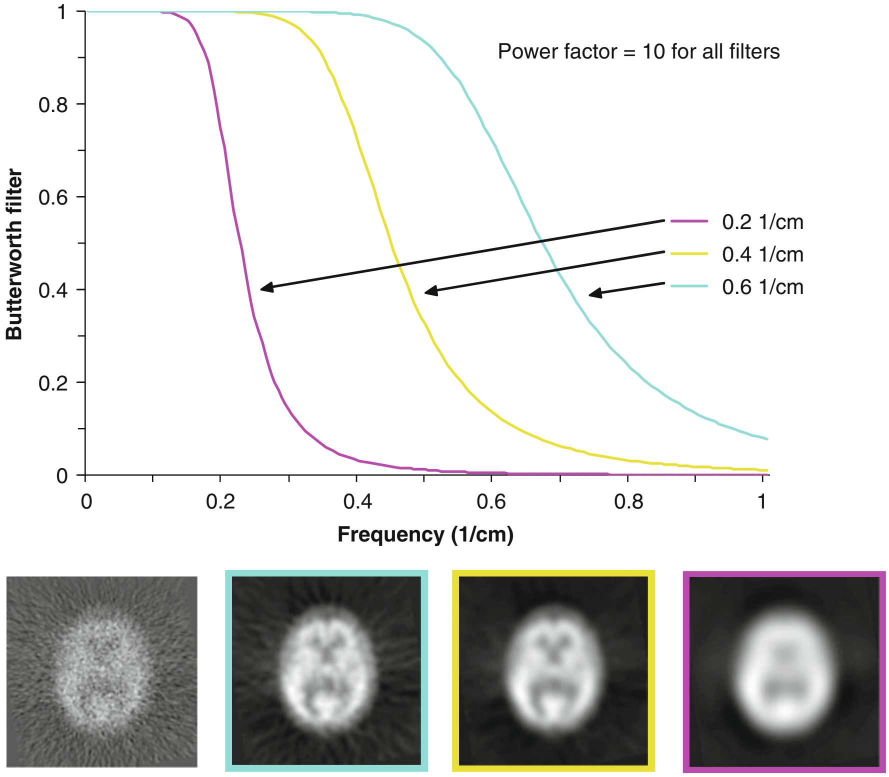

No filter 0.6 cycles/cm 0.4 cycles/cm 0.2 cycles/cm

Fig. 12.15 Butterworth filter with three different critical frequencies and a constant power factor of 10: 0.6 cycles/cm, 0.4 cycles/cm, and 0.2 cycles/cm. Note how changing the critical frequency (like changing the cutoff frequency of the Hanning filter) effectually shifts the filter up and down the Frequency axis of the plot

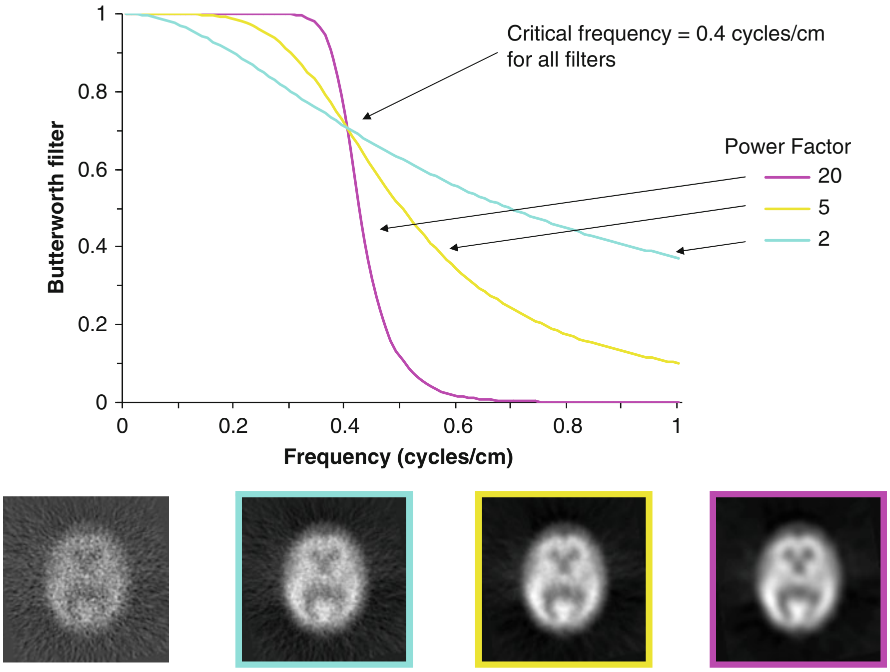

No filter PF = 2 PF = 5 PF = 20

Fig. 12.16 Butterworth filter with three different power factors and a constant critical frequency of 0.4 cycles/cm: PF = 20, PF = 5, and PF = 2. Note how increasing the power factor increases the steepness of the filter's roll-off toward 0

FFT or Fast Fourier Transform). The frequency information that is returned by the Fourier Transform can also be viewed as a series of sine or cosine waves. For instance, the square wave shown in Fig. 12.18, panel A has frequency components that are shown in panel B. These components can also be shown as a series of sine waves, as seen in panel C. If one were to do a point-by-point addition of all of the sine waves, the resultant waveform would be equal to the square wave shown in panel A. Another term for each of these individual sine waves, or frequencies, is harmonics. The lowest frequency is referred to as the first harmonic, the second lowest frequency as the second harmonic, and so on. In reality, the first harmonic is the best "fit" of a sine wave to the original data. As one adds more harmonics, the "fit" becomes more and more like the original curve. There are several different ways to display this frequency information to the user, we will be showing the information in graph style, where each point on the x-axis represents a specific frequency, and the height of each point represents how much of that frequency is in the image (as in Fig. 12.18b. Let's illustrate this with several examples.

Figure 12.19a shows an image consisting of eight increasing frequencies, ranging from.0039 cycles/pixel to 0.5 cycles/pixel. Panel B shows the corresponding graph of these eight frequencies (also referred to as the input frequency spectrum or just input spectrum). Panel C shows a graphical representation of a uniform filter; note that this graph uses the same units as the input frequency spectrum. In frequency space, filtering is simply the multiplication of the input frequency spectrum, point-by-point, with the filter. In other words, every frequency point in the input spectrum is multiplied by its corresponding point in the filter graph. For this example, since the uniform filter has a value of "1" for each frequency, each input frequency is multiplied by 1 and copied to the output frequency spectrum graph (see panel D). We can now convert this filtered output frequency spectrum into an image by running the inverse of the Fourier transform; the resultant image is seen in panel E. Since the output spectrum is identical to the input spectrum, it is not surprising that the filtered image (output image) is identical to the input image.

We can now add what we learned about different filters in the previous section to the filtering process. Figure 12.20, panels A and B show the same input image and spectrum as Fig. 12.19. However, the filter in Fig. 12.20c is now that of a Hanning filter with

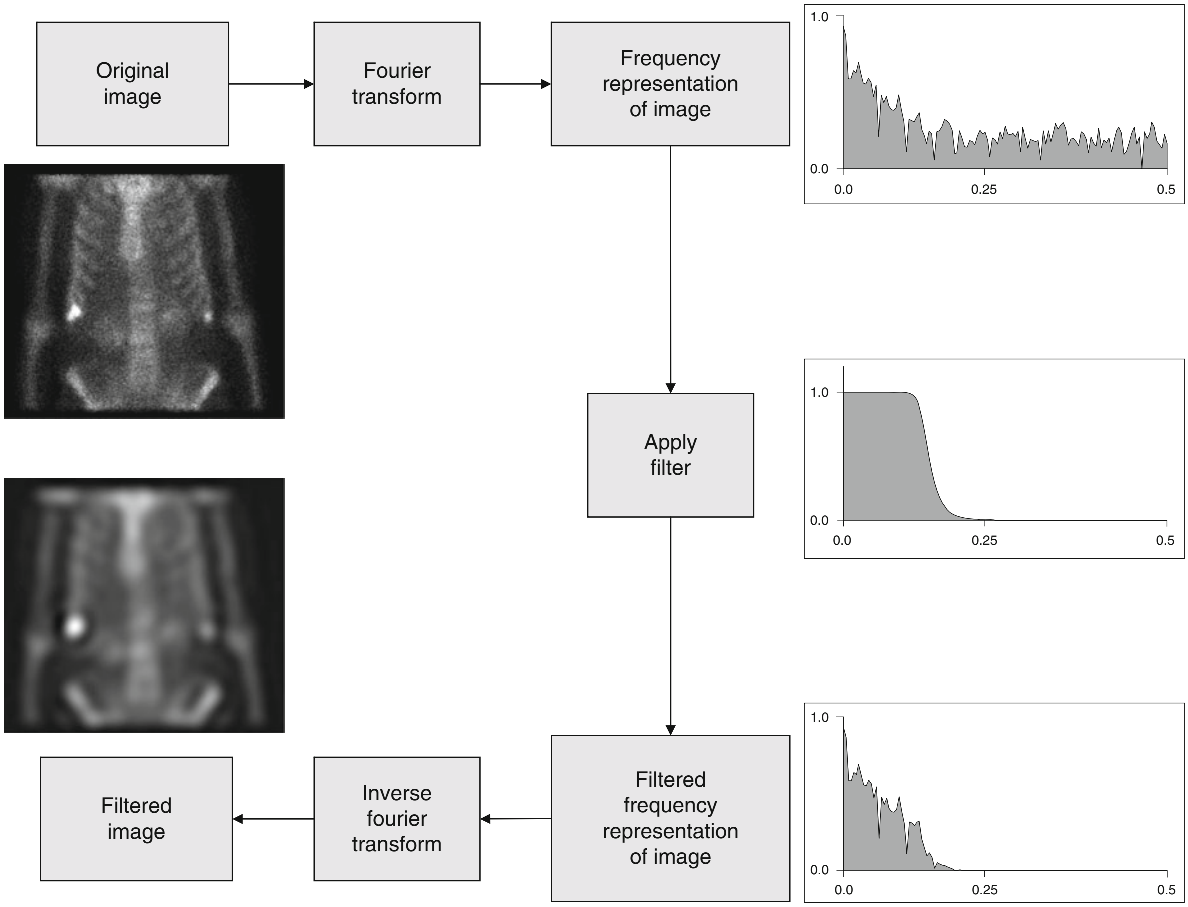

Fig. 12.17 Diagram of the filtering process. The original image is converted to Frequency Space using the Fourier transform, where an image filter is applied. The resultant frequency spectrum is converted back to image space using the inverse Fourier Transform

a cutoff value of 0.5 Nyquist. The output spectrum in Panel D shows the results of multiplying the input spectrum by this filter, and the corresponding image is shown in panel E. Note that because of this filter's sharp roll-off, the frequency at 0.5 Nyquist is completely removed from the image, and the other frequencies have been somewhat reduced. Figure 12.21 shows an additional example using a Butterworth filter with a critical frequency of 0.03125 Nyquist and an order of 5. Note that this filter has completely removed the frequency at 0.5 Nyquist, and has all but completely removed the frequencies at 0.125 and 0.25 Nyquist as seen in the output spectrum and the resultant filtered image.

As a final example of filtering, let's look at a real case involving a Planar Bone Scan (Fig. 12.22). In Panel A you can see the original image. Panel B shows the input spectrum of the Bone scan (note the rather continuous nature of the input spectrum, indicating that frequencies of all values are present). Panel C shows a Butterworth filter with a critical frequency of 0.125 Nyquist and a Power Factor of 20. The resultant spectrum (Panel D) and image (Panel E) show the results of this filter.

12.6.3.3 Fourier Transform in Curve Fitting

We can use the "fitting" property of the Fourier Transform to our advantage when looking at various physiologic phenomena. For instance, when calculating an ejection fraction, the time-volume curve generally looks much like a time-shifted sine wave (see Fig. 12.23). We can use the first harmonic of the

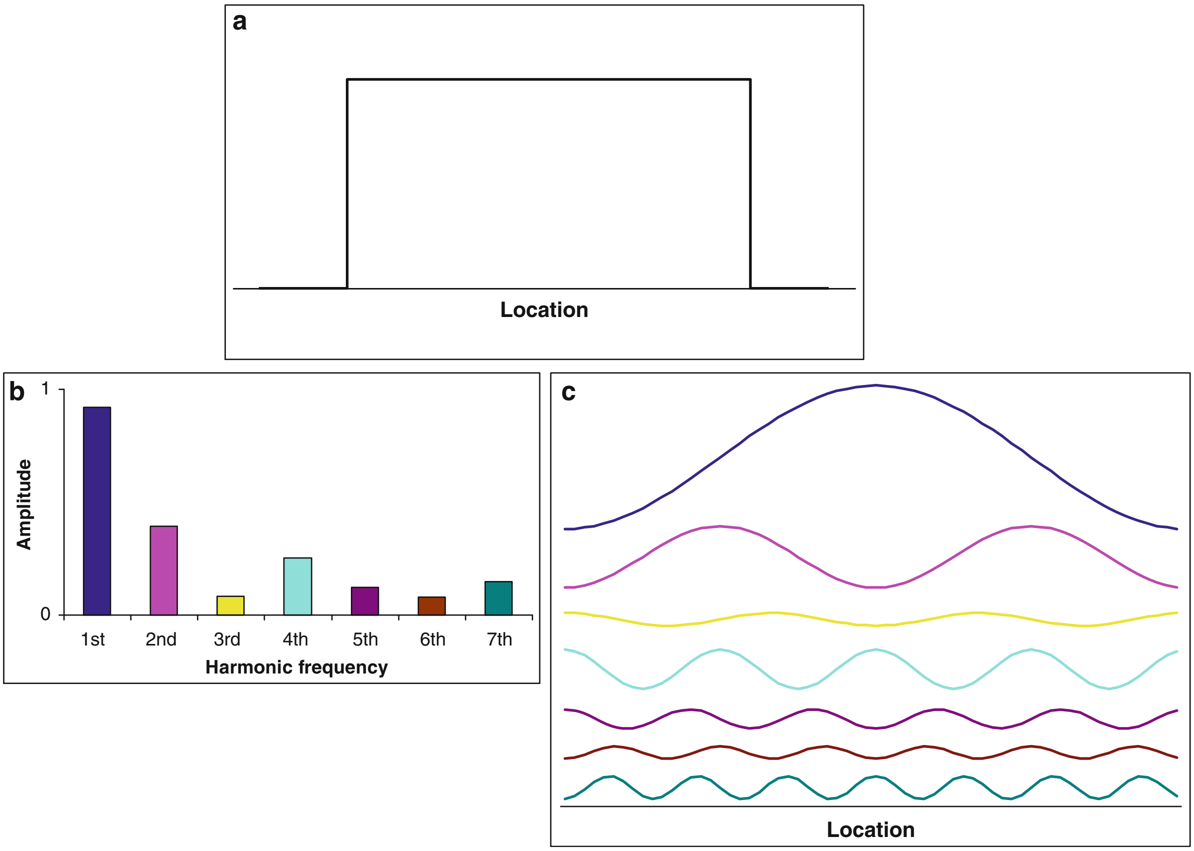

Fig. 12.18 Decomposition of square wave into frequency components using the Fourier transform. (**a**) Original square wave. (**b**) Fourier transform frequency spectrum of the square wave shown in (**a**). (**c**) Another way of showing the Fourier Transform spectrum as sine waves

Fourier Transform of the time–volume curve to actually "fit" this time–volume curve to a sine wave. This is of particular interest when there are few time points or where the data is particularly noisy [15]. As you can see in Fig. 12.23, the first harmonic of the Fourier Transform is much smoother than the curve of the original eight points. In addition, recent research has shown that the temporal resolution of an 8 or 16-frame study is equivalent to a 64-frame study, when the original study has been "fitted" to the first, second, or third harmonics of a Fourier Transform [16].

There are at least two exciting Cardiology applications of fitting time-volume curves using the Fourier Transform: (1) Determination of Wall Thickening and (2) Phase Analysis. Though space will not permit a detailed treatment of these two applications, we shall describe them briefly below.

12.6.3.4 Fourier Transform in the Determination of Wall Thickening

It has been shown that if the size of an object being measured is less than or equal to twice the resolution of the imaging equipment (twice the FWHM), then the counts in the object increase linearly with its thickness [17]. This property of linear count increases can be used to measure the "thickening" of the myocardial wall in myocardial perfusion gated SPECT studies [18]. A typical SPECT system has an intrinsic resolution of about 10 mm, and the average end-diastolic width of the walls of the left ventricle are about 1 cm [19]; this falls into the category of the object size being less than twice the FWHM of the imaging system. If we extract the counts of the left-ventricular myocardium from a myocardial perfusion gated SPECT

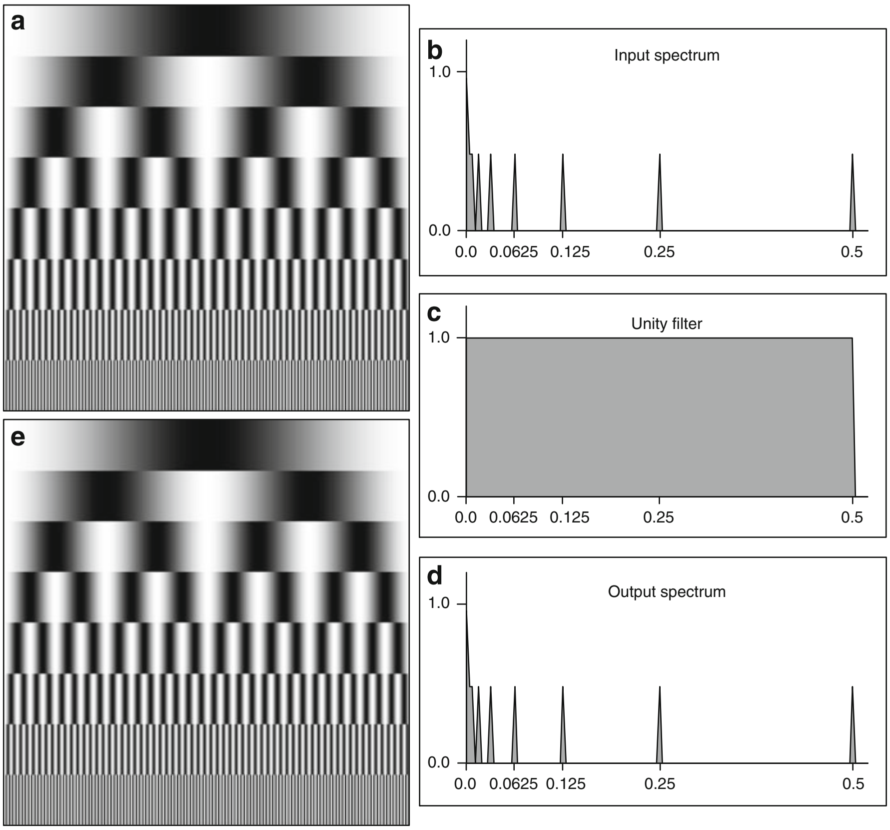

Fig. 12.19 Example of filtering. (**a**) Original image with eight frequencies. (**b**) FFT spectrum of original image. (**c**) Unity filter (one everywhere). (**d**) FFT spectrum after application of filter. (**e**) Filtered image

study, and look at the counts of each point in the myocardium as it goes through the cardiac cycle, we will get a series of time-count curves similar to the curve shown in Fig. 12.23. If we fit each of these curves using the first harmonic of the Fourier Transform (also shown in Fig. 12.23), then we can estimate the percentage of thickening of the myocardium between end-systole and end-diastole. Assuming that the counts are linearly proportional to thickness, then the percentage of increase of counts from end-diastole to end-systole should be linearly proportional to the increase in wall thickness. For example, if the amplitude of the sine wave is 13, and the DC component is 135, then the percentage thickening could be calculated as (2 × amplitude)/(DC – amplitude) and would be 21%. Though this number could be calculated directly from the counts extracted from the myocardium, it is less affected by noise or number of acquired frames/cardiac cycle if the Fourier Transform is used. By calculating this percentage thickening for every point in the myocardium, we can come up with a global assessment of how the myocardium is thickening throughout the cardiac cycle. This assessment can be important, as areas that are not thickening (which implies impaired function) may be an independent indicator of dead or dying myocardial tissue. It can also be used to rule-out false positives. For instance, an area of the myocardium that does not exhibit many counts at stress and at rest (compared to the surrounding tissue) would normally be viewed as dead myocardial tissue. However, if this same area is seen to “thicken,” i.e., there are count changes between end-diastole and end-systole, then the count reductions are probably due to some kind of artifact (Breast or Diaphragm attenuation), and the tissue is more than likely alive and well [20].

12.6.3.5 Fourier Transform in the Calculation of Phase Analysis

Another benefit of using the Fourier Transform is that we can also measure the onset of thickening (or onset of mechanical contraction); this is sometimes referred

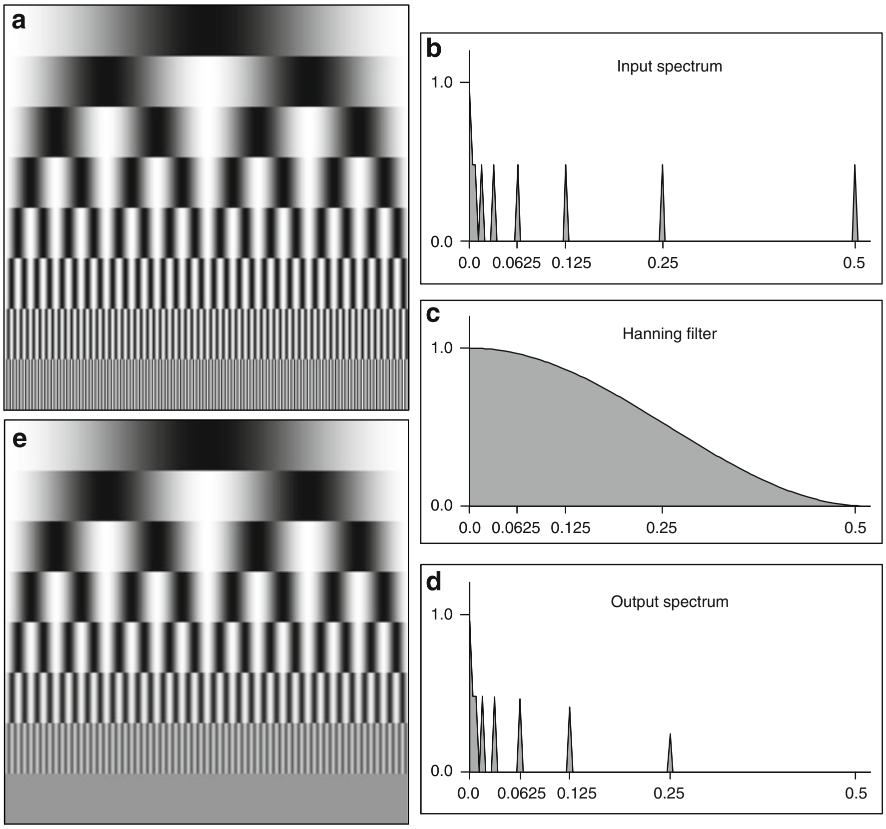

Fig. 12.20 Example of Hanning filter (**a**) Original image with eight frequencies. (**b**) FFT Spectrum of Original Image. (**c**) Hanning filter, cutoff frequency = 0.5 Nyquist. (**d**) FFT Spectrum after application of filter. (**e**) Filtered image

to as Phase Analysis. If we again look at Fig. 12.23, the point at which the thickening curve crosses the DC component of the Fourier Transform is known as the phase shift or just phase of the curve. A sine wave with no phase shift would have started at the DC crossing point, which would be a phase of 0. As the sine wave is shifted to the right, the phase increases indicating a delay in the contraction of the ventricle. If we calculate this phase shift for every point in the myocardium (much like we calculate the percentage thickening for every point in the myocardium) we can use this distribution of phases to diagnose various cardiac conditions. There are several ways to look at this distribution of phases, two of which are most common: (1) as a histogram and (2) as an image. Figure 12.24 shows the output screen of a normal Multiple Gated Acquisition (MUGA) study. Notice that the phase image is all one color and the histogram is very narrow, both indicating that all sections of the ventricle are essentially contracting at the same time. Figure 12.25 shows the output screen of a MUGA for a patient with an apical aneurysm. Notice that the phase image is no longer a single color and the histogram has broadened a little. More recently, phase analysis information extracted from gated SPECT or PET studies [15, 21, 22] has proven useful in evaluating patients being considered for Cardiac Resynchronization Therapy (CRT) [23–26]. Figure 12.26 shows the normalized phase polar maps and histograms from two different studies, a normal study (top) and an abnormal study (bottom). Note how the normal study has a tall, narrow, well-defined peak, indicating that the left ventricle is very synchronous. In contrast, notice how the abnormal study has a very broad diffuse distribution, indicating that the left ventricle is severely dyssynchronous.

12.7 Region of Interest Analysis

Generally, quantitative measures from nuclear medicine images are obtained only from certain important regions; these are called regions of interest (ROIs).

Fig. 12.21 Example of Butterworth filter (**a**) Original image with eight frequencies. (**b**) FFT spectrum of original image. (**c**) Butterworth filter, critical frequency = 0.03125 Nyquist, Power Factor = 5. (**d**) FFT spectrum after application of filter. (**e**) Filtered image

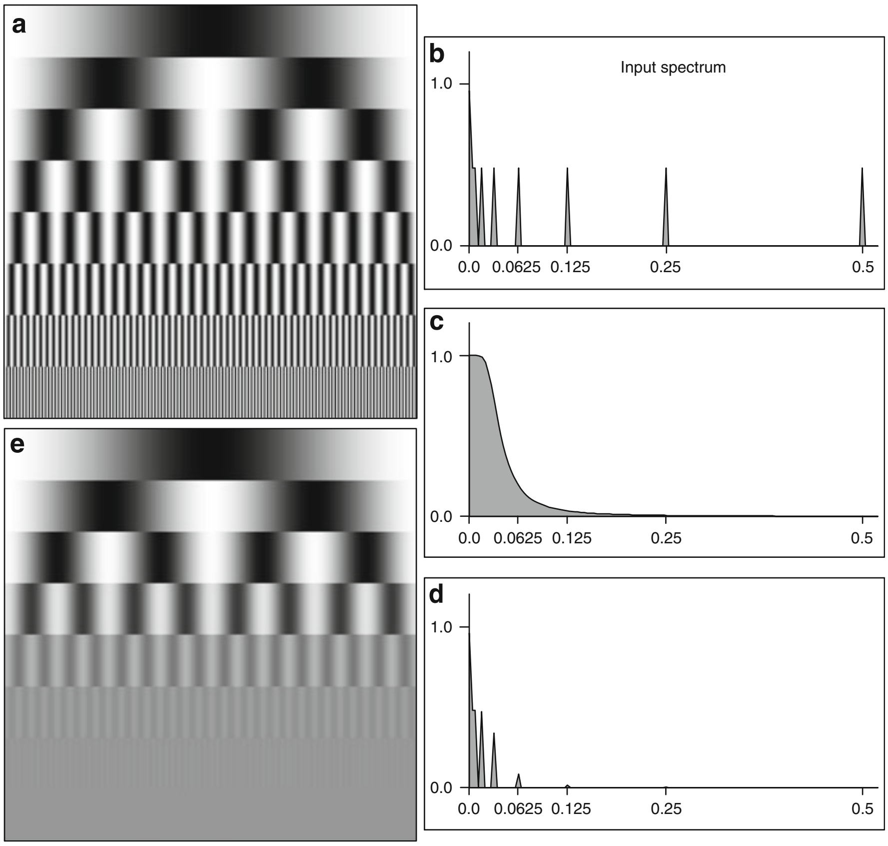

If the region is 3-dimensional (3D), the term "volume of interest" is sometimes used (VOI). A ROI can be, for example, a tumor in an oncologic image, the kidney in a renal image, or the myocardium in a cardiac image. Measurements may be taken from the count values within the region; for example, the mean and maximum count values can be computed. They may also be taken from the shape of the boundary of the ROI; for example, the number of pixels within it, or its size may be determined. Obtaining an appropriate, or accurate, ROI is one of the most basic steps in image quantitation. In some cases, geometric ROIs, such as an oval or circle, may be appropriate. It is always possible to interactively trace the region of interest, and this is still probably the most common approach. An expert user can use his or her knowledge of the expected shape of the organ along with the appearance of the image to obtain a good boundary. These two approaches are shown in Fig. 12.27. However, interactive processes are associated with significant inter and intraobserver variability, they are time consuming, and there are not always experts available to perform the task. For that reason, automatic methods may be helpful, if only to aid the user in performing the interactive tracing.

12.7.1 Background Subtraction

In planar images, counts through the entire body are acquired, from soft tissue both in front of and behind the organ of interest. Accurate quantification of the counts within the organ may be obtained only if these counts can be eliminated from the organ ROI; this is called background subtraction. The most common approach for background subtraction is to find a region in the image that is near to, but not in the organ ROI, with no obvious additional high-count structures.

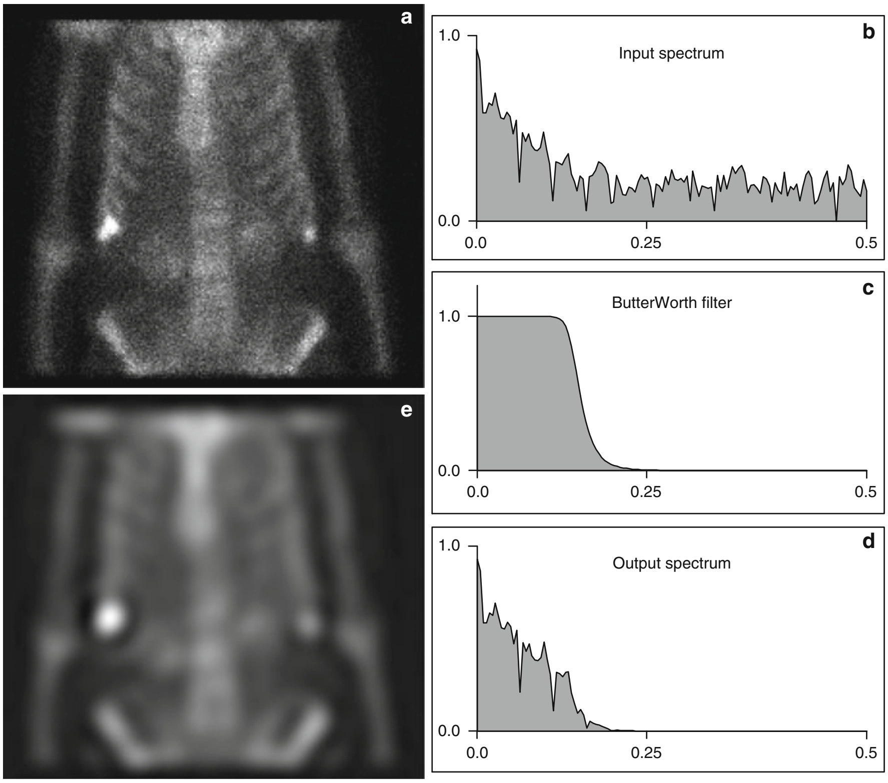

Fig. 12.22 Example of Butterworth filter with planar bone scan. (**a**) Original planar bone scan. (**b**) FFT spectrum of original image. (**c**) Butterworth filter.125 Nyquist, Power Factor = 20. (**d**) FFT spectrum after application of filter. (**e**) Filtered image

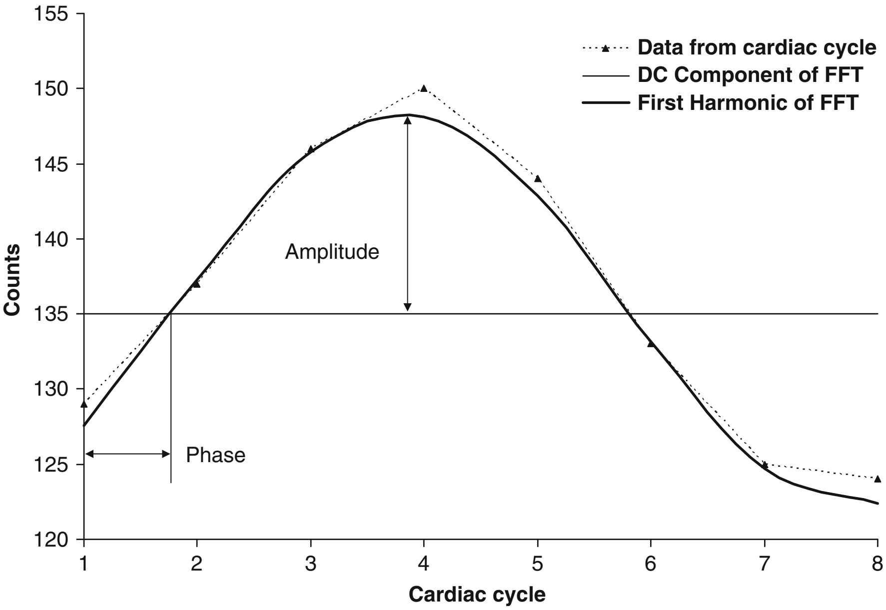

Fig. 12.23 Example of Fourier curve fitting. Eight time points from a midventricular gated short-axis slice are shown in the *dotted lines*. The result of fitting these eight points to the first harmonic of the FFT are shown in the *dark solid line*. The *light solid line* shows the DC component (or average) of the FFT

A background region can be drawn interactively or automatically in this area. The mean counts per pixel are determined, and this value is subtracted from every value in the image. More complicated background regions can be drawn as well; for example, an oval shell, or doughnut, can be used to surround the

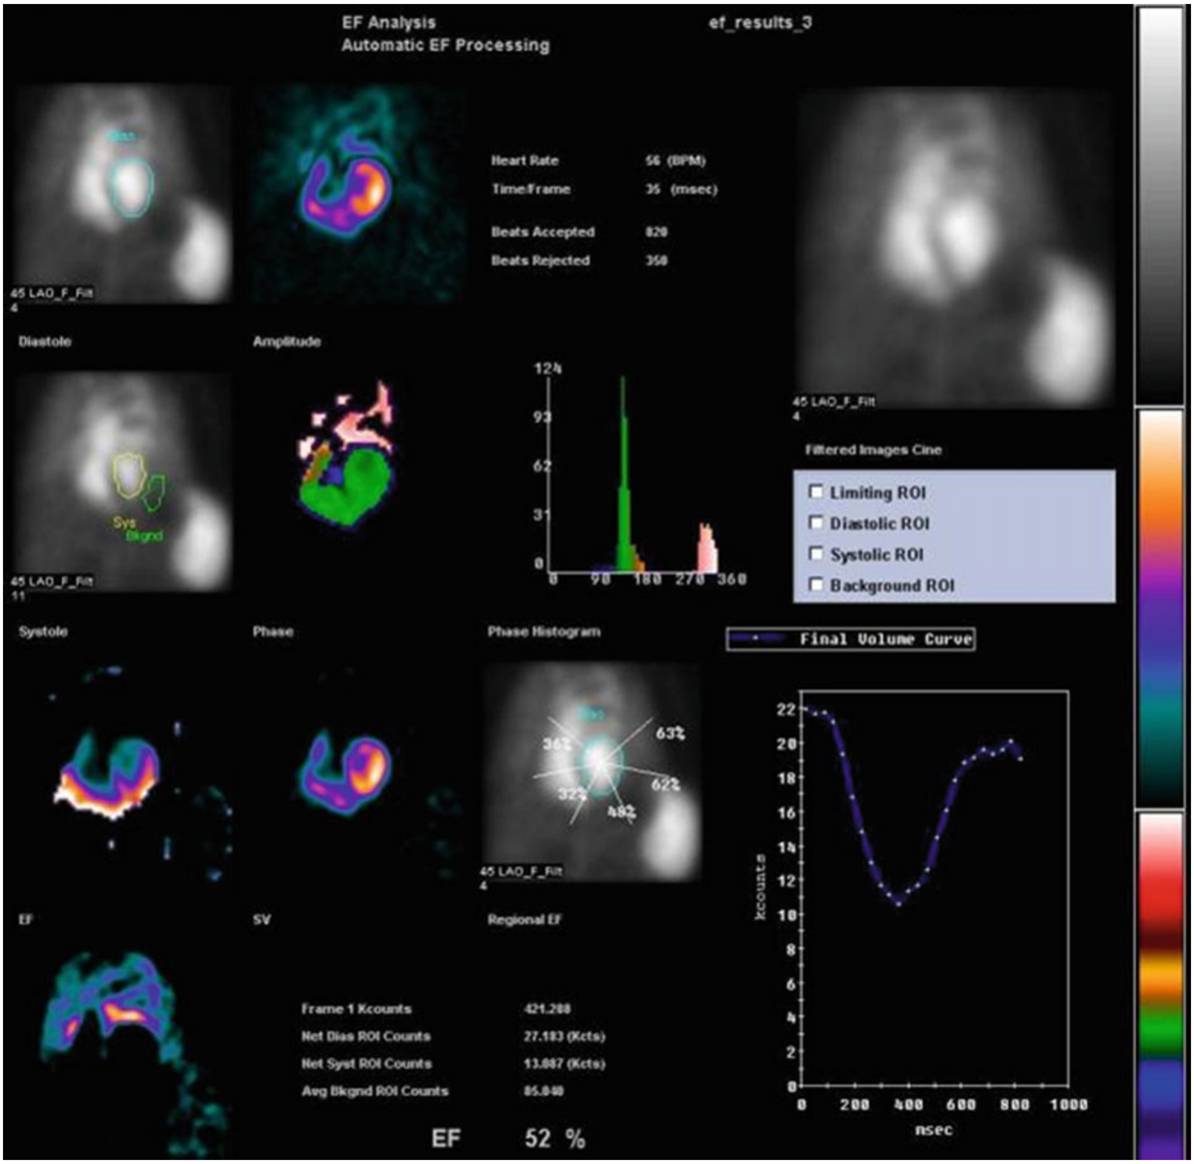

Heart Rate 56 (BPM)
Time/Frame 35 (msec)

Beats Accepted 820
Beats Rejected 350

☐ Limiting ROI
☐ Diastolic ROI
☐ Systolic ROI
☐ Background ROI

Frame 1 Kcounts 421.288
Net Dias ROI Counts 27.183 (Kcts)
Net Syst ROI Counts 13.087 (Kcts)
Avg Bkgnd ROI Counts 85.840

EF 52 %

Fig. 12.24 Normal MUGA example

kidney in renal images. In this case, when the background region extends from one side of the organ ROI to the other, values within the background may be interpolated from the edge values to account for a varying background. The resulting 2-D map of background values is then subtracted pixel by pixel from the organ ROI. An example of this is shown in Fig. 12.28; there, the background was automatically generated in a renal processing program.

12.7.2 Time Activity Curves

In much of nuclear imaging, the changes in radiotracer concentration over time due to mechanical or physiological processes can be more important than visualization of anatomy. The primary tool for analyzing these changes are Time Activity Curves (sometimes abbreviated as TACs), which represent the counts in a

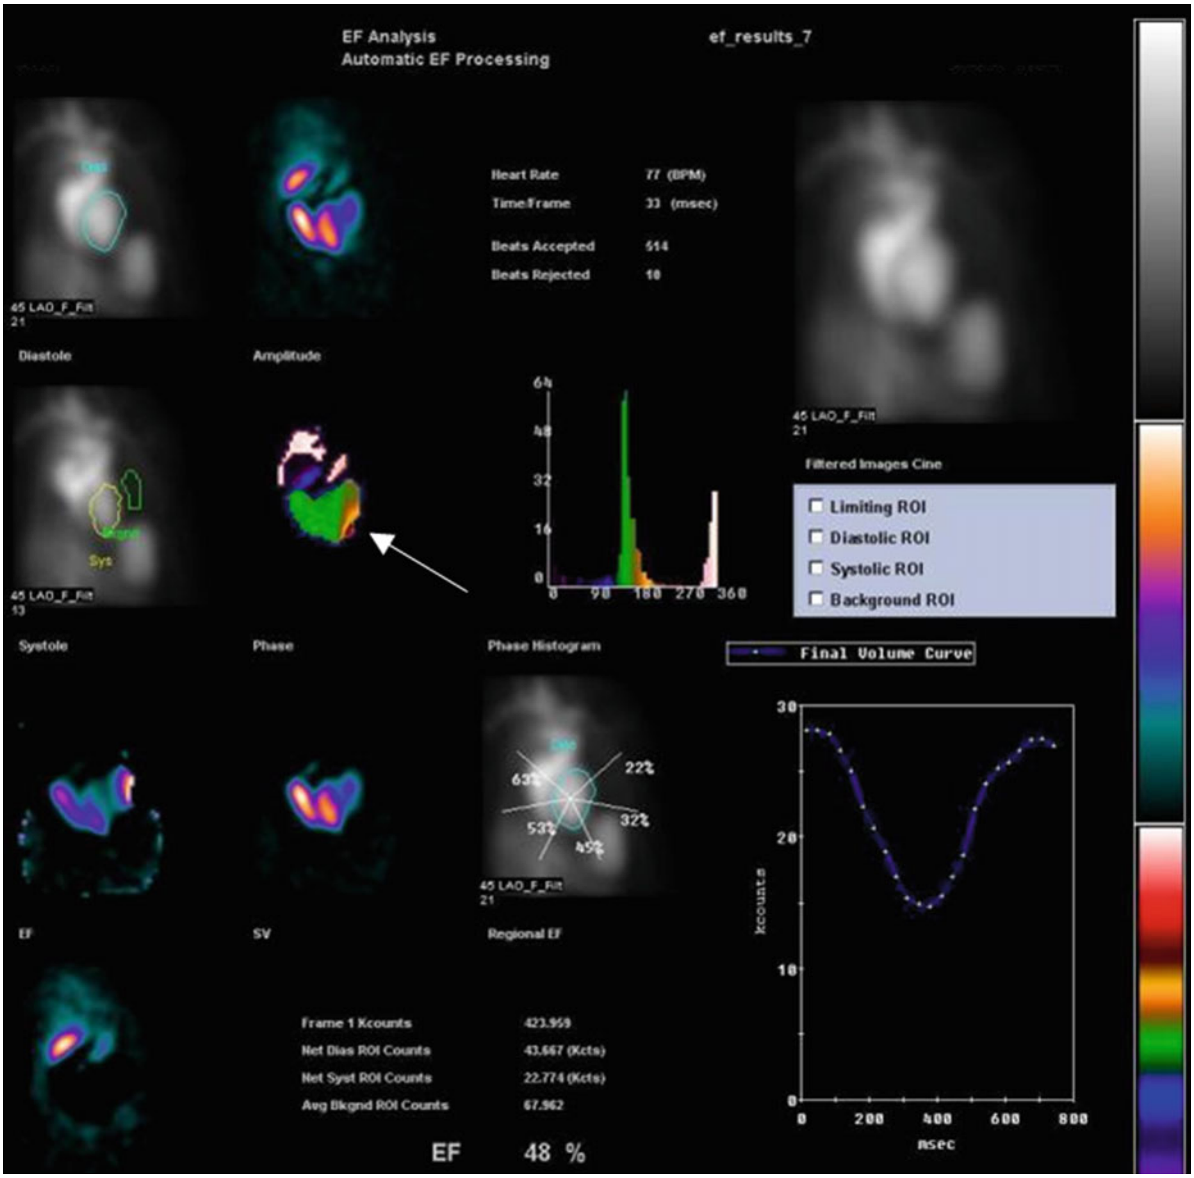

Heart Rate 77 (BPM)
Time/Frame 33 (msec)
Beats Accepted 514
Beats Rejected 10

Frame 1 Kcounts 423.959
Net Dias ROI Counts 43.667 (Kcts)
Net Syst ROI Counts 22.774 (Kcts)
Avg Bkgnd ROI Counts 67.962

EF 48 %

Fig. 12.25 MUGA example showing an apical aneurysm (*arrow*)

region of interest of a dynamic image graphed over time. Examples of widely used time activity curve analysis include estimation of filling and clearance of radiotracer from the kidneys in renograms and of left ventricular volume changes of gated blood pool studies of the heart.

Renogram curves are time activity curves of a radiopharmaceutical as it transits the kidney. Used for the diagnoses of renovascular disease, the curves are used to evaluate the kidney uptake and washout of radiopharmaceuticals with key factors being the time to the peak height of the curve after injection (Tmax), the time it takes for the activity in the kidney to fall to 50% of its peak value and the relative uptake values between the left and right kidneys (T1/2) [27]. Since Tmax is much shorter than T1/2,the dynamic renogram is usually acquired in two phases with 1–3 s frames from the first 60 s and 1–3 min intervals for the remainder of the study. The rapid frame rate of the first phase allows the evaluation of the rapid uptake of the

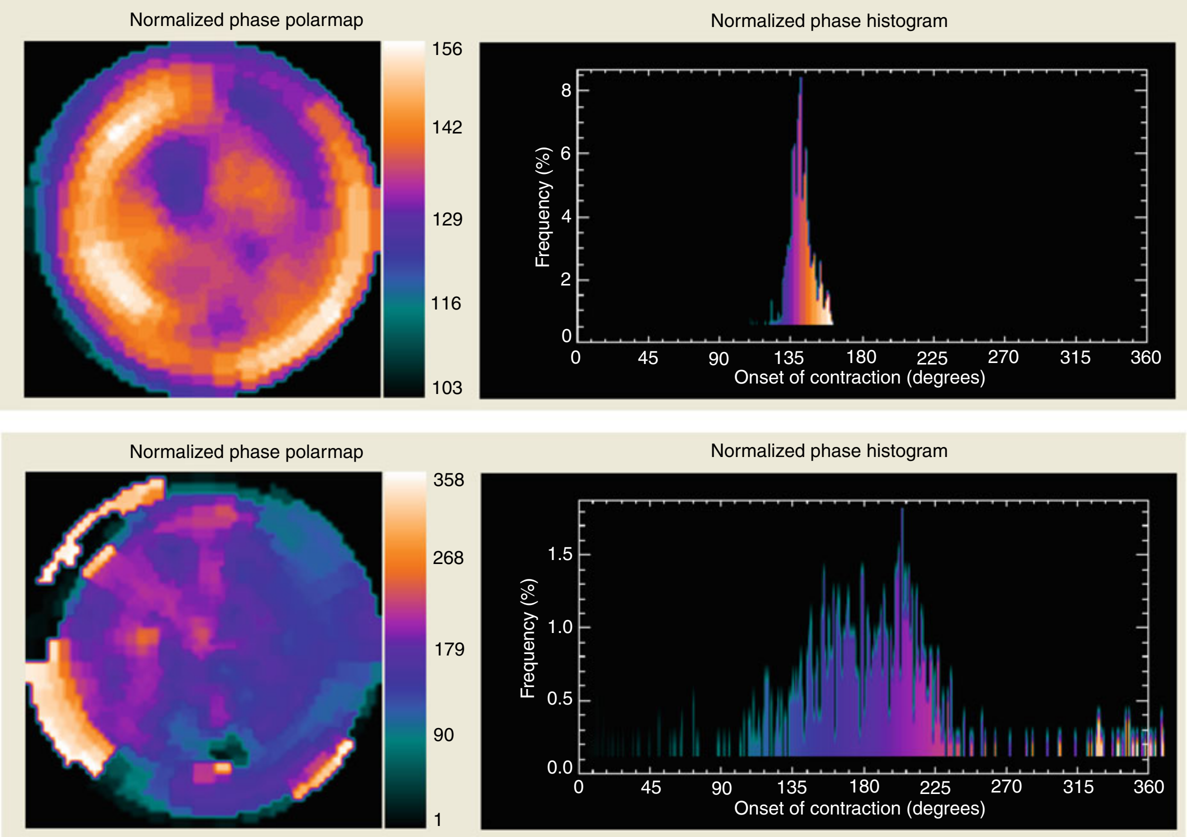

Fig. 12.26 Examples of phase analysis from gated myocardial perfusion SPECT studies. *Top*: normalized phase polar map and histogram from a Normal study, note how tall and narrow the histogram is, indicating that the ventricle is very synchronous. *Bottom*: normalized phase polar map and histogram from an abnormal study, note how spread out the histogram is, indicating severe myocardial dyssynchrony

Fig. 12.27 Examples of a geometric region of interest (*left*) and a free-hand drawn region of interest (*right*) placed around the left ventricle in a radionuclide ventriculogram, or MUGA study

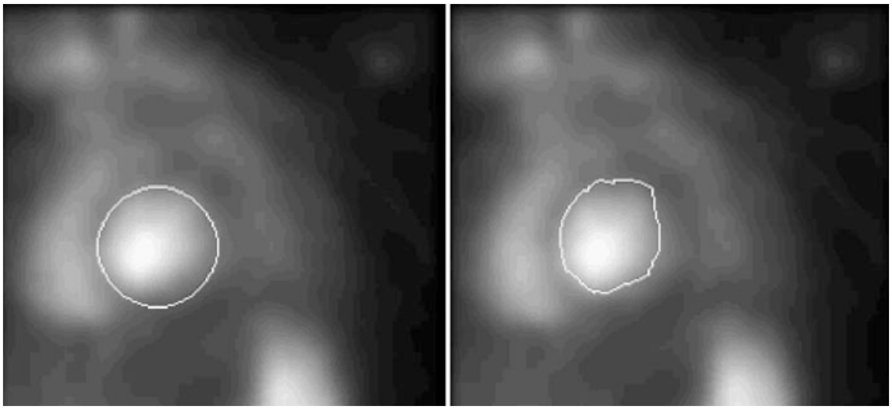

radiopharmaceutical as the kidney clears it from the blood. The second phase uses longer frames to evaluate the washout of the radiopharmaceutical from the kidney. The key ROIs are those placed over the kidneys and background (which may be elliptical around the whole kidney or crescent shaped as described earlier),

Fig. 12.28 An example of automatic ROIs placed on a renal image. The kidney ROIs are determined by a complex algorithm using expected location, image intensities, and edge detection in polar coordinates. Note the rather crescent-shaped background regions placed at the lateral edge of each kidney. Their shapes are based on the detected kidney ROIs, and are a single pixel thick

but ROIs over the bladder and aorta may also provide useful information [28, 29]. The generation and analysis of renogram curves is shown in Fig. 12.29.

Gated blood pool studies (equilibrium radionuclide angiography) of the heart provide a great deal of information about regional heart motion and other factors, but the single most important calculation is that of the ejection fraction calculated from a time activity curve of the left ventricle [27]. The EKG R-to-R interval is usually divided into 24 time segments, each with a corresponding scintigram representing the distribution of blood (which has been labeled with a radiopharmaceutical) at that time in the cardiac cycle. ROIs are drawn over the left ventricle of the heart at end diastole (and preferably at each of the intervals) [30] and counts are extracted. Displayed as a curve, the graph represents the change in volume of the left ventricle at that time in the heart's cycle. The assumption is made that the counts extracted from the ROI are proportional to the volume of blood in the ventricle. To improve the accuracy of this assumption, counts from the ROI that may originate in front of or behind the heart are estimated using a background ROI drawn over an area near the heart but not over areas of significant blood pool. The counts per pixel in the background ROI are calculated, and subtracted from the counts in the volume curve. Selection of a proper background ROI is critical to the calculation as oversubtracting background can falsely elevate the ejection fraction as shown in Fig. 12.30.

Regardless of how time activity curves are generated, there are a number of techniques that maybe used to analyze the data they contain. If the frames are acquired very quickly or there is little activity in the ROI, it may be necessary to use curve-smoothing algorithms to reduce the noise in the curve. If the underlying processes behind the changes in activity represented by the time activity curve are understood, curve fitting techniques may be used to glean further information from the curve.

12.8 Image Segmentation

12.8.1 Thresholding

It may be possible to find a pixel value above which everything is in the region of interest, and below which everything is background. Such a pixel value is called a threshold. Thresholds can be set interactively, but a few methods for obtaining them automatically have been proposed. In some cases, setting a single threshold at a known percentage of the maximum pixel intensity may suffice. Multiple thresholds, for example, one every 10%, can be used to create isocontours; this may give the user an idea of which threshold is the correct one to define the ROI. These approaches are shown in Fig. 12.31.

Another simple approach is to use a histogram of pixel values, so that a graph of number of pixels vs. pixel value is obtained. In some images, for example, brain images, the histogram will be bimodal. A high peak will exist at low pixel values, indicating background pixels; a second peak will occur at higher pixel values, indicating brain pixels. The valley between the two peaks can be used as a threshold to separate brain from background. In other images, where more than one structure is higher in intensity, it may be possible for the user to select which structure is that which is important. The result of thresholding is an image where all areas above the threshold are set to 1, and all areas below are set to 0. An example of this type of approach on a single slice of a brain perfusion image is

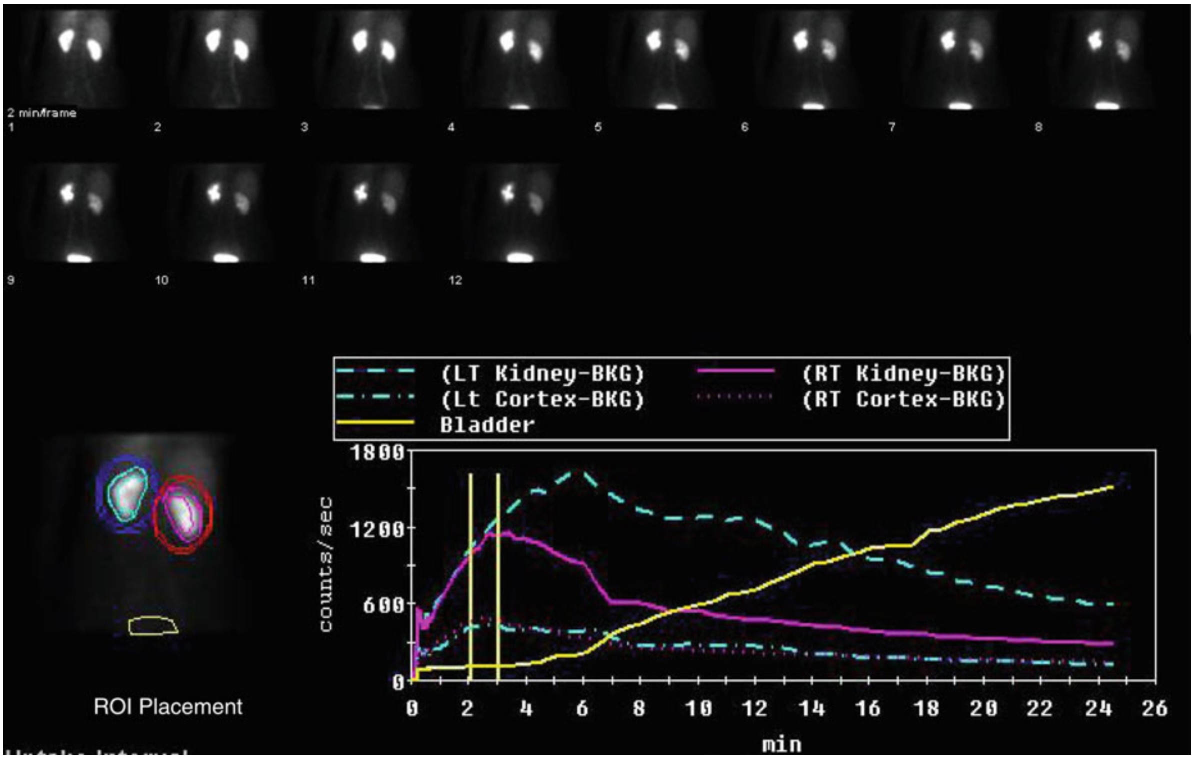

Fig. 12.29 ^{99m}Tc MAG3 renogram curves for each kidney are extracted from ROIs placed over the kidney, the renal cortex (along the outside rim of the kidney), and background. An ROI is also placed over the bladder. This image shows the clearance phase of the renogram, and each frame represents a 2 min interval. The uptake phase of this renogram is shown in Fig. 12.5. With 2 s per frame, counts from the uptake phase are shown in the first 30 s of the curve in this figure. Counts extracted from the kidney ROIs are background corrected before plotting. Counts extracted from the bladder ROI increase throughout the study as the tracer is cleared from the blood by the kidneys. (The authors would like to acknowledge Raghuveer K. Halkar, MD, Emory University, Atlanta, GA, for his assistance with this figure.)

shown in Fig. 12.32. Finding the boundary of the region containing 1s can be performed using the edge tracking method described below, for example.

While most thresholding operations are too simplistic to be widely applicable in defining ROIs, they are often used as an important step in a more complicated processing algorithm. One application where thresholding is particularly important is tumor delineation [31–33]. Again, additional processing is usually applied; however, given that there is rarely any a priori information about tumor shape or location, image intensity is often the most important factor in defining its boundaries.

12.8.2 Edge Detection

Edges of structures in images are generally defined points of quick intensity changes, or high intensity gradients. A portion of the image that falls rapidly from a high value to a low value is often taken to be the edge of a structure. High-pass filtering, or image sharpening, brings out these regions of changing intensity. However, in a complete image slice, there are usually many edges not associated with the region of interest. Conversely, some edges of the ROI may not actually have strong intensity changes associated with them, particularly, if there are low pixel values associated with some physiological process or abnormality. Some method of following the correct boundary is necessary. The problem is, given a potential point on the boundary, or seed, to successively follow the contour and create a closed boundary. This is generally posed as a search algorithm; that is, finding the best contour given some specific constraints that act as a cost function that is to be optimized. Moving from the seed pixel onward, a boundary "cost" is accumulated

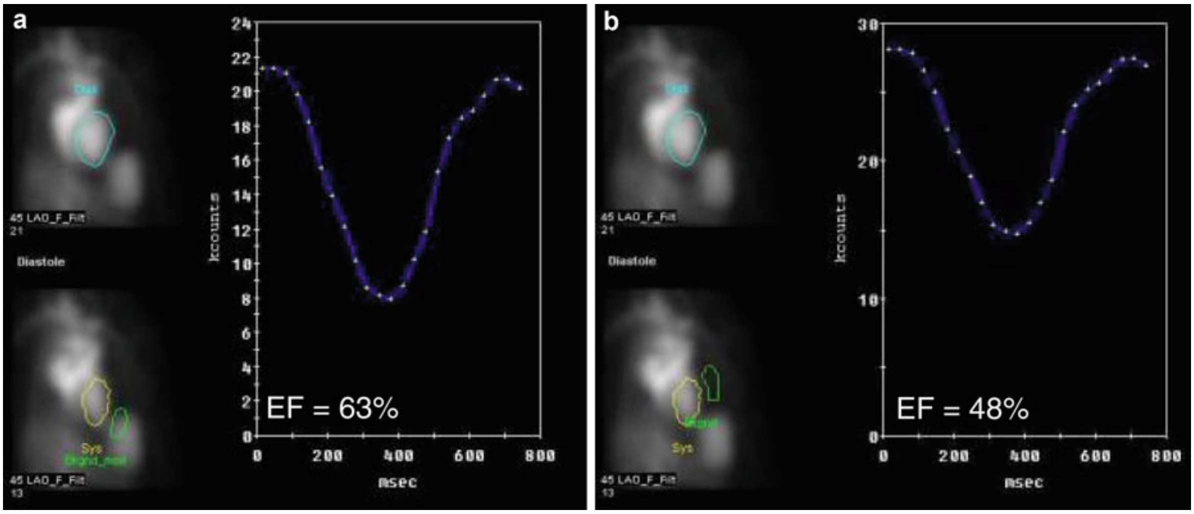

Fig. 12.30 Gated blood pool volume curves are extracted from ROIs drawn over the blood pool of the left ventricle of the ERNA study shown in Fig. 12.5. The volume curve represents background corrected counts at each of the 24 frames or segments of the cardiac cycle. Care must be taken in the placement of the background ROI, since the counts from this region are used to estimate the nonblood pool counts in ventricular ROI. In panel (**a**) the background ROI is placed in a standard location, but overlaps the spleen which has an unusually high concentration of activity. The resulting oversubtraction of background gives an artifactually high estimate of the ejection fraction (EF). In panel (**b**) the background ROI is placed in a more appropriate location for this patient and a much lower EF is calculated. (The authors would like to acknowledge Raghuveer K. Halkar, MD, Emory University, Atlanta, GA, for his assistance with this figure.)

Fig. 12.31 Example of threshold-based regions of interest placed in the same study as Fig. 12.1. A single threshold of 50% (*left*) creates multiple regions of interest; however, the user can indicate which is important. Use of multiple thresholds (*right*) provides isocontours that allow the user to choose more accurately the correct threshold to outline the heart

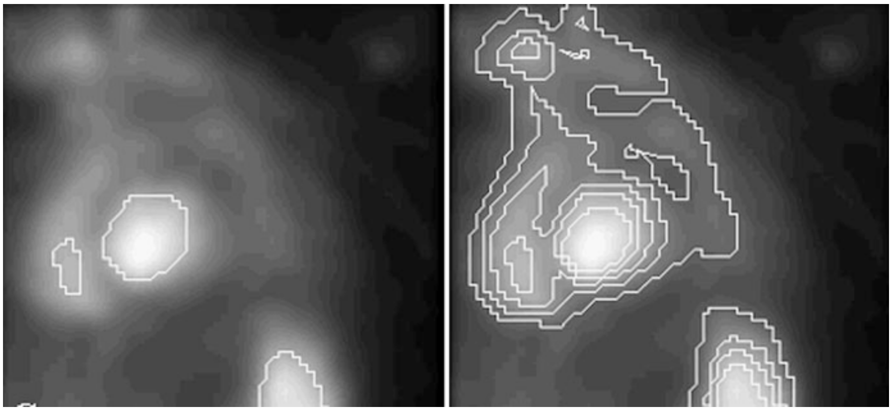

associated with the possible neighboring pixels that may be chosen as the next boundary point. Typical cost functions might include edge magnitude and direction, curvature of a segment, or closeness to a known average contour. One sophisticated approach is described generally for 2D medical image boundary detection [34]; however, note that edge following in 3D becomes much more complicated, and is infrequently used.

12.8.3 Model-Based Approaches

If the expected shape of the region of interest is known, then it may be possible to locate the ROI by fitting a model of the object to the image. This model may be obtained, for example, by hand tracing numerous training images and finding the average shape of the ROI. The model may be as simple as a new image

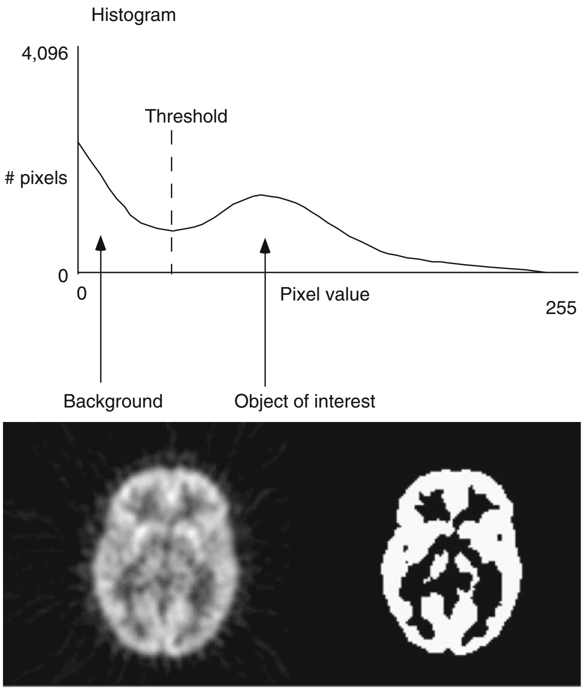

Fig. 12.32 A histogram of image intensities in a slice of a PET FDG brain image (*top*) shows a bimodal distribution. One peak occurs in the background pixel intensities, and another is seen in the gray matter pixel intensities. Choosing a pixel intensity between these two peaks allows the background to be separated from the object of interest, in this case, the gray matter. The original image and the result of the thresholding are seen at the bottom of the figure

containing nothing but this average shape with its values set to the average values obtained from the training set. In one simple approach, registration techniques may be used to fit the model to the new image containing the object. The model is shifted, rotated, and scaled over all possible locations in the new image, and a measure of the match is computed at each. It may even be nonlinearly warped to match the new image. A common metric for evaluating the goodness of the match is the sum of the squared differences between the template and the image. This difference should be small when the model matches well with the image. In a second more complicated approach, the boundaries of the model may be iteratively adapted so that it better fits the object of interest. Constraints on the boundary may be used so that they do not stretch too far beyond the original model and that they maintain a smooth surface. In fact, fitting of the model may also be posed as a registration problem. At any rate, the goal of fitting the model to the image in this manner is to be able to use the boundaries of the transformed model as the boundaries of the ROI in the new image. This approach is becoming more and more common in 3D, as it is able to encode complicated 3D shapes as well as variances in image intensities. When shape is included, the model may be called an active shape model; when it includes image intensity, it may be called an active appearance model.

Most model-based approaches have focused on the brain and heart [35–39]. Often, they were developed as part of a first step in finding abnormal regions of perfusion or metabolism. The model is used to segment a set of normal studies, and then normal ranges within the model ROIs are determined. When a test subject is segmented with the model, the values within the model ROIs can be compared to the normal values in order to automatically determine any regions of hypometabolism or hypoperfusion. Figure 12.33 shows how this approach has been used in quantitation of brain FDG images.

12.8.4 Application-Specific Approaches

Most automatic quantitative programs use very application-specific techniques to isolate the region of interest. This is simply because encoding more specific information into the boundary detection algorithm tends to make it more accurate and more robust. One very common method for automated boundary detection in cardiac imaging, for example, is to perform the operation in polar (for 2D), spherical, or cylindrical (for 3D) coordinates. In this approach, the origin of the coordinate system is placed inside the region of interest. This may be done interactively or automatically. Searching along rays that extend from this origin at set angles, points of high intensity or of high gradient are determined. High-intensity points might be used when searching for the midmyocardium in perfusion images, while high-gradient points might be used when searching for the edge of the ventricular chamber in blood pool images. Because the radius of the boundary should not change greatly between neighboring rays, the search can be constrained, and the boundary can be forced to be smooth. Searches can be limited by placing constraints on how far the ray can extend, and filtering the radii after the boundary detection can further enforce smoothness constraints.

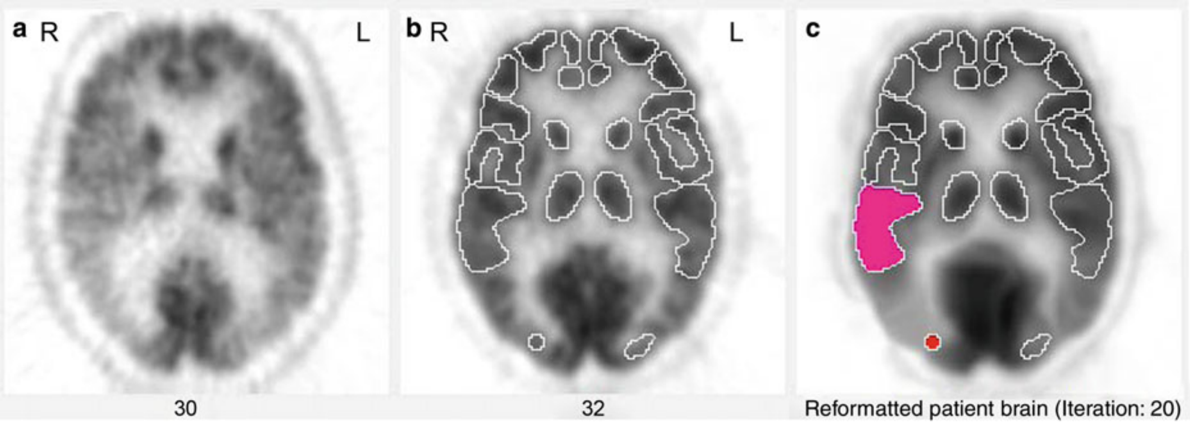

Fig. 12.33 An example of a model-based approach for defining ROIs on a PET FDG image of the brain. (**a**) One slice of a patient's FDG brain image. (**b**) One slice of the "standard," or model, normal FDG image. Regions of interest have already been defined on this standard. (**c**) The patient's image is aligned to the standard, and the ROIs from the standard can then be placed over the patient's brain

This sort of approach is used frequently in organs that are primarily convex, such as the kidney, the brain surface, and the heart [40, 41]. In addition, note that this approach provides a simple way to find the boundary of a thresholded object, given that the object is mostly convex. Figure 12.28 shows an example of how this approach has been applied to find kidney ROIs as a first step in quantitation of renal function [42].

12.9 Three Dimensional Displays

12.9.1 Surface Rendering

Three-dimensional graphics techniques can be used to display both the structure and function of organs visualized in nuclear medicine. If the organ of interest has been accurately segmented using edge detection techniques, then the resulting boundaries can be used to create a 3D surface onto which information about detected counts can be overlaid. The 3D surface may be generated by connecting points of the detected boundaries into triangles, or in some cases, quadrilaterals. Various algorithms exist to perform such triangulations. If the result of edge detection is a set of points, the theoretically most efficient triangularization is called a "Delauney" triangularization [43]. If the result of the edge detection is a binary image, perhaps obtained by thresholding, then an efficient algorithm for generating the triangles of the surface, such as the marching cubes algorithm, can be used [44]. However, if a systematic approach has been used for the edge detection, often a standard triangularization will be sufficient. For example, if the boundary can be divided into the same number of points in each slice of the organ, then connecting the points into triangles is simple. Color can be assigned to each vertex of the triangles in numerous ways. In brain imaging, perfusion or metabolism may be estimated by sampling perpendicularly to the detected surface at each triangle vertex and determining the count value just inside of it. The LV epicardial surface can be color coded with the maximum values obtained at the middle of the myocardium; this is straightforward if both epicardial surface detection and midmyocardial maximum count determination are done using the same spherical coordinate system, for example. Once a set of triangles with vertex colors has been generated, standard software programs are used to display them. Such displays can routinely be rotated in real time and viewed from any angle with current computer power. Sophisticated effects, including highlights, reflection, and transparency may be easily applied. The mathematics behind these effects, as well as a much more detailed discussion of surface rendering can be found in the book by Foley and Van Dam [45]. Such displays have the advantage of showing the actual size and shape of the organ, and the extent and location of

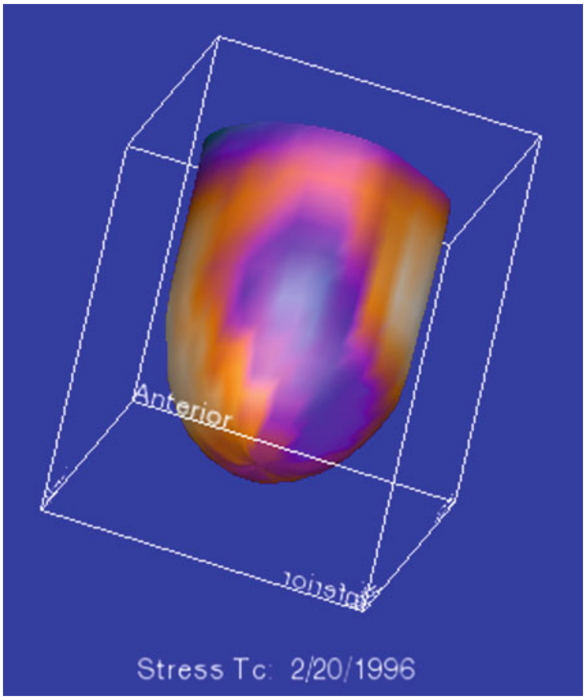

Fig. 12.34 A surface-rendered 3D display created from a SPECT myocardial perfusion study. Once the epicardial and endocardial surfaces have been detected (i.e., a myocardial ROI has been defined), the *x*, *y*, and *z* values of the epicardial surface can be connected into a triangle mesh. The colors used to shade each triangle are based on the maximum values of the perfusion image between the epicardial and endocardial surface at each epicardial point

any defect. Some studies have shown that the 3D models displaying left ventricular perfusion are more accurate for evaluating the size and location of perfusion defects than polar maps [46] or slice-by-slice displays [47]. Figure 12.34 shows a typical surface rendering of the left ventricle, color coded for myocardial perfusion.

12.9.2 Maximum Intensity Projection

A second approach to creating 3D displays from cardiac images generates the myocardial boundaries directly from image voxels without explicit boundary detection. The most useful technique employed in nuclear medicine is maximum intensity projection (MIP) developed by Wallis et al. [48]. Maximum intensity projection involves rotating the 3D tomographic volume into a desired viewing angle and extracting the maximum pixel along each row and column of the rotated volume onto a 2D image plane. To enhance the 3D effect of the image, the volume can be depth-weighted, so that pixels extracted from the front of the volume are scaled more highly than those extracted from the rear of the volume, even if their original intensities were equal. This emphasizes structures in the front of the volume. Since the maximum pixel is always extracted, this type of volume rendering is very useful in blood-pool imaging [49] as well as other hot-spot imaging procedures. Although not directly applicable to cold-spot imaging, a variation of this technique has been successfully applied to liver-spleen imaging and brain perfusion imaging [50]. This variation consists of blurring the original image using a 3D filter and subtracting the blurred image from the original image. The resultant dataset is then rendered using MIP as described above.

12.9.3 Volume Rendering

While maximum intensity projection techniques may be considered a simple type of volume rendering, generally, the term implies a much more complicated algorithm. Like maximum intensity projection, however, volume rendering techniques do not need explicit boundary detection or creation of geometric surfaces. While volume rendering techniques are more useful when used with CT data, it is useful to understand the methods given the current interest in combined PET/CT and SPECT/CT scanners. Volume rendering is a 3D graphics technique first described by Drebin et al. [51] which provides very realistic-appearing visualizations. For a volume of data, such as a stack of 2D tomographic slices, one assigns to every pixel value a transparency, a color, and a reflectivity. The visual process is then simulated by recreating the physical process of light traveling through or bouncing off the pixels that it encounters as it travels through the volume. For any given volume of data, the main problems involve the correct assignment of the pixel properties. Selecting a pixel intensity threshold above which everything is opaque, and below which everything is transparent, for example, will provide a 3D image with hard surfaces at that threshold. In nuclear medicine, the use of such a threshold often results in normal pixels being assigned a value of 1, and abnormal pixels being assigned a value of 0, so that the 3D display depicts a "hole" in the surface. Such techniques have been used primarily in brain images; however, they have also been applied to cardiac images. Unfortunately, very few medical images are accurately segmented with such simplistic threshold techniques, and note that the size of the abnormality, or hole, is completely dependent on choosing the right threshold. CT images are in fact probably the best adapted to this simple method of assigning pixel properties, since Hounsfield units are related to the electron density of particular tissues, and thus, bone can be made opaque and white, while muscle tissue can be made translucent and red, for example. In fact, a basic 3D display method similar to volume rendering was described quite early for use in CT images by Hoehne et al. [52]. While volume rendering of bone in CT images is generally quite robust, most 3D displays of soft tissue generated from CT require large amounts of preprocessing to remove unwanted structures or improve the segmentation. The accuracy of this preprocessing is generally related to the quality of the images, and user intervention may be required to produce a useable 3D display. In addition, large amounts of memory and processing power are required for volume rendering of large datasets. Nevertheless, the realism, high resolution, and flexibility of volume renderings give them an advantage over surface renderings in many cases; this can be appreciated in Fig. 12.35.

Fig. 12.35 A cardiac CT study may be volume rendered by assigning colors and transparencies to the Hounsfield units within it, and then reproducing the effect of light being transmitted and reflected from it. In this case, sophisticated segmentation and additional user interaction must be used to eliminate noncardiac structures from the CT prior to 3D rendering

12.10 Principles of Image Registration

As mentioned above in the model-based approach for segmentation, alignment of a presegmented model with a specific patient's image can be used to find regions or volumes of interest in the new image. In addition, alignment (also called registration or fusion) is often used with nuclear medicine for other important purposes. Low-resolution nuclear medicine images are aligned with higher-resolution anatomic images of the same patient to help identify the exact location of a hot or cold spot in the nuclear image. In this application, it may be sufficient to transform the test image using only linear transformations. Also, in order to use many quantitative programs, the patient's specific nuclear image must be put into a standard space. More and more, this standardization implies that the test image is being aligned to a standard, or atlas, image. In this application, frequently, a nonlinear warping transformation is needed.

Registration is generally posed as an iterative optimization problem. The differences between two images are to be minimized over some transformation. Thus, two separate pieces of software are needed: one to measure the difference, or cost function, and a second to determine how to change the image with respect to the standard to reduce this difference function at each iteration.

While there are many options for computing the cost function, a very popular approach is called the mutual information criterion [53, 54]. Mutual information can be described as the joint histogram between the two images. When two images are perfectly aligned, the joint histogram should ideally exist only along the line of identity. As the images are moved away from each other, the joint histogram also spreads out. Mutual information is a measure of the spread of the joint histogram. An example of this can be seen in Fig. 12.36.

For linear registration, that is, registration that includes only translation, scale, and rotation, generally closed form optimization methods are used to iteratively find the transformation that optimizes mutual information between two images. One such method is the Newton–Raphson method, which finds the roots of the derivative of the cost function, which occur at function minima. Woods et al. used this method in early work aligning PET brain images to MRI, and an example is shown in Fig. 12.37.

For nonlinear registration, or warping, the challenge is to find a good match between the images while making sure that the warping is smooth. This can be achieved by adding a regularizer, or smoothness constraint. For example, the warping can be constrained to be elastic, so that small deformations are easily achieved, but larger movements are "penalized," as they require more "force," similar to any elastic deformation. This is the approach taken by Bajcsy et al. in early work aligning MR brain images [55]. Other constraints based on the known physical properties of the objects being aligned may be used to help find a good solution to the warping problem. For example, it is possible to constrain all deformations to be volume-preserving, so that pixels can change shape but must keep their original volumes, as suggested by Haber et al. [56]. Minimization of these more complicated functions may be attained by posing them as Euler-Lagrange equations, whose solution is a local optimum. An example of warping a brain image to match an atlas was shown in Fig. 12.33.

However, nonlinear warping has some major drawbacks. First, it can be quite time-consuming, especially for large images. One additional source of computational complexity is the nonlinear interpolation required. For example, a method such as thin plate splines may be required to interpolate multiple unevenly spaced points to new, unevenly spaced pixel locations [57]. Another problem is that if the proper constraints are not used in the warping algorithm, physically impossible deformations may occur. Parts of the image may "fold" onto one another, or more than one point may be mapped into a single location. Finally, it is very difficult to evaluate the

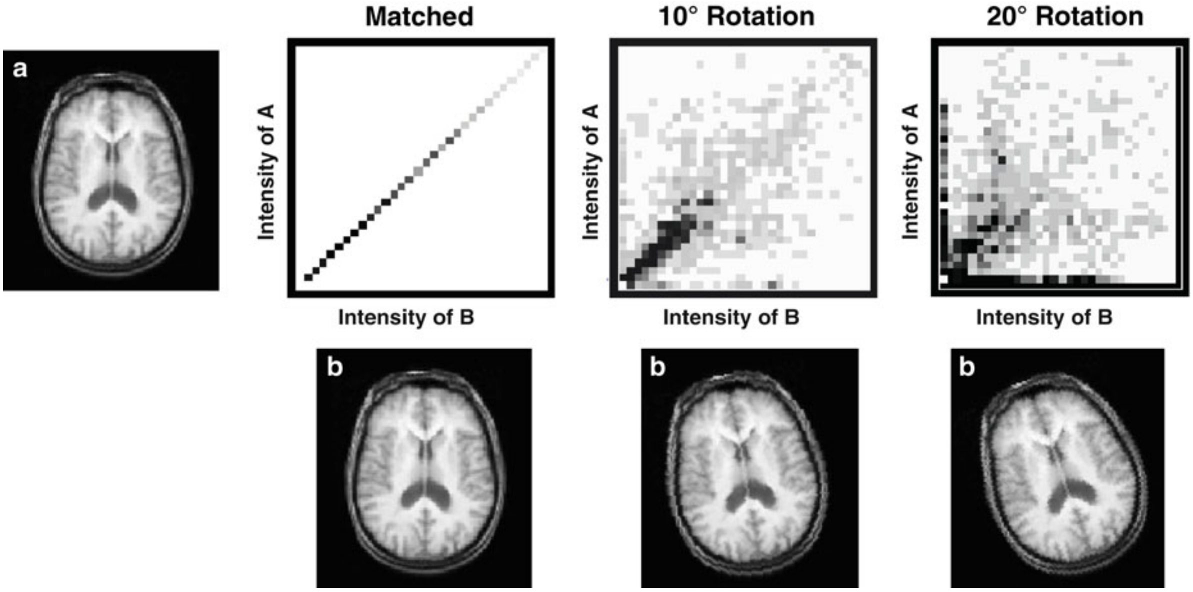

Fig. 12.36 Mutual information criteria for image registration. Image (**b**) is to be aligned with Image (**a**). At each rotation of (**b**) the joint histogram is created. When the images are in exact alignment, the joint histogram has values only on the line of identity. As the images are rotated more and more out of alignment, the joint histogram becomes more diffuse. Mutual information is a measure of the spread of the joint histogram

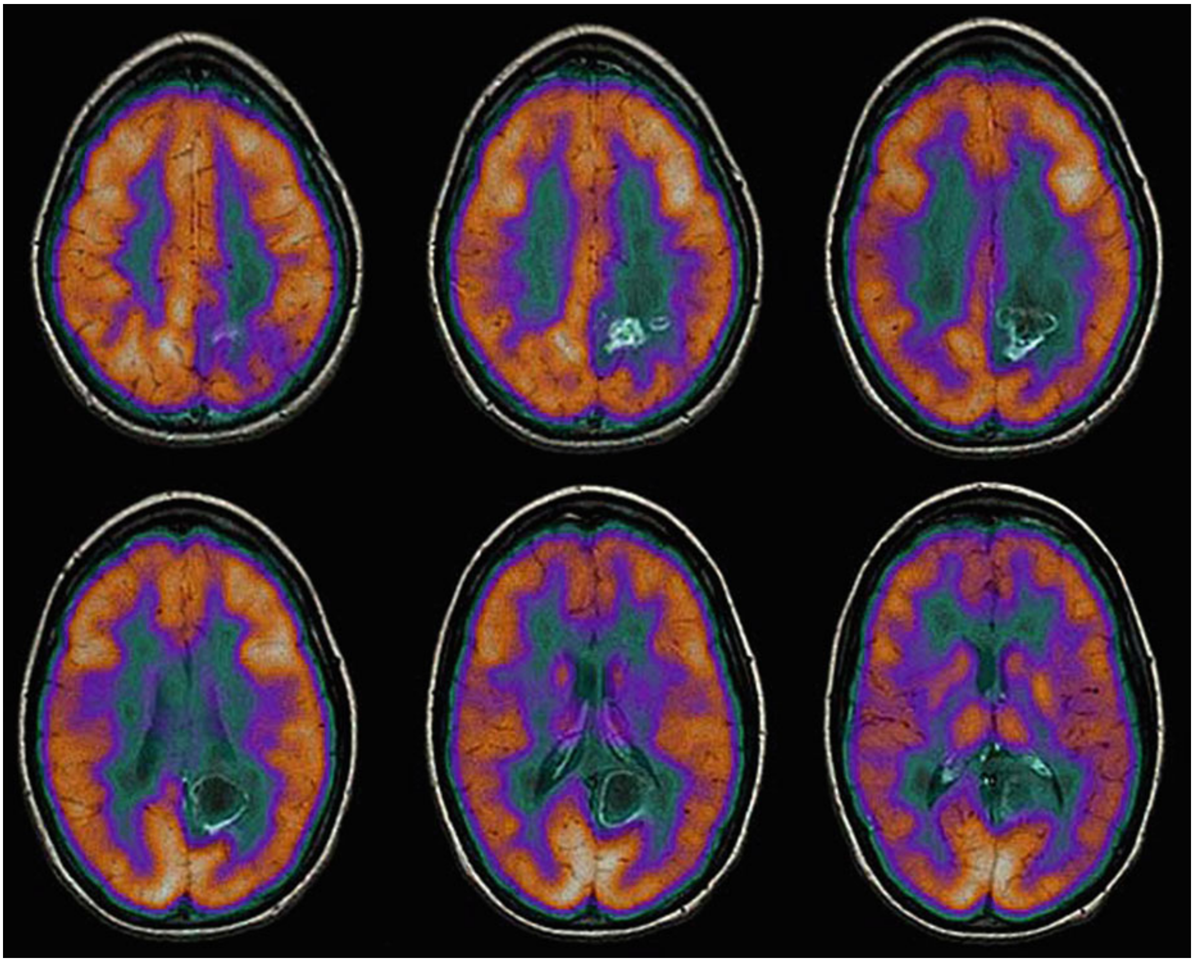

Fig. 12.37 An example of using linear registration to align a PET brain image to an MRI of the same patient. A function based on the ratio of the image pixel values was used as the difference measure, and this was minimized using a Newton–Raphson method. In this image, the PET is overlaid in color over the grayscale MRI

results of a warping algorithm. What may look like a reasonable deformation may not be accurate, as there is often no "ground truth." The user should be sure to accept nonlinear alignments only after careful assessment.

12.11 Image Normalization

As mentioned in the above section on Image Presentation, images that are acquired in Nuclear Medicine are generally constructed from the counts that are received at each individual pixel. This means that the counts in each pixel are relative (or related) to each other and the actual activity is being measured in the object; however, without additional processing, these counts are not absolute in nature. Furthermore, there may be some physiologic phenomena that occur that obscures the object of interest.

12.11.1 Extraobject Activity

In Cardiac SPECT, it is not uncommon for the gallbladder or a loop of bowel to have more activity than the heart, thereby causing the heart to appear reduced in the images. Figure 12.38 shows a Cardiac SPECT study where a loop of bowel (seen in the bottom of the image) is much brighter than any of the pixels in the heart, thereby causing the counts in the heart to seem very low. To correct this problem, the study needs to be "normalized" to the counts in the myocardium, and not the counts in the entire image. This normalization (or scaling) process simply rescales the image such that the maximum displayed value is now equal to the maximum in the heart, instead of the maximum in the image or volume. In this example, the maximum activity in the bowel was 839 counts and the maximum activity in the heart was 530 counts. Figure 12.39 shows the same image scaled to the maximum of the myocardium, instead of the maximum of the volume. Notice how the loop of bowel is now very bright but the myocardium looks much more normal. This phenomena can be seen in many different imaging scenarios, such as the bone scan shown previously (Fig. 12.4). Notice how hot the bladder is compared to the rest of the image in

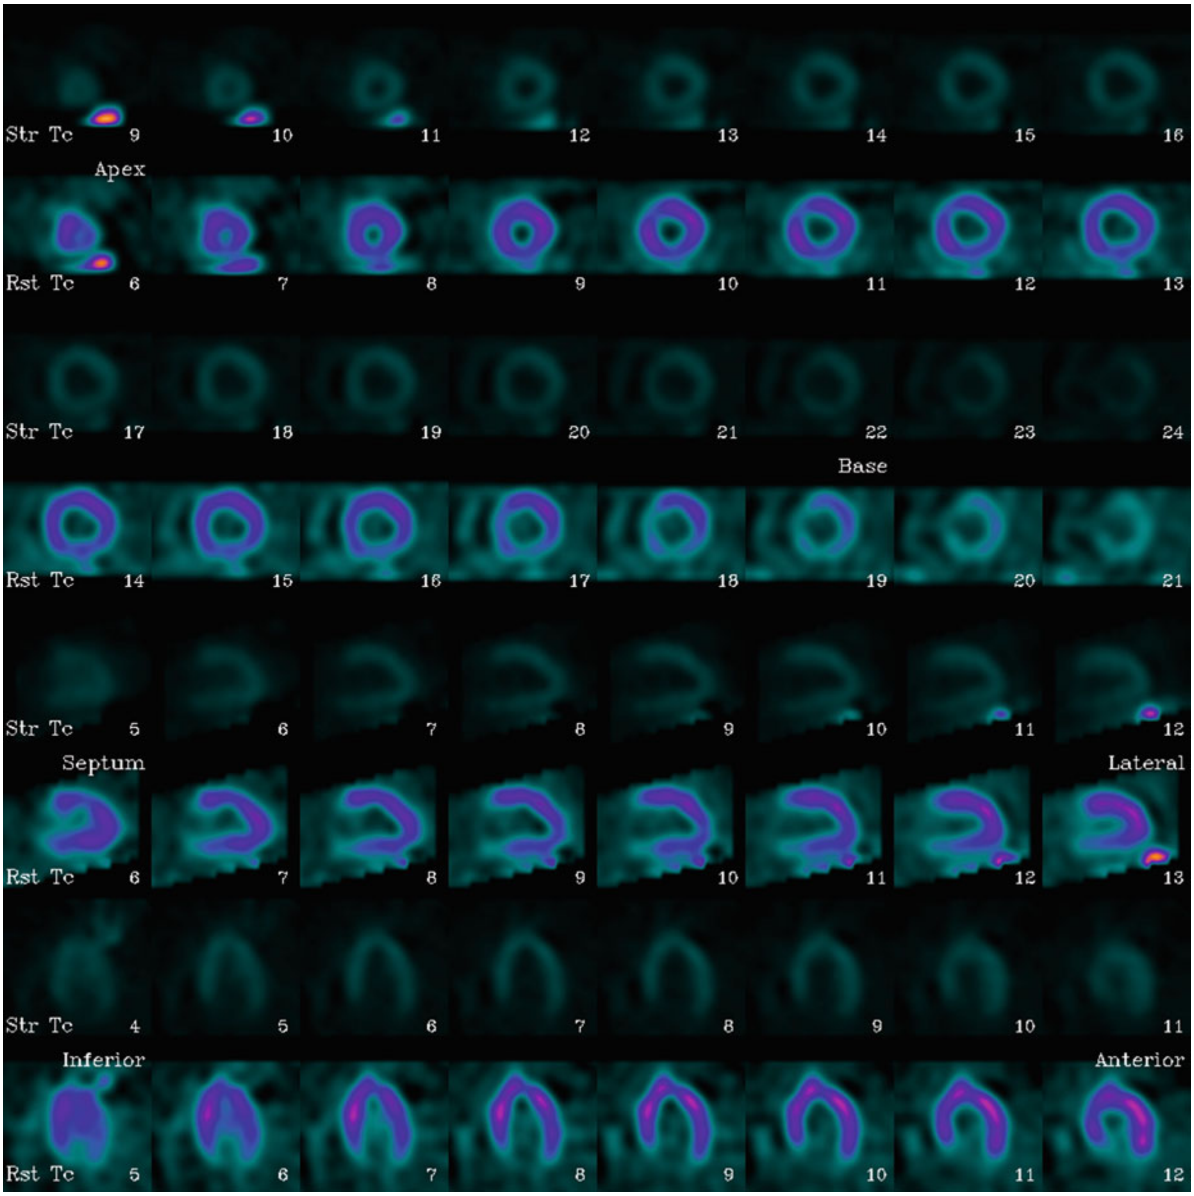

Fig. 12.38 Example of normalization, scaled to the maximum pixel in the volume

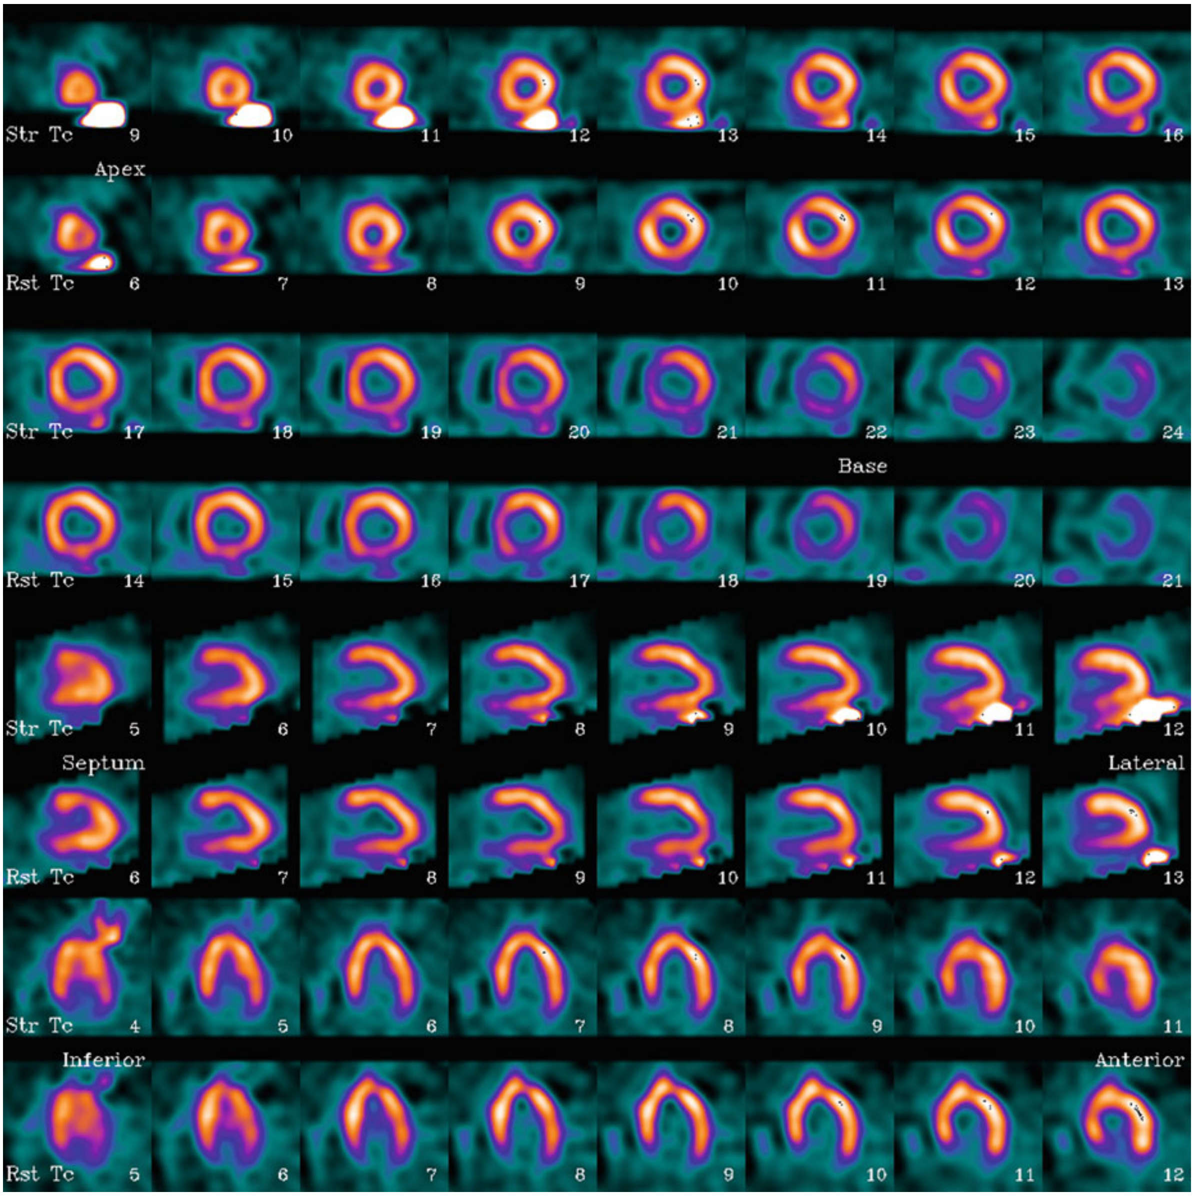

Fig. 12.39 Example of normalization, scaled to the maximum activity in the myocardium

the left-hand image. There are several ways to correct for this: (a) masking out the hot bladder and renormalizing to the maximum count in the image, (b) drawing a region-of-interest around most of the image (but excluding the bladder) and renormalizing to the maximum count in the ROI and (c) using the color table to renormalize the image by setting a new maximum value for display. In this case, the display maximum was reset to 100 in the right-hand image, thereby allowing the clinical details in the image to be seen.

12.11.2 Absolute Normalization/Quantification

Another area of interest and research over the years has been in the area of absolute quantification, or converting the relative counts from the image into absolute measures of activity. Briefly, this involves correcting the image for attenuation, scatter, and resolution as well as measuring the counts in the arterial blood supply and then some form of modeling to predict the activity in the target organ. As you can see, this can be a complicated problem, but it can have big rewards. For instance, in Cardiac SPECT, if a patient has balanced disease (equal disease in all of the coronary territories), then all of the heart will look "relatively" the same which could easily be interpreted as a normal scan. However, if we could get an absolute measure of the activity in the heart, then we would be able to tell that the activity was below normal and that the scan was actually abnormal, though in "relative" terms it looked normal [58–60].

12.11.3 Normalization for Database Quantification

Because of the "relative" nature of the counts acquired in Nuclear Medicine, one of the approaches to overcome this is to compare a patient's scan to a database of normal volunteers. This technique is used in many imaging scenarios such as Cardiac, Renal, Neuro, and others. In order for this technique to work correctly, the patient's scan needs to be normalized or scaled to the same range as the data in the normal database. For instance, Fig. 12.40a shows the average or mean counts for the male and female normal files for a typical 1 day ^{99m}Tc Sestamibi myocardial SPECT scan. In order to compare a patient's scan to this normal file, it must first be scaled into the same range as the normal file, using the same region of activity. For instance, if the normal file had a maximum count of 1,000, which occurred in the anterior wall, and the anterior wall of the patient's scan had a maximum count of 1,200, then the patient's scan would need to be multiplied by 0.83 before comparison to the normal file. Figure 12.40 shows an abnormal patient after comparison to the normal file, in which abnormal areas have been turned black (note the large anterior and inferior wall defects) [61–63].

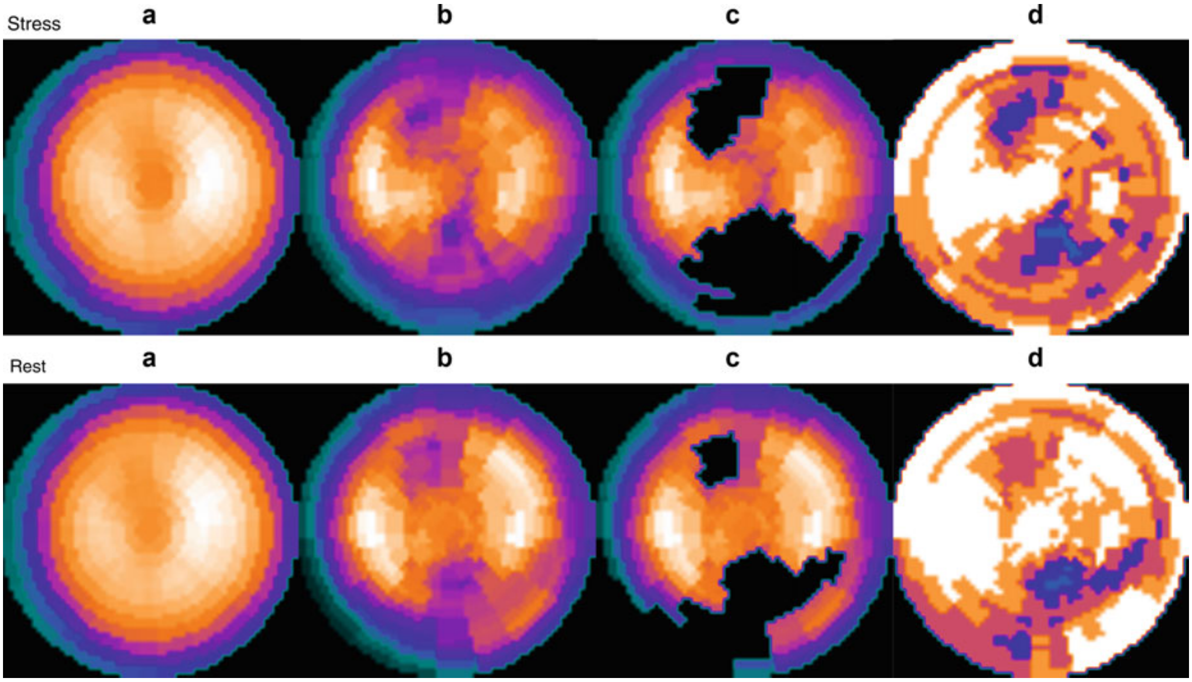

Fig. 12.40 Example of SPECT database comparison (**a**) Stress and rest female normal file for Tc-99m 1 day sestamibi. (**b**) Stress and rest raw polar maps (counts extracted from the stress and rest short-axis datasets). (**c**) Stress and rest blackout polar maps after comparison to the normal database; blacked-out areas show regions of significant hypoperfusion. (**d**) Stress and rest standard deviation polar maps, *darker areas* show regions that are more standard deviations from the normal file than lighter areas

12.12 Conclusion

Image processing is an integral part of the daily nuclear medicine routine and many physiological as well as functional parameters can be extracted from the different nuclear medicine procedures by various processing tools. Not only is image processing able to convey this desired information, but also a variety of display tools exist to assist in image interpretation and patient diagnosis. Hopefully, this overview has given you an appreciation for how difficult it would be to work with any kind of imaging modality without the use of computers. In fact, computers are a very necessary part of any imaging modality and will only become more important as we continue to push the envelope of imaging.

References

1. Cassen B, Curtis L, Reed CW (1949) A sensitive directional gamma ray detector. Technical Report #UCLA-49 (OSTI ID: 4434981), University of California, Los Angeles
2. Anger HO (1957) A new instrument for mapping gamma-ray emitters. Biology and Medicine Quarterly Report for October, November, December 1956, Report #UCRL-3653 (OSTI ID: 4354301), p. 51
3. Kuhl DE, Edwards RQ (1963) Image separation radioisotope scanning. Radiology 80:653–662
4. Jaszczak RJ (2006) The early years of single photon emission computed tomography (SPECT): an anthology of selected reminiscences. Phys Med Biol 51:R99–R115
5. Anger HO (1964) Scintillation camera with multichannel collimators. J Nucl Med 5:515–531
6. Gotway MB, Leung JW, Gooding GA, Litt HI, Reddy GP, Morita ET, Webb WR, Clark OH, Higgins CB (2002) Hyperfunctioning parathyroid tissue: spectrum of appearances on noninvasive imaging. AJR Am J Roentgenol 179:495–502
7. Hoffman EJ, Huang SC, Phelps ME (1979) Quantitation in positron emission computed tomography: 1. Effect of object size. J Comput Assist Tomogr 3:299–308
8. Demirkaya O, Al Mazrou R (2007) Performance test data analysis of scintillation cameras. IEEE Trans Nucl Sci 54:1506–1515
9. Anderson S (2005) Collins English dictionary, 7th edn. HarperCollins, Glasgow
10. Galt JR, Garcia EV, Nowak DJ (1986) Filtering in Frequency Space. J Nucl Med Technol 14:152–160
11. Hansen CL (2002) Digital image processing for clinicians, part I: basics of image formation. J Nucl Cardiol 9:343–349
12. Hansen CL (2002) Digital image processing for clinicians, part II: filtering. J Nucl Cardiol 9:429–437
13. Hansen CL (2002) Digital image processing for clinicians, part III: SPECT reconstruction. J Nucl Cardiol 9:542–549
14. Zubal IG, Wisniewski G (1997) Understanding Fourier space and filter selection. J Nucl Cardiol 4:234–243
15. Cooke CD, Garcia EV, Cullom SJ, Faber TL, Pettigrew RI (1994) Determining the accuracy of calculating systolic wall thickening using a fast Fourier transform approximation: A simulation study based on canine and patient data. J Nucl Med 35:1185–1192
16. Chen J, Faber TL, Cooke CD, Garcia EV (2008) Temporal resolution of multiharmonic phase analysis of ECG-gated myocardial perfusion SPECT studies. J Nucl Cardiol 15:383–391
17. Galt JR, Garcia EV, Robbins WL (1990) Effects of myocardial wall thickness on spect quantification. IEEE Trans Med Imaging 9:144–150
18. Faber TL, Cooke CD, Folks RD, Vansant JP, Nichols KJ, DePuey EG, Pettigrew RI, Garcia EV (1999) Left ventricular function and perfusion from gated SPECT perfusion images: an integrated method. J Nucl Med 40:650–659
19. Pflugfelder PW, Sechtem UP, White RD, Higgins CB (1988) Quantification of regional myocardial function by rapid cine MR imaging. Am J Roentgenol 150:523–529
20. DePuey EG, Rozanski A (1995) Using gated technetium-99m-sestamibi SPECT to characterize fixed myocardial defects as infarct or artifact. J Nucl Med 36:952–955
21. Chen J, Garcia EV, Folks RD, Cooke CD, Faber TL, Tauxe EL, Iskandrian AE (2005) Onset of left ventricular mechanical contraction as determined by phase analysis of ECG-gated myocardial perfusion SPECT imaging: development of a diagnostic tool for assessment of cardiac mechanical dyssynchrony. J Nucl Cardiol 12:687–695
22. Chen J, Henneman MM, Trimble MA, Bax JJ, Borges-Neto S, Iskandrian AE, Nichols KJ, Garcia EV (2008) Assessment of left ventricular mechanical dyssynchrony by phase analysis of ECG-gated SPECT myocardial perfusion imaging. J Nucl Cardiol 15:127–136
23. Henneman MM, Chen J, Dibbets-Schneider P, Stokkel MP, Bleeker GB, Ypenburg C, van der Wall EE, Schalij MJ, Garcia EV, Bax JJ (2007) Can LV dyssynchrony as assessed with phase analysis on gated myocardial perfusion SPECT predict response to CRT? J Nucl Med 48:1104–1111
24. Marsan NA, Henneman MM, Chen J, Ypenburg C, Dibbets P, Ghio S, Bleeker GB, Stokkel MP, van der Wall EE, Tavazzi L, Garcia EV, Bax JJ (2008) Left ventricular dyssynchrony assessed by two three-dimensional imaging modalities: phase analysis of gated myocardial perfusion SPECT and tri-plane tissue Doppler imaging. Eur J Nucl Med Mol Imaging 35:166–173
25. Henneman MM, Chen J, Ypenburg C, Dibbets P, Bleeker GB, Boersma E, Stokkel MP, van der Wall EE, Garcia EV, Bax JJ (2007) Phase analysis of gated myocardial perfusion single-photon emission computed tomography compared with tissue Doppler imaging for the assessment of left ventricular dyssynchrony. J Am Coll Cardiol 49:1708–1714
26. Trimble MA, Borges-Neto S, Honeycutt EF, Shaw LK, Pagnanelli R, Chen J, Iskandrian AE, Garcia EV, Velazquez EJ (2008) Evaluation of mechanical dyssynchrony and myocardial perfusion using phase analysis of gated SPECT imaging in patients with left ventricular dysfunction. J Nucl Cardiol 15:663–670

27. Taylor A, Schuster DM, Alazraki NP (2006) A clinician's guide to nuclear medicine, 2nd edn. Society of Nuclear Medicine, Reston
28. Taylor A Jr, Corrigan PL, Galt J, Garcia EV, Folks R, Jones M, Manatunga A, Eshima D (1995) Measuring technetium-99m-MAG3 clearance with an improved camera-based method. J Nucl Med 36:1689–1695
29. Taylor AT Jr, Fletcher JW, Nally JV Jr, Blaufox MD, Dubovsky EV, Fine EJ, Kahn D, Morton KA, Russell CD, Sfakianakis GN, Aurell M, Dondi M, Fommei E, Geyskes G, Granerus G, Oei HY (1998) Procedure guideline for diagnosis of renovascular hypertension. Society of Nuclear Medicine. J Nucl Med 39:1297–1302
30. Corbett JR, Akinboboye OO, Bacharach SL, Borer JS, Botvinick EH, DePuey EG, Ficaro EP, Hansen CL, Henzlova MJ, Van Kriekinge S (2006) Equilibrium radionuclide angiocardiography. J Nucl Cardiol 13:e56–e79
31. Erdi YE, Mawlawi O, Larson SM, Imbriaco M, Yeung H, Finn R, Humm JL (1997) Segmentation of lung lesion volume by adaptive positron emission tomography image thresholding. Cancer 80:2505–2509
32. Jentzen W, Freudenberg L, Eising EG, Heinze M, Brandau W, Bockisch A (2007) Segmentation of PET volumes by iterative image thresholding. J Nucl Med 48:108–114
33. Brambilla M, Matheoud R, Secco C, Loi G, Krengli M, Inglese E (2008) Threshold segmentation for PET target volume delineation in radiation treatment planning: the role of target-to-background ratio and target size. Med Phys 35:1207–1213
34. Mortensen E, Morse B, Barrett W, Udupa J (1992) Adaptive boundary detection using live-wire 2-dimensional dynamic-programming. In: Proceedings of the Computers in Cardiology, pp 635–638
35. Declerck J, Feldmar J, Goris ML, Betting F (1997) Automatic registration and alignment on a template of cardiac stress and rest reoriented SPECT images. IEEE Trans Med Imaging 16:727–737
36. Slomka PJ, Hurwitz GA, Stephenson J, Cradduck T (1995) Automated alignment and sizing of myocardial stress and rest scans to three-dimensional normal templates using an image registration algorithm (see comment). J Nucl Med 36:1115–1122
37. Mykkanen J, Tohka J, Luoma J, Ruotsalainen U (2005) Automatic extraction of brain surface and mid-sagittal plane from PET images applying deformable models. Comput Meth Programs Biomed 79:1–17
38. Minoshima S, Koeppe RA, Frey KA, Kuhl DE (1994) Anatomic standardization: linear scaling and nonlinear warping of functional brain images. J Nucl Med 35:1528–1537
39. Minoshima S, Frey KA, Koeppe RA, Foster NL, Kuhl DE (1995) A diagnostic approach in Alzheimer's disease using three-dimensional stereotactic surface projections of fluorine-18-FDG PET. J Nucl Med 36:1238–1248
40. Garcia EV, Cooke CD, Van Train KF, Folks RD, Peifer JW, DePuey EG, Maddahi J, Alazraki NP, Galt JR, Ezquerra NF, Ziffer JA, Areeda JS, Berman DS (1990) Technical aspects of myocardial SPECT imaging with technetium-99m sestamibi. Am J Cardiol 66:23E–31E
41. Germano G, Kavanagh PB, Waechter P, Areeda J, Van Kriekinge S, Sharir T, Lewin HC, Berman DS (2000) A new algorithm for the quantitation of myocardial perfusion SPECT. I: technical principles and reproducibility (see comment). J Nucl Med 41:712–719
42. Garcia E, Folks R, Pak S, Taylor A (2008) Automatic definition of renal regions-of-interests (ROIs) from MAG3 renograms in patients with suspected renal obstruction. J Nucl Med (Meeting Abstracts) 49:386P
43. Delaunay B (1934) Sur la sphere vide. A memoire de Georges Voronoi. Izv Akad Nauk SSSR, Otdelenie Matematicheskih i Estestvennyh Nauk 7:793–800
44. Lorensen WE, Cline HE (1987) Marching cubes: a high resolution 3D surface construction algorithm. SIGGRAPH Comput Graph 21:163–169
45. Foley JD, Phillips RL, Hughes JF, van Dam A, Feiner SK (1994) Introduction to computer graphics. Addison-Wesley, Longman
46. Cooke CD, Vansant JP, Krawczynska EG, Faber TL, Garcia EV (1997) Clinical validation of three-dimensional color-modulated displays of myocardial perfusion. J Nucl Cardiol 4:108–116
47. Santana CA, Garcia EV, Vansant JP, Krawczynska EG, Folks RD, Cooke CD, Faber TL (2000) Three-dimensional color-modulated display of myocardial SPECT perfusion distributions accurately assesses coronary artery disease. J Nucl Med 41:1941–1946
48. Wallis JW, Miller TR (1990) Volume rendering in three-dimensional display of SPECT images (see comments). J Nucl Med 31:1421–1428
49. Miller TR, Wallis JW, Sampathkumaran KS (1989) Three-dimensional display of gated cardiac blood-pool studies (see comments). J Nucl Med 30:2036–2041
50. Wallis JW, Miller TR (1991) Display of cold lesions in volume rendering of SPECT studies. J Nucl Med 32:985
51. Drebin RA, Carpenter L, Hanrahan P (1988) Volume rendering. In: Proceedings of the 15th annual conference on Computer graphics and interactive techniques. ACM
52. Hoehne KH, Delapaz RL, Bernstein R, Taylor RC (1987) Combined surface display and reformatting for the three-dimensional analysis of tomographic data. Invest Radiol 22:658–664
53. Viola P, Wells WM (1995) Alignment by maximization of mutual information. In: Proceedings of the Fifth International Conference on Computer Vision (ICCV 95), June 20–23, Massachusetts Institute of Technology, Cambridge, MA. IEEE Computer Society, Washington, DC, pp 16–23
54. Collignon A, Maes F, Delaere D, Vandermeulen D, Suetens P, Marchal G (1995) Automated multi-modality image registration based on information theory. Inform Process Med Imaging 3:263–274
55. Bajcsy R, Kovacic S (1989) Multiresolution Elastic Matching. Comput Vision Graph Image Processing 46:1–21
56. Haber E, Modersitzki J (2004) Numerical methods for volume preserving image registration. Inverse Prob 20: 1621–1638
57. Bookstein FL (1989) Principal warps – thin-plate splines and the decomposition of deformations. IEEE Trans Pattern Anal Mach Intell 11:567–585
58. Kuhle WG, Porenta G, Huang SC, Buxton D, Gambhir SS, Hansen H, Phelps ME, Schelbert HR (1992) Quantification of regional myocardial blood flow using 13N-ammonia and reoriented dynamic positron emission tomographic imaging. Circulation 86:1004–1017

59. Hutchins GD, Schwaiger M, Rosenspire KC, Krivokapich J, Schelbert H, Kuhl DE (1990) Noninvasive quantification of regional blood flow in the human heart using N-13 ammonia and dynamic positron emission tomographic imaging. J Am Coll Cardiol 15:1032–1042
60. Kaufmann PA, Camici PG (2005) Myocardial blood flow measurement by PET: technical aspects and clinical applications. (erratum appears in J Nucl Med. 2005 46(2):291). J Nucl Med 46:75–88
61. Van Train KF, Areeda JS, Garcia EV, Cooke CD, Maddahi J, Kiat H, Germano G, Silagan G, Folks RD, Berman DS (1993) Quantitative same-day rest-stress technetium-99m-sestamibi SPECT: definition and validation of stress normal limits and criteria for abnormality. J Nucl Med 34:1494–1502
62. Van Train KF, Garcia EV, Maddahi J, Areeda JS, Cooke CD, Kiat H, Silagan G, Folks RD, Friedman J, Matzer L, Germano G, Bateman T, Ziffer JA, DePuey EG, Fink-Bennett D, Cloninger K, Berman DS (1994) Multicenter trial validation for quantitative analysis of same-day rest- stress technetium-99m-sestamibi myocardial tomograms. J Nucl Med 35:609–618
63. Santana CA, Folks RD, Garcia EV, Verdes L, Sanyal R, Hainer J, Di Carli MF, Esteves FP (2007) Quantitative (82) Rb PET/CT: development and validation of myocardial perfusion database. J Nucl Med 48:1122–1128