

PET: Physics, Instrumentation, and Scanners

Simon R. Cherry and Magnus Dahlbom

Positron emission tomography (PET) is a nuclear imaging technique that uses the unique decay characteristics of radionuclides that decay by positron emission. These radionuclides are produced in a cyclotron and are then used to label compounds of biological interest. The labeled compound (typically 10^{13} – 10^{15} labeled molecules) is introduced into the body, usually by intravenous injection, and is distributed in tissues in a manner determined by its biochemical properties. When the radioactive atom on a particular molecule decays, a positron is ejected from the nucleus, ultimately leading to the emission of high-energy photons that have a good probability of escaping from the body. A PET scanner consists of a set of detectors that surround the object to be imaged and are designed to convert these high-energy photons into an electrical signal that can be fed to subsequent electronics. In a typical PET scan, 10^6 to 10^9 events (decays) will be detected. These events are corrected for a number of factors and then reconstructed into a tomographic image using mathematical algorithms. The output of the reconstruction process is a three-dimensional (3-D) image volume, where the signal intensity in any particular image voxel* is proportional to the amount of the radionuclide (and, hence, the amount of the labeled molecule to which it is attached) in that voxel. Thus, PET images allow the spatial distribution of radiolabeled tracers to be mapped quantitatively in a living human. By taking a time sequence of images, the tissue concentration of the radiolabeled molecules as a function of time is measured, and with appropriate mathematical modeling, the rate of specific biological processes can be determined.

This book is designed to give the reader a solid understanding of the physics and instrumentation aspects of PET, including how PET data are collected and formed into an image. We begin with a review of the basic physics underlying PET and discuss in detail the detector technology used in modern PET scanners. The manner in which PET data are acquired is described, and the many correction factors that must be applied to ensure that the data are quantitative are

*A voxel is a volume element in a three-dimensional image array. It is analogous to a pixel in a two-dimensional image array.

introduced. The methods by which PET data are reconstructed into a three-dimensional image volume are explained, along with some of the approaches that are used to analyze and quantify the resultant images. Finally, a variety of modern PET imaging systems are discussed, including those designed for clinical service and research and small-animal imaging, along with methods for evaluating the performance of these systems.

PHYSICS OF POSITRON EMISSION AND ANNIHILATION

Basic nuclear physics and positron emission

The nucleus of an atom is composed of two different types of *nucleons*, known as *protons* and *neutrons*. These particles have similar masses but differ in that a proton has positive charge, whereas a neutron is uncharged. A cloud of negatively charged *electrons* surrounds the nucleus. In an uncharged atom, the number of electrons equals the number of protons. The basic properties of protons, electrons, and neutrons are listed in Table 1. The number of protons in an atom is known as the *atomic number* (often denoted as Z) and defines the element to which the atom belongs. The total number of nucleons is known as the *mass number*, often denoted by A . Atoms with the same Z , but different values of A , are *isotopes* of the element corresponding to atomic number Z . Nuclei usually are defined by the following notation:

$${}^A_ZX \text{ or } {}^AX \quad (1)$$

where X is the one- or two-letter symbol for the element with atomic number Z (e.g., Fe for iron and C for carbon), and A is the mass number. For example, ${}^{18}\text{F}$ is an isotope of fluorine and consists of 9 protons (because it is fluorine) and 9 neutrons. Sometimes, this isotope will also be written as F-18 or fluorine-18.

EXAMPLE 1

How many neutrons and protons are in the nucleus of ${}^{13}\text{N}$?

ANSWER

Consulting a periodic table of the elements reveals that nitrogen has an atomic number of 7 and therefore, 7 protons. The mass number of this isotope is 13, so the number of neutrons must be $(13 - 7) = 6$.

EXAMPLE 2

How would an atom with 29 protons and 35 neutrons be written in the notation of Equation 1.

ANSWER

Referring to a periodic table of the elements shows that the element corresponding to $Z = 29$ is copper (symbol Cu). The total number of nucleons is $(29 + 35) = 64$. Therefore, this nucleus is ${}^{64}\text{Cu}$.

The nucleus is held together by two opposing forces. The strong force is an attractive force between nucleons and is balanced by the repulsive coulomb (elec-

TABLE 1. Mass and Charge Properties of Nucleons, Electrons, and Positrons

	Proton (p)	Neutron (n)	Electron (e^-)	Positron (e^+)
Mass	1.67×10^{-27} kg	1.67×10^{-27} kg	9.1×10^{-31} kg	9.1×10^{-31} kg
Charge	$+1.6 \times 10^{-19}$ C	0	-1.6×10^{-19} C	$+1.6 \times 10^{-19}$ C

Based on data from Handbook of Physics and Chemistry, 71st Edition, Ed: D.R. Lide, CRC Press, Boca Raton, FL, 1991.

trical) force between the positively charged protons. If a nucleus has either an excess number of protons or neutrons, it is unstable and prone to radioactive decay, leading to a change in the number of protons or neutrons in the nucleus and a more stable configuration. Nuclei that decay in this manner are known as *radionuclides*. For a specific element with atomic number Z , isotopes that are unstable and which undergo radioactive decay are known as *radioisotopes* of that element.

One common method by which nuclei with an excess of protons may decay is through *positron emission* (also known as β^+ or *beta-plus decay*). Essentially, a proton in the nucleus of the atom is converted into a neutron (n) and a positron (e^+). The positron is the antiparticle to the electron with the same mass but opposite electric charge (see Table 1). The positron is ejected from the nucleus, along with a neutrino (ν) that is not detected. An example of a radionuclide that decays by positron emission is ^{11}C :



The net energy released during positron emission is shared between the daughter nucleus, the positron, and the neutrino. Positrons are therefore emitted with a range of energies, from zero up to a maximum *endpoint energy* E_{max} . This endpoint energy is determined by the difference in atomic masses between the parent atom and the daughter atom, taking into account gamma-ray emission from excited states that may occur if the transition is not between the ground states of the two nuclei. The mean kinetic energy of the emitted positrons is approximately $0.33 \times E_{\text{max}}$. Decay by positron emission is the basis for PET imaging.

Proton-rich radionuclides also can decay by a process known as *electron capture*. Here, the nucleus captures an orbital electron and converts a proton into a neutron, thus decreasing the atomic number Z by one. Once again, a neutrino is released. An example of electron capture would be the decay of ^{125}I :



Electron capture decay can lead to emission of x-rays (filling of the orbital vacancy created by the captured electron) or gamma-rays (electron capture leaves the nucleus in an excited state with further decay to the ground state by emission of one or more gamma-rays). These emissions may also be used for in vivo imaging but do not share the unique properties of decay by positron emission which are explained in the section on Annihilation (p. 5). Decay by electron capture and positron emission compete with one another, with positron emission usually being the dominant process in low Z nuclei, and electron capture being more likely in higher Z nuclei. Radionuclides that decay predominantly by positron emission are preferred for PET imaging.

TABLE 2. Select List of Radionuclides That Decay by Positron Emission and Are Relevant to PET Imaging

Radionuclide	Half-life	$E_{max}(MeV)$	β^+ Branching Fraction
^{11}C	20.4 min	0.96	1.00
^{13}N	9.97 min	1.20	1.00
^{15}O	122 s	1.73	1.00
^{18}F	109.8 min	0.63	0.97
^{22}Na	2.60 y	0.55	0.90
^{62}Cu	9.74 min	2.93	0.97
^{64}Cu	12.7 h	0.65	0.29
^{68}Ga	67.6 min	1.89	0.89
^{76}Br	16.2 h	Various	0.56
^{82}Rb	1.27 min	2.60, 3.38	0.96
^{124}I	4.17 d	1.53, 2.14	0.23

Based on data from Table of Nuclides: www2.bnl.gov/ton (accessed October 17th, 2002)

Many radionuclides decay by positron emission. Table 2 presents a selection of these radionuclides that are commonly encountered in relation to PET imaging. Included in the table are the maximum kinetic energy of the emitted positrons, E_{max} , and the fraction of decays that occur by positron emission. The energy of the emissions from radioactive decay are normally given in units of electron volts (eV), which is a more convenient unit than standard Système International (SI) energy units for handling the relatively small energies involved. One electron volt is defined as being equal to the energy acquired by an electron when it is accelerated through a potential difference of one volt. The conversion to joules, the SI unit for energy is:

$$1 \text{ eV} = 1.6 \times 10^{-19} \text{ J} \quad (4)$$

For PET imaging, units of kiloelectron volts ($1 \text{ keV} = 10^3 \text{ eV}$) and megaelectron volts ($1 \text{ MeV} = 10^6 \text{ eV}$) are commonly used.

Table 2 also lists the half-life of the radionuclides. A sample of identical radioactive atoms will decay in an exponential fashion, and the half-life is the time required for half the atoms in the sample to decay. The relationship between the activity A of a sample at time t , and the half-life, $T_{1/2}$, is given by:

$$A(t) = A(0)\exp(-\ln 2 \times t/T_{1/2}) \quad (5)$$

where $A(0)$ is the activity of the sample at time 0. Activity is measured in units of the number of disintegrations per second:

$$1 \text{ bequerel (Bq)} = 1 \text{ disintegration per second} \quad (6)$$

In the United States, traditional units of the curie (Ci) and millicurie ($1 \text{ mCi} = 10^{-3} \text{ Ci}$) are still frequently used. The conversion is:

$$1 \text{ mCi} = 37 \times 10^6 \text{ Bq} = 37 \text{ MBq} \quad (7a)$$

or

$$1 \text{ MBq} = 27 \times 10^{-6} \text{ Ci} = 27 \mu\text{Ci} \quad (7b)$$

For more information on the physics of positron emission, see the textbook by Evans.¹

EXAMPLE 3

A sample of ^{18}F is measured at 10:40 AM and has an activity of 30 MBq. It is injected into a patient at 11:30 AM. How much activity was injected?

ANSWER

From Table 2, the half-life of ^{18}F is 109.8 minutes. The time elapsed between measurement of the sample and injection is 50 minutes. Using Equation 5, the activity at the time of injection is:

$$A(t) = 30 \text{ MBq} \times \exp(-0.693 \times 50/109.8) = 21.9 \text{ MBq}$$

Annihilation

The positron that is ejected following β^+ decay has a very short lifetime in electron-rich material such as tissue. It rapidly loses its kinetic energy in inelastic interactions with atomic electrons in the tissue, and once most of its energy is dissipated (typically within 10^{-1} to 10^{-2} cm, depending on its energy), it will combine with an electron and form a hydrogen-like state known as *positronium*. In the analogy to hydrogen, the proton that forms the nucleus in a hydrogen atom is substituted by a positron. This state lasts only about 10^{-10} seconds before a process known as *annihilation* occurs, where the mass of the electron and the positron is converted into electromagnetic energy. Because the positron and electron are almost at rest when this occurs, the energy released comes largely from the mass of the particles and can be computed from Einstein's mass-energy equivalence as:

$$E = mc^2 = m_e c^2 + m_p c^2 \quad (8)$$

where m_e is the mass of the electron, m_p is the mass of the positron, and c is the speed of light (3×10^8 m/s). Inserting the values from Table 1, and using Equation 8 and the conversion in Equation 4, the energy released can be shown to be 1.022 MeV.

The energy is released in the form of high-energy photons. As the positron and electron are almost at rest when the annihilation occurs, the net momentum is close to zero. Because momentum as well as energy must be conserved, it is not in general possible for annihilation to result in the emission of a single photon; otherwise, a net momentum would occur in the direction of that photon. Instead, two photons are emitted simultaneously in opposite directions (180° apart), carrying an energy equal to $1.022 \text{ MeV}/2$, or 511 keV, ensuring that both energy and momentum are conserved. This process is shown schematically in Figure 1. Higher order annihilation, in which more than 2 photons are emitted, is also possible, but only occurs in about 0.003% of the annihilations.

The annihilation process has a number of very important properties that are advantageous for imaging and lead directly to the concept of PET. First, the annihilation photons are very energetic (they fall in the gamma-ray region of the electromagnetic spectrum and are roughly a factor of ten higher in energy than diagnostic x-rays), which means they have a good chance of escaping the body for external detection. It is, therefore, the annihilation photons that are detected in PET imaging, not the positrons (which are absorbed locally). Second, two photons are emitted with a precise geometric relationship. If both photons can





FIGURE 1. The process of positron emission and subsequent positron-electron annihilation results in two 511 keV annihilation photons emitted 180° apart. The site of annihilation is usually very close to the point of positron emission because the emitted positrons rapidly lose their energy in tissue (see Figure 5).

be detected and localized externally, the line joining the detected locations passes directly through the point of annihilation (Figure 2A). This was originally referred to as *electronic collimation*.² Because the point of annihilation is very close to the point of positron emission, this also gives a good indication (again to within a line) of where the radioactive atom was in the body. Contrast this with radioactive decay schemes that result in emission of a single photon. Although a single detector can be used, the detection and localization of a single photon tells nothing about where it came from in the body (Figure 2B). The direction of the photon can only be determined by the using a form of absorptive collimation, which only allows photons emitted in a certain direction to impinge on the detector (Figure 2C). This reduces the number of events that are detected for a given amount of radioactivity in the body by at least 1 to 2 orders of magnitude compared with electronic collimation. Electronic collimation also allows events to be collected from many different directions simultaneously leading to the capability of rapid tomographic imaging (see Image Reconstruction, p. 70). Third, all positron-emitting radionuclides, independent of the element involved, or the energy of the emitted positrons, ultimately lead to the emission of two back-to-back 511 keV photons; that is, a PET scanner can be designed and optimized for imaging all positron-emitting radionuclides at this single energy. One drawback to this, however, is that it is not possible to perform dual-radionuclide studies with PET and distinguish between the radionuclides based on the energy of the emissions. Because the annihilation photons fall in the gamma-ray region of the electromagnetic spectrum, the terms photons and gamma-rays are often used interchangeably when referring to the annihilation photons. Annihilation photons is technically the correct term because the radiation does not

FIGURE 2. (A) Radionuclides that decay by positron emission result in two annihilation photons emitted 180° apart. If both photons are detected, the detection locations define (to within the distance traveled by the positron prior to annihilation) a line along which the decaying atom was located. (B) Radionuclides that decay by emitting single photons provide no positional information, as a detected event could originate from anywhere in the sample volume. (C) For single photon imaging, physical collimation can be used to absorb all photons except those that are incident on the detector from one particular direction (in this case perpendicular to the detector face), defining a line of origin just like the coincident 511-keV photons do following positron emission. To achieve this localization, however, the radiation from the majority of decays has been absorbed and does not contribute to image formation, leading to the detection of many fewer events for a given amount of radioactivity in the object. Absorptive collimation of this kind is the approach used in planar nuclear medicine imaging and in single photon emission computed tomography (SPECT).



arise directly from the nucleus. However, the properties of annihilation photons are absolutely identical to a 511-keV gamma-ray—the difference in terminology reflects their different origins.

The annihilation process forms the basis for PET imaging. A PET scanner is designed to detect and localize the simultaneous back-to-back annihilation photons that are emitted following decay of a radionuclide by positron emission (Figure 3). In a typical PET scan, many millions of these photon pairs will be detected from a compound that is tagged with a positron-emitting radionuclide and which has been injected into the body.

As described above, the detection of the annihilation photons only localizes the location of the radioactive atom to within a line joining the detecting positions. Two approaches can then be used to form an image that reflects the actual locations of the radioactive atoms and, therefore, the compound to which it is attached. The first approach is conceptually the most simple, but is rarely used. It involves measuring the difference in arrival time of the two photons at the detectors. Obviously, if an annihilation occurs closer to detector 1 than detector 2, then the annihilation photon directed towards detector 1 will arrive at that detector earlier than the annihilation photon directed towards detector 2. The relationship between the difference in arrival time of the two annihilation photons, Δt , and the location d of the annihilation with respect to a point exactly half-way between the two detectors, is given by:

$$d = \frac{\Delta t \times c}{2} \tag{9}$$

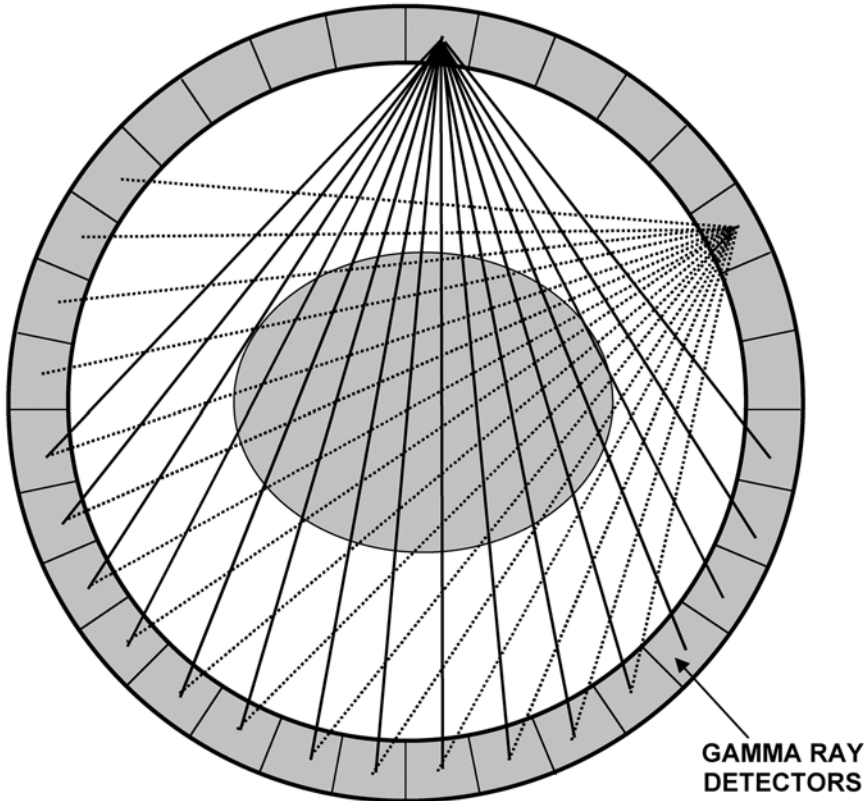


FIGURE 3. Schematic drawing of a PET scanner consisting of a ring of high-energy photon (gamma-ray) detectors. A ring geometry is shown, but other possibilities include polygonal assemblies of panels and opposing rotating planar detectors. The detectors are designed to record as many of the annihilation photons as possible and to locate the line along which the decay occurred by determining the two interaction vertices. Each detector is in electronic coincidence with a fan of detectors on the opposite side of the ring, so the object is simultaneously sampled from many different angles. For clarity, the measured lines of response for just two detectors are shown in this figure. Typically, 10^6 to 10^9 events (detections of annihilation photon pairs) are needed in a PET scan to reconstruct a statistically meaningful image of the distribution of radioactivity in the body.

where c is the speed of light (30 cm/ns). In practice, this method, known as *time of flight*, is very difficult and costly to implement because of the very small time differences involved. Even a timing resolution as fine as 100 ps would only yield a positional resolution of ~ 1.5 cm. With currently available detector technology, the best timing resolution that can be achieved is on the order of a few hundred picoseconds. Therefore, time-of-flight approaches do not yield the desired accuracy of a few millimeters, and no PET scanners are currently manufactured using this technique. The approach that is used almost universally involves the concept of *computed tomography*. By measuring the total radioactivity along lines that pass at many different angles through the object, mathematical algorithms can be used to compute cross-sectional images that reflect the concentration of the positron-emitting radionuclide in tissues throughout the body. This is discussed in Image Reconstruction (p. 70).

Positron range and nonlinearity

There are two effects in PET imaging systems that lead to errors in determining the line along which a positron-emitting radionuclide is to be found. These effects place some finite limits on the spatial resolution attainable with PET and manifest themselves as a blurring of the reconstructed images.

The first of these effects is positron range. As shown in Figure 4 (top), this is the distance from the site of positron emission to the site of annihilation. A PET scanner detects the annihilation photons which define the line along which the annihilation takes place, not the line along which the decaying atom is located. Because the positrons follow a tortuous path in tissue, undergoing multiple direction-changing interactions with electrons prior to annihilation, the total path length the positron travels is considerably longer than the positron range. From the perspective of PET imaging, it is the perpendicular distance from the emission site to the line defined by the annihilation photons that matters and which causes mispositioning.

As described earlier, radionuclides differ in the energy of emitted positrons. Some radionuclides emit, on average, higher energy positrons than others, making the positron range effect radionuclide-dependent. Figure 5 shows the annihilation locations for positron emission from a point source emitter located at the center of a block of tissue-equivalent material. Notice the broader distribution for oxygen-15 (a high energy positron emitter with $E_{\max} = 1.72$ MeV) compared to fluorine-18 ($E_{\max} = 0.64$ MeV). Profiles through these distributions reveal that they are nonGaussian in nature and are best fitted by exponential functions. Several groups have either measured,³ computed,⁴ or simulated⁵ these distributions. Although the trends are similar, some disagreement between these studies is noted on the exact width and shape of the distribution. The blurring effect on the final PET image, however, clearly ranges from a few tenths of a millimeter up to several millimeters, depending on the radionuclide and its E_{\max} .

Positron range limits the ultimate resolution attainable by PET. Studies have shown the ability to reduce positron range, particularly in radionuclides with large E_{\max} , by using strong magnetic fields.⁶⁻⁸ However, this is not currently practical to implement in the complex setting of a PET system. The positron range distribution may also in theory be deconvolved from the PET image.^{9,10} In practice, the data rarely, if ever, have the statistical quality (sufficient number of events) to make this advantageous, as deconvolution leads to noise amplification. A better approach may be to incorporate positron range distribution information into iterative reconstruction algorithms (Iterative Reconstruction Methods, p. 86), which should lead to improvements in image resolution that are consistent with the statistical quality of the data when using positron-emitters with a high E_{\max} . To put this discussion in perspective, it should also be pointed out that positron range is not a major limiting factor in PET imaging at the present time, except perhaps in animal studies of the very highest resolution using positron emitters with relatively high values (> 1.5 MeV) of E_{\max} .

The second effect comes from the fact that the positron and electron are not completely at rest when they annihilate. The small net momentum of these particles means that the annihilation photons will not be at exactly 180° and will, in fact, be emitted with a distribution of angles around 180° . This is known as


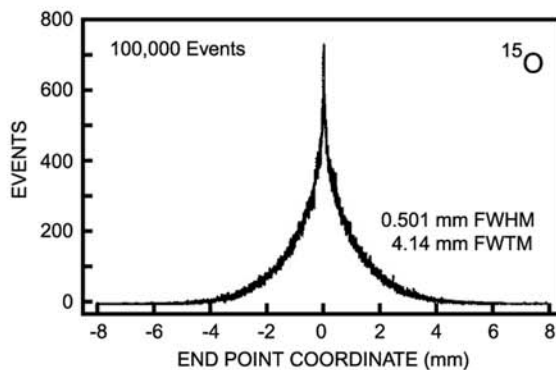
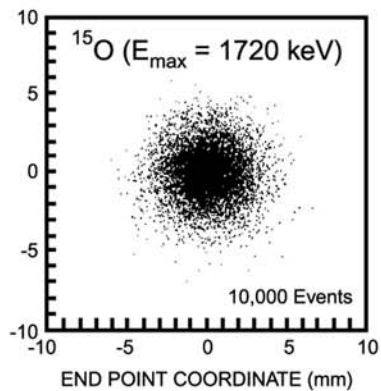








FIGURE 4. Error in determining the location of the emitting nucleus due to positron range (top) and noncolinearity (bottom). The positron range error is dependent on the energy of the emitted positrons. Noncolinearity is independent of radionuclide, and the error is determined by the separation of the detectors. The deviation from noncolinearity is highly exaggerated in the figure; the average angular deviation from 180° is about $\pm 0.25^\circ$. (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)

FIGURE 5. A: Simulations for several PET radionuclides showing the distribution of positron annihilation sites in water for positrons emitted at the center of the image (position 0.0 mm). B: Profiles through the simulated distributions showing measured FWHM and FWTM of the distributions. Abbreviations: FWHM, full width at half maximum; FWTM, full width at tenth maximum. (Reproduced with permission from Levin C, Hoffman EJ. *Phys Med Biol* 1999, 44: 781–799.)



A.

B.

noncolinearity. This effect is independent of radionuclide because the positrons must lose most of their energy before they can annihilate; hence, the initial energy is irrelevant. The distribution of emitted angles is roughly Gaussian in shape, with a full width at half maximum (FWHM)* of $\sim 0.5^\circ$. After detecting the annihilation photons, PET assumes that the emission was exactly back to back, resulting in a small error in locating the line of annihilation (Figure 4 bottom). Assuming a Gaussian distribution and using the fact that the angles are small, the blurring effect due to noncolinearity, Δ_{nc} , can be estimated as:

$$\Delta_{nc} = 0.0022 \times D \quad (10)$$

where D is the diameter of the PET scanner. The error increases linearly as the diameter of the PET scanner increases. Once again, the effect is relatively small compared with the detector resolution in most clinical PET scanners. In PET scanners used for animals, D generally is small, and as illustrated in Example 4, noncolinearity is not a major limiting factor at the present time.

EXAMPLE 4

Calculate the blurring due to photon noncolinearity in an 80-cm diameter PET scanner designed for imaging humans and in a 15-cm diameter PET scanner designed for imaging small animals.

ANSWER

From Equation 10, the blurring is calculated as:

80-cm human scanner:

$$\Delta_{nc} = 0.0022 \times D = 0.0022 \times 800 \text{ mm} = 1.76 \text{ mm}$$

15-cm small-animal scanner:

$$\Delta_{nc} = 0.0022 \times D = 0.0022 \times 150 \text{ mm} = 0.33 \text{ mm}$$

511 keV PHOTON INTERACTIONS IN MATTER


It is important to understand how the 511-keV photons emitted following annihilation interact with the tissue surrounding them, with the detector material of the PET scanner, and with materials such as lead and tungsten that may be used for shielding or slice collimation purposes. The *photoelectric effect* and *Compton scattering* are two major mechanisms by which 511-keV photons interact with matter.

Photoelectric interactions

Figure 6 summarizes interaction by the photoelectric effect. A 511-keV photon will interact with an atom as a whole in the surrounding medium and is completely absorbed by transferring its energy to an orbital electron. This electron is given enough energy to escape the atom but is quickly absorbed in solids and

*FWHM is often used to characterize a distribution that is Gaussian or nearly Gaussian and involves measuring the width of the distribution at the point where it reaches half the maximum amplitude. A related measure, full width at tenth maximum (FWTM), identifies the width of the distribution where it reaches one tenth of its maximum amplitude.

FIGURE 6. Schematic representation of the photoelectric effect. The incident photon transfers all of its energy to an electron which is ejected from the atom but which is itself absorbed by material nearby. (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)



liquids. An x-ray with an energy equal to the binding energy of the electron is also generated as the vacancy in the electron shell of the atom is filled. These x-rays typically have energies of tens of keV and are also quickly absorbed in the medium. The net result of a photoelectric interaction in a reasonably dense liquid or solid is the complete absorption of the original photon with all 511 keV of energy deposited locally (within a sphere of a few hundred microns in diameter) in the material. The probability of photoelectric absorption per unit distance in a medium strongly depends on the atomic number of the medium in which the photon is propagating. At 511 keV, it is roughly proportional to $Z^{3,4}$.

Compton scattering interactions

Compton scattering interactions are shown in Figure 7. Here, the 511-keV photon scatters off a free or loosely bound electron in the medium, transferring some of its energy to the electron and changing direction in the process. Imposing conservation of momentum and energy leads to a simple relationship¹ between the energy of the original photon (E), the energy of the scattered photon (E_{sc}) and the angle through which it is scattered, θ :

$$E_{sc} = \frac{m_e c^2}{\frac{m_e c^2}{E} + 1 - \cos\theta} \quad (11)$$

In this equation, m_e is the mass of the electron and c is the speed of light (2.998×10^8 m/s). Using units of electron volts for energy, the term $m_e c^2$ is equal to 511




FIGURE 7. Schematic representation of Compton scattering in which the incident photon transfers part of its energy to an electron, causing it to change direction. The scattered photon carries considerable energy and can have a long range in materials such as tissue. (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)

keV. In PET, the incoming photon of interest has an energy level of 511 keV; the equation, therefore, reduces further to:

$$E_{sc}(\text{keV}) = \frac{511}{2 - \cos \theta} \quad (12)$$

The recoil energy that is transferred to the electron, E_{re} , which is dissipated in the medium, is equal to $E - E_{sc}$:

$$E_{re} = E - E_{sc} = E \times \frac{(1 - \cos \theta)}{\left(\frac{m_e c^2}{E} + 1 - \cos \theta\right)} \quad (13)$$

Again substituting $E = 511$ keV, this reduces to:

$$E_{re}(\text{keV}) = 511 \times \frac{1 - \cos \theta}{2 - \cos \theta} \quad (14)$$

The maximum energy that can be imparted to the electron (and, therefore, the medium) occurs when the photon is scattered through 180° . The probability of Compton scattering per unit length of absorbing medium is linearly proportional to the atomic number of the medium.

EXAMPLE 5

Calculate the minimum energy of a 511-keV photon after it has undergone Compton scattering. What is the energy given up to the recoil electron?

ANSWER

The minimum energy will occur when the maximum energy is given to the electron. This occurs for a scattering angle of 180° . From Equation 12, the energy of the scattered photon will be:

$$E_{sc} = 511 \text{ keV} / (2 - \cos 180^\circ) = 170 \text{ keV}$$

The energy given to the electron is simply $E - E_{sc} = 511 \text{ keV} - 170 \text{ keV} = 340 \text{ keV}$.

The angular distribution of the scattered photons is given by the Klein-Nischnina equation.¹ It is independent of the scattering medium but strongly dependent on the energy of the photons. The angular distribution for a range of energies is shown in Figure 8.

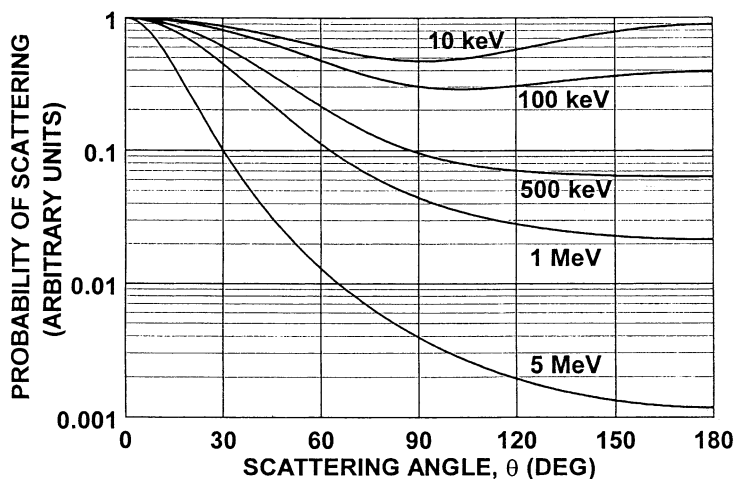
Interaction cross-sections in various materials

The interaction (absorption or scattering) of 511-keV photons by matter can be described with a simple exponential relationship:

$$I(x) = I(0)\exp(-\mu x) \quad (15)$$

where $I(0)$ is the 511-keV photon flux impinging on the medium, x is the thickness of the medium, and $I(x)$ is the flux of 511-keV photons that passes through the medium without interaction. The parameter μ is the linear attenuation coefficient and is the probability per unit distance that an interaction will occur.

FIGURE 8. Relative probability of Compton scattering (normalized per unit of solid angle) versus scattering angle. At 511 keV, small-angle forward scatter is most likely. (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)



For 511-keV photons, it is largely made up of components due to photoelectric absorption and Compton scattering, such that:

$$\mu \approx \mu_{\text{compton}} + \mu_{\text{photoelectric}} \quad (16)$$

For PET imaging, three media are of potential interest: tissue in the body, the detector material, and any material used for shielding or collimation. The attenuation coefficients for soft tissue, bone, for a typical detector material (bismuth germanate or BGO) and for lead and tungsten are shown in Table 3. The *half-value thickness*, the thickness of material that is required to cause half of the 511-keV photons to interact, is also given.

The annihilation photons must pass through the body so that they can be detected. The dominant form of interaction for 511-keV photons in tissue (Table 3) is Compton scattering. Therefore, photon interactions in the body attenuate the signal by redirecting annihilation photons that would have struck a particular detector pair. The angular correlation between the annihilation photons is randomized by the scattering process, so if the redirected photons still escape the body and are detected in the PET scanner, they will be incorrectly located. This results in a background of scattered events in the images. Unfortunately, even at the high energies of the annihilation photons, substantial numbers are scattered in the body because the Compton scatter cross-section is quite high. Example 6 illustrates this point. A number of steps are taken to correct for the

TABLE 3. Linear Attenuation Coefficients for Soft Tissue, Bone, Bismuth Germanate (a Detector Material), Lead, and Tungsten at 511 keV.

Material	μ_{Compton} (cm^{-1})	$\mu_{\text{photoelectric}}$ (cm^{-1})	μ (cm^{-1})	Half-value thickness (cm)
Soft tissue	~0.096	~0.00002	~0.096	7.2
Bone	~0.169	~0.001	~0.17	4.1
Bismuth germanate (B60)	0.51	0.40	0.96	0.76
Lead	0.76	0.89	1.78	0.42
Tungsten	1.31	1.09	2.59	0.29

attenuation due to scatter and to minimize and remove the scatter background as explained in Attenuation Correction (p. 56) and Scatter Correction (p. 63). This is, in part, based on the fact that photons that undergo Compton interaction in the body will lose energy (Equation 12) and can thus be rejected by a detector if they have sufficient energy discrimination.

EXAMPLE 6

Determine the probability that a 511-keV photon emitted 7.5 cm deep inside the brain is Compton scattered. How does this change if the photon is emitted from a point in the liver, 20 cm from the surface of the body?

ANSWER

From Table 3, the attenuation coefficient due to Compton scattering for 511-keV photons in tissue is 0.096 cm^{-1} . Using Equation 15 we find:

$$I(x)/I(0) = \exp(-\mu x)$$

Brain: $x = 7.5 \text{ cm}$, $I(x)/I(0) = 0.49$ (49% of photons escape unscattered)

Liver: $x = 20 \text{ cm}$, $I(x)/I(0) = 0.15$ (15% of photons escape unscattered)

This implies that relatively large numbers of photons are scattered. When using PET, we require that both photons are unscattered when they reach the detectors. The probability that neither photon scatters is roughly the square of the probabilities above, assuming both photons pass through equal amounts of tissue.

The function of the PET scanner is to detect those 511-keV photons that escape the body without interacting. The detector material should, therefore, be something that has a high probability of stopping these photons, that is, very dense materials, with large values of μ . An example of such a material is BGO. It is also preferable that the detector has as high a ratio between photoelectric and Compton interactions as possible. Photoelectric interactions are preferred in a detector because they result in all of the energy being deposited locally. Compton scattering can result in multiple interactions within a detector or interactions in adjacent detectors. It can, therefore, be difficult to unambiguously define the location of the interaction.

EXAMPLE 7

What thickness of bismuth germanate detector material would be required to cause 90% of the incoming photon flux to interact?

ANSWER

From Table 3, the attenuation coefficient for 511-keV photons in bismuth germanate is 0.96 cm^{-1} . Using Equation 15 and setting $I(x)/I(0) = 0.1$ (only 10% transmitted, 90% interact) we find:

$$I(x)/I(0) = 0.1 = \exp(-0.96 \times x)$$

$$x = \ln(0.1) / -0.96 = 2.4 \text{ cm}$$

This demonstrates that detectors will need to be several centimeters thick to be highly efficient at stopping 511-keV photons.

Finally, we may want to shield the detectors from radioactivity that is outside the imaging volume. In addition, some multislice PET scanners use collimators in the axial direction to define the slices of the object that is being imaged. These axial collimators can also be an effective way to decrease the detection of photons that scatter in the body. The idea of collimator or shielding material is to absorb any photons that are incident on them, and so materials with the very highest attenuation values (subject to requirements for ease of machining and cost) are used. Lead and tungsten are two commonly used materials in this regard.

511 keV PHOTON DETECTORS

A PET scanner is comprised of a set of two or more detectors. To obtain the best quality image for a given injected dose of radioactivity, the detectors must have a very high *efficiency* for detecting 511-keV photons that impinge on their surface (the more photon pairs that are detected, the better the signal-to-noise in the image) and must also give precise information on the spatial location of the interaction (this relates directly to the *spatial resolution* of the images). The latter is generally achieved in one of two ways, either by using arrays of small detector elements, in which case the precision of localization is related to the size of the detector elements, or by using a larger area detector that has position-sensing capability built in. It is also important to be able to determine when a photon struck the detectors, so that the time of all detected events can be compared to determine which ones arrived closely enough in time to correspond to an annihilation pair. The ability of a pair of detectors to determine the time difference in arrival of the annihilation photons is known as the *timing resolution* and is typically on the order of 2 to 6 ns. A typical timing window that is used in PET scanners so as not to accidentally reject annihilation photon pairs is typically 2 to 3 times the timing resolution, leading to values in the range of 4 to 18 ns. Finally, the detectors should indicate the energy of the incoming annihilation photon such that those that have scattered in the body (and have thus lost energy as explained in the section on Compton Scattering Interactions (p. 13) can be rejected. The ability of the detector to determine the energy of the photon is known as the *energy resolution*.

Scintillation detectors are widely used gamma-ray detectors that form the basis for almost all PET scanners in use today. These detectors consist of a dense crystalline *scintillator* material that serves as an interacting medium for gamma-rays and high-energy photons and which emits visible light when energy is deposited inside of them. This light is then subsequently detected by some form of visible light photon detector and converted into an electrical current (Figure 9). This section reviews the components of scintillation detectors and shows how they are used as detectors in PET scanners. Other technologies for gamma-ray detection (and therefore annihilation photon detection) are also briefly reviewed.

Scintillators

Scintillators are transparent materials that have the property of emitting light in the visible region of the spectrum when energy from particles or high-energy

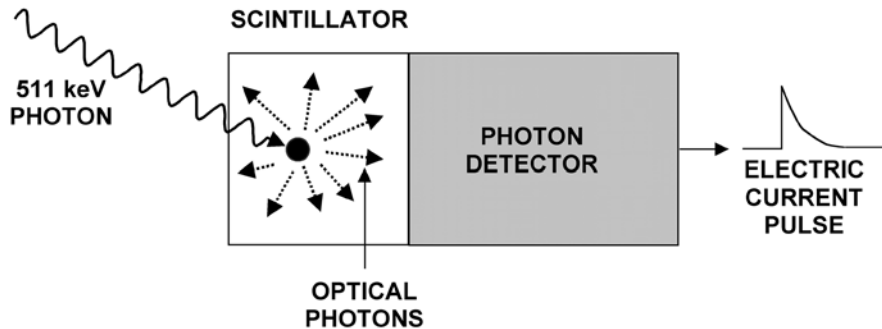


FIGURE 9. Basic components of a scintillation detector. The incident annihilation photon interacts in the scintillator (either by photoelectric or Compton interaction). Each annihilation photon that interacts produces a single pulse in the detector, with the amplitude of the pulse being determined by the number of scintillation photons reaching the photon detector and any amplification inherent in the photon detector.

photons are deposited in them. The light is emitted isotropically and the amount of light emitted is proportional to the amount of energy that is deposited in the material. These materials can be organic or inorganic compounds and can come in both solid and liquid forms. An excellent review of scintillators can be found in the textbook by Knoll.¹¹ Scintillators are characterized by their stopping power, their brightness, the wavelength of the emitted light, and the time over which the light is produced—important considerations in choosing a scintillator for a particular application. For the purposes of PET imaging, the scintillator must be a dense material that can stop a large fraction of the incident 511-keV photons. For this reason, dense, inorganic, solid scintillators are the scintillators of choice. Table 4 lists some of the properties of scintillator materials suitable for detecting gamma-rays in the 100 to 1000 keV range. As demonstrated by Example 7, a thickness of several centimeters of these scintillators is required to effectively stop a large percentage of incident 511-keV photons. It is also apparent, that even in the most dense available scintillators, Compton interactions are more likely to occur than photoelectric interactions at 511 keV.

While stopping power is a major factor in the choice of a scintillator for PET, other considerations are also important. The brightness of the scintillator (the number of light photons produced per 511 keV interaction) is important because the integrated light signal from the scintillator (converted by the subsequent photon detector from photons into electrons) is used in several different ways. In many detectors, the relative amplitudes of the signals seen by adjacent light sensors viewing a piece of scintillator are used to determine the location of the interaction. The integrated light signal is used as a direct measure of the energy deposited in the scintillator; therefore, by placing a lower threshold on the output, it is possible to reject low-energy photons that have scattered in the body. In both cases, a major source of noise in the measurement (leading to errors in positioning or energy) are statistical fluctuations in the number of scintillation photons detected. These fluctuations are governed by Poisson counting statistics and reduce as $1/\sqrt{N}$ where N is the number of scintillation photons that are detected.

TABLE 4. Properties of Scintillator Materials Useful for Gamma-Ray Detection at 511 keV

<i>Scintillator</i>	<i>Density (g/cc)</i>	<i>Light output (photons per 511 keV)</i>	<i>Decay time (ns)</i>	<i>Index of refraction</i>	<i>Linear attenuation at 511 keV (cm⁻¹)</i>	<i>Ratio between photoelectric and Compton</i>
Sodium iodide [NaI(Tl)]	3.67	19400	230	1.85	0.34	0.22
Bismuth Germanate (BGO)	7.13	4200	300	2.15	0.96	0.78
Lutetium Oxyorthosilicate (LSO:Ce)	7.40	~13000	~47	1.82	0.88	0.52
Gadolinium Oxyorthosilicate (GSO:Ce)	6.71	~4600	~56	1.85	0.70	0.35
Barium Fluoride (BaF2)	4.89	700, 4900	0.6, 630	1.56	0.45	0.24
Yttrium Aluminum Perovskite (YAP:Ce)	5.37	~9200	~27	1.95	0.46	0.05

Based on data (Table 8.3) from Knoll GF (Radiation Detection and Measurements, 3rd Edition, 2000, Wiley, New York, 2000), light output and decay time data for cerium doped scintillators such as LSO, GSO, and YAP are approximate and can vary by tens of percent depending on cerium concentration, impurities, and growing conditions.

Because PET imaging involves the coincident detection of the two annihilation photons, it will be important to have an accurate assessment of exactly when a photon interacts in a detector. The accuracy of timing is determined in large part both by the decay time of the scintillator and its brightness. A fast, bright scintillator will produce a signal with less timing variation than a slow, dim scintillator. This observation is based on an analysis of the spread of the average arrival times of the first scintillation photons at the photodetector. It is these first photons which trigger the start of a pulse and is the earliest time point that can be detected. Finally, the index of refraction of the scintillator is also important as this determines how efficiently optical photons can be transmitted from the scintillator to the photodetector. Large mismatches in index result in significant internal reflection at the scintillator/photodetector boundary and reduce light transmission to the photodetector.

When discussing PET detectors based on scintillators detectors, it is very important to clearly distinguish between the high-energy annihilation photons (511 keV) that are absorbed by the scintillator and the burst of low-energy optical photons (energy of a few eV) that are subsequently emitted by the scintillator and converted into an electric current. Each interaction of a 511-keV photon in the scintillator ultimately produces a single electrical pulse. The number of visible light photons generated in the scintillator and detected by the photon detector determines the amplitude of that pulse.

Photomultiplier tubes

The vast majority of commercially available PET scanners use *photomultiplier tubes* (PMTs)¹² as the photon detector to convert scintillation light into an elec-




FIGURE 10. A photomultiplier tube (PMT) consists of a series of dynodes (electrodes) each of which is held at a greater voltage with a resistor chain. Each dynode is coated with an emissive material in an evacuated glass tube. The inner surface of the entrance window (the photocathode) is also coated with an emissive material. Light photons striking the photocathode can release electrons into the tube, and these electrons are accelerated by a potential difference to the first dynode. Each electron has sufficient energy upon striking the first dynode to release further electrons, which, in turn, are accelerated to the second dynode. After 10 dynode stages, each original electron produced at the photocathode has been amplified into approximately 10^6 electrons, producing a sizeable current at the PMT output (anode). (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)

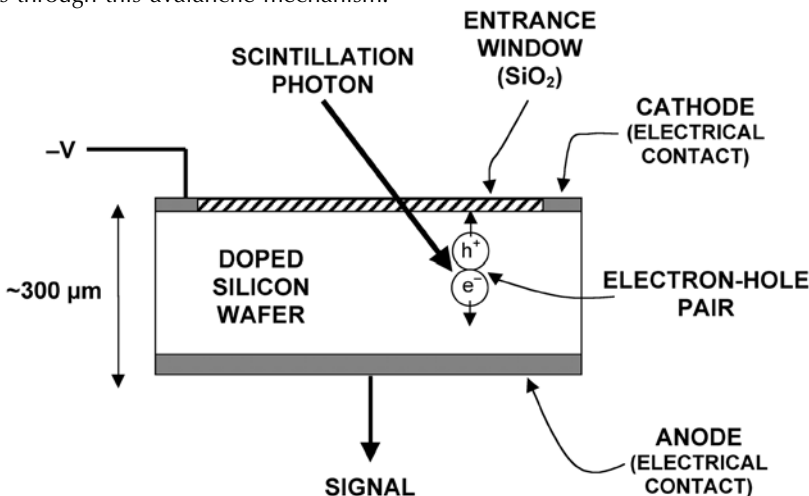
trical current. A cross-section through a typical photomultiplier tube is shown in Figure 10. Light from the scintillator is transmitted through the glass entrance window of the PMT and excites the *photocathode*. The photocathode is made from a thin layer of material that can easily liberate electrons as energy is deposited in it. Each light photon from the scintillator has roughly a 15% to 25% chance (depending on wavelength) to liberate an electron. This probability is called the *quantum efficiency* of the PMT. A high potential difference accelerates the electron from the photocathode and directs it to strike a positively charged electrode called the first *dynode*. This dynode is also coated with an emissive material that readily releases electrons, and each impinging electron has acquired sufficient energy to release on the order of 3 to 4 secondary electrons from the dynode. These electrons are in turn accelerated to the second dynode and so forth, ultimately creating an avalanche of photoelectrons. After 10 stages of amplification, each initial electron has created on the order of 10^6 electrons, which, occurring over a period of a few nanoseconds, lead to an easily detectable current in the milliamp range. PMTs come in a wide range of shapes and sizes and also are available as multichannel and position sensitive models. Most PET scanners use round or square single-channel PMTs in the range of 1 to 5 cm in diameter. The advantages of PMTs are their high *gain* (amplification), which leads to high signal-to-noise pulses, their stability and ruggedness, and their fast response (the output pulse from a PMT rises in approximately a nanosecond for

a step function input of light into the PMT). The disadvantages are that they are quite bulky and fairly expensive.

Solid state photodetectors

An alternative to PMTs are photon detectors based on the *silicon photodiode*.¹³ A simple photodiode consists of a thin piece of silicon (typically a few hundred microns thick), which has been carefully doped with impurities to create a favorable electric field profile in the material (Figure 11). A small voltage of 10^2 to 10^3 V is applied across the silicon diode. When a scintillation light photon interacts in the silicon, it often has sufficient energy to liberate an electron from the lattice structure of the silicon. The vacancy it leaves behind, known as a hole, has the properties of a net positive charge. Under the applied electric field, the electron drifts towards the anode (positively charged electrode) and the hole drifts towards the cathode (negatively charged electrode), constituting an electric current that can be measured. The quantum efficiency of photodiodes is approximately 60% to 80%, providing a much more efficient conversion of photons to electrons than is possible with PMTs. However, photodiodes have no internal gain, producing only one detected electron-hole pair per scintillation photon. This leads to a signal that is roughly 10^6 times weaker than a PMT signal, reducing the signal-to-noise of the pulses and degrading the ability to determine the energy deposited in the scintillator. The low signal-to-noise also necessitates the use of long integration times in the electronics, reducing the ability to time the arrival of the pulses (an important aspect for PET). Therefore, photodiodes are generally not suitable for use in PET.

FIGURE 11. Schematic cross-section of a typical photodiode. Scintillation photons enter through the entrance window and have sufficient energy to liberate an electron-hole pair in the silicon. The electric field profile in the silicon moves the electrons and holes towards the anode and cathode, respectively, creating a current that can be measured. Each light photon produces at most one electron, so the signal levels are very low. Avalanche photodiodes are very similar in structure, except the applied voltage is much higher, providing electrons with sufficient energy to create further electron-hole pairs in the silicon. Each light photon produces a signal of up to 10^2 to 10^3 electrons through this avalanche mechanism.



A modification of the photodiode leads to a device known as the *avalanche photodiode* (APD).¹⁴ Here, the voltage applied across the photodiode is much higher and creates a situation where an electron gains enough energy between collisions in the silicon to release further electrons. This leads to an avalanche effect, similar to that seen in the photomultiplier tube. The gain in these devices critically depends on factors such as the applied voltage and temperature. Special care needs to be taken in the fabrication and operation of these devices to obtain stable results. However, gains of 10^2 to 10^3 are typical, yielding improved signal-to-noise over photodiodes. Once again, the quantum efficiency is in the 60% to 80% range. When combined with the relatively high gain, this leads to roughly equivalent performance in terms of energy and timing performance compared with PMT-based detectors. APDs are now available both as single-channel units (ranging in size from 1 mm to over 2 cm in diameter) and as multi-element arrays. APDs allow for a more compact PET scanner design and may in the future replace PMTs as the photon detector of choice.

Block detector

The majority of dedicated PET scanners in use in the early part of the 21st century have detectors based on the block design proposed originally by Casey and Nutt.¹⁵ A schematic of the *block detector* is shown in Figure 12. A relatively large block of scintillator material (typically 4×4 cm in area by 3 cm deep) is segmented into an array of smaller detector elements (typically 8×8). The saw cuts are filled with a white reflective material that helps to optically isolate individual elements within the block. The scintillator block is coupled to four single-channel PMTs. The depth of the saw cuts is empirically determined to share scintillation light in a linear fashion between the four PMTs as a function of the position of the annihilation photon interaction within the block. For example, if an annihilation photon interacts in the corner detector element, the deep cuts ensure that virtually all the scintillation light photons that are produced from the interaction end up in the PMT sitting directly underneath that element. Alternatively, an event interacting towards the middle of the block, where the cuts are shallower, results in a roughly equal spread of scintillation light among all four PMTs. By careful design of the depth of the cuts, and with sufficient scintillation light, interactions in each detector element will produce a unique distribution of scintillation light and, therefore, signals on the four PMTs.

In practice, an X and Y coordinate is calculated for each annihilation photon that interacts in the block detector based on:

$$X = (S_A + S_B - S_C - S_D) / (S_A + S_B + S_C + S_D) \quad (17a)$$

and

$$Y = (S_A + S_C - S_B - S_D) / (S_A + S_B + S_C + S_D) \quad (17b)$$

where S_A , S_B , S_C and S_D are the four PMT signals shown in Figure 12. Figure 13 shows the result of an experiment in which the surface of a block detector is uniformly irradiated with 511-keV annihilation photons with the event locations histogrammed into a two-dimensional (2-D) image based on the calculated X,Y locations. This real measurement shows considerable spatial distortions in the array of spots, which is due to the fact that it is not possible to





FIGURE 12. Schematic drawing of a typical PET block detector. A block of scintillator is segmented into an 8×8 array using a diamond saw. White reflective material is used in the saw cuts to optically isolate elements. Depth of the saw cuts determines the spread of scintillation light onto four single-channel photomultiplier tubes. By looking at the ratio of signals in the four PMTs, the detector element in which an annihilation photon interacted can be determined. (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)

FIGURE 13. Image resulting from flood irradiating the front surface of a block detector with 511 keV photons, applying Equation 17 to the resulting PMT signals and displaying an image of a histogram of the X , Y signals. This measurement is known as a flood histogram. The individual detector elements in this 8×8 detector module can be visualized (the edge crystals are hard to see, as they are binned at the extreme edges of the image). Each spot is of a finite size due to the limited number of scintillation light photons contributing to the signal used to calculate X and Y . (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)



design the cuts such that the response is completely linear across the whole detector face. Therefore, a lookup table is created from these *flood histograms* relating each calculated position X,Y to each of the 64 elements in the detector. It is also apparent that the spots are of a finite size and overlap to a certain degree. This is due to statistical fluctuations in the PMT signals used to calculate X,Y , which in turn is caused by the limited number of scintillation photons produced and subsequently detected after a 511-keV annihilation photon interacts in the detector. These fluctuations ultimately limit the size and number of detector elements that can be decoded using four PMTs. In the particular example shown in Figure 13, the 8×8 array of elements are visualized relatively clearly, but had the detector been segmented into a 16×16 array of elements, it is highly unlikely that the individual elements could have been resolved from each other.

The block detector is a very cost-effective approach to PET, as it allows on the order of 64 crystals to be decoded from just four PMTs. Because the photodetectors are one of the most expensive components of a PET scanner, this 16:1 multiplexing is the key to developing PET scanners with thousands of detector elements at a reasonable cost. A large number of detector elements implies good solid angle coverage (improving the chances of detecting annihilation photon pairs that are being emitted in all directions). The block design also leads to detector elements that are smaller than the PMTs themselves. Smaller detector elements allow gamma-ray interactions to be better localized leading to improved spatial resolution as well. The spatial resolution of the block detector is primarily determined by the width of the detector elements (assuming scintillation light is sufficient to resolve each of the elements). This width is commonly 3 to 5 mm in current generation block detectors designed for clinical PET scanners.

A further extension of the block detector design leads to the concept of *quadrant sharing*.¹⁶ In this case, larger PMTs are used and each scintillation block is placed on the corner quadrants of four PMTs as shown in Figure 14. Four PMTs are still being used to decode each block, but each PMT now actually serves four different scintillator blocks. The block in which the interaction occurs is determined by which four PMTs show a significant signal, and the location of the signal within the block is determined from Equation 17. This approach, when extended to large area detector panels, leads to almost another fourfold reduction in the number of PMTs required per detector element, giving a total multiplexing approximately 64:1, and lower overall detector cost. Alternatively, this approach can be used to decode smaller detector elements for the same sized photomultiplier tubes used in the original block detector. One drawback of the quadrant sharing approach is that it requires that blocks be structured into large planar panels and also that there is one half of a PMT width at each end of the panel which is not usable. This unusable space results in relatively large gaps between panels when they are assembled in a hexagonal or octagonal geometry to form a scanner. Because the detector is no longer modular, repair and replacement is also more difficult. However, because of the dramatic cost-saving in this approach, it is being implemented in commercially available PET systems. PET scanners that use block detector and quadrant-sharing approaches are described in more detail in the PET System Design section (p. 107).

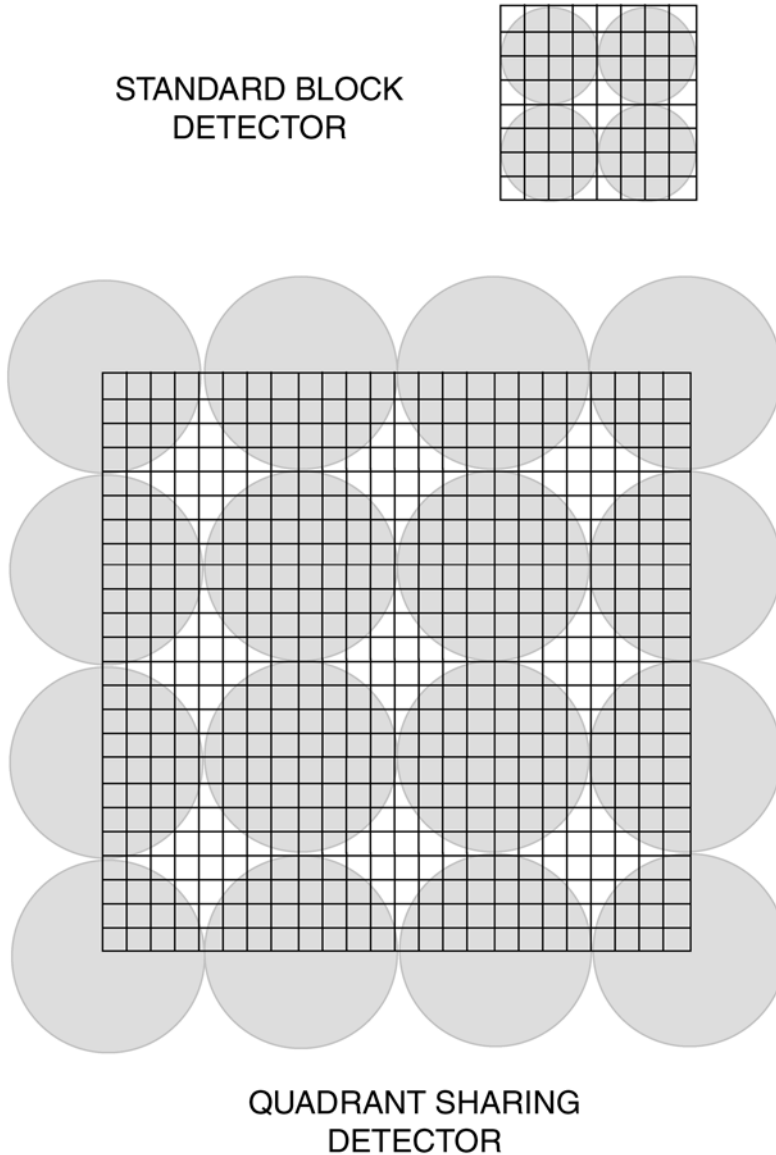


FIGURE 14. Concept of quadrant sharing that enables detector elements to be decoded using a smaller number of larger diameter PM tubes. For a given detector element size, this approach can reduce the number of PM tubes by almost a factor of four if large panels are constructed. (Adapted with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)

Continuous gamma camera detector

The other major approach to constructing a PET detector is based on a large-area, continuous plate of NaI(Tl) scintillator coupled to a matrix of PMTs as shown in Figure 15. This detector is essentially the same that is used in conventional nuclear medicine gamma cameras, although when the detector is designed specifically for PET, a thicker crystal is used to provide sufficient efficiency at 511 keV.¹⁷ The location of an interaction is determined by the




FIGURE 15. Photograph of a large-area NaI(Tl) detector designed for PET applications. The scintillator plate is 50 cm long by 15 cm wide by 2.5 cm thick and is read by thirty, 5-cm diameter PM tubes. Six of these detectors have been used in an hexagonal array to form a PET scanner. (Photograph courtesy of Dr. Joel Karp, University of Pennsylvania.)

distribution of scintillation light among the PMTs (Figure 16), with the signal from each PMT being digitized and then appropriately weighted such that the position determined from the PMT outputs is linearly related to the position of interaction. The position information provided from these detectors is con-




FIGURE 16. Schematic cross-section through a continuous gamma camera detector designed for 511-keV annihilation photons. A thick continuous sheet of NaI(Tl) scintillator is viewed by an array of PM tubes. The NaI(Tl) crystal is covered in reflective material on the sides and back to help direct more scintillation light towards the PM tubes and is hermetically sealed in a thin metal case, with a glass front window to allow scintillation light to reach the PM tubes.

tinuous. The binning of the position data can be chosen to satisfy sampling criteria for image reconstruction (see Limitations of Filtered Back-projection, p. 80). This differs from the block detector in which position information is determined by which individual detector element produces a signal. The block detector, therefore, produces discrete position information, with a sampling interval equal to the center-to-center spacing of the detector elements.

The performance of continuous detectors critically depends on the number of scintillation photons detected, as this directly impacts the spatial resolution of the detector by determining the signal-to-noise of the PMT signals that are used to calculate the position of interaction on the detector face. Therefore, this approach has only been successfully used with very high light output scintillators such as NaI(Tl). The thickness of the crystal (and, therefore, the efficiency for stopping 511-keV annihilation photons) is limited to approximately 25 mm, as the spatial resolution degrades with increasing crystal thickness. This is due to the fact that the scintillation light will spread over a larger area before reaching the PMTs, producing a lower amplitude signal across a larger number of PMTs. For a constant noise level in the PMTs and electronics, the signal-to-noise of the signals being used to calculate the position is therefore poorer and the positioning accuracy is degraded. With a 10-mm thick NaI(Tl) crystal, it is possible to achieve an intrinsic spatial resolution as high as ~ 3 mm; at 25 mm thickness this degrades to 4 to 5 mm (Dr. Joel Karp, unpublished observation). Special efforts must also be made with the electronics to allow this large-area detector to handle multiple events occurring in different parts of the detector at the same time. Otherwise, detector *dead time* (the time required to process an event before another event can be properly recorded) becomes a limiting factor in overall performance.

These detectors generally are large flat plates (typically 30–50 cm in size). Significant deadspace occurs at the edges of the detectors due to the need to hermetically encapsulate NaI(Tl) which is highly hygroscopic. Furthermore, the edges of the scintillator plates yield poor spatial resolution due to alteration of the shape of the light distribution by scintillation light that interacts with the edges of the crystal. Curved NaI(Tl) scintillator plates have become available and allow for large-area PET detectors to be constructed as segments of a ring, although regions of poor spatial resolution and gaps in active detector area still remain at the interface between segments. A system for brain imaging has even been made from a single annular NaI(Tl) crystal, eliminating deadspace completely.¹⁸ However, this is a more expensive approach, and crystal failure would result in the loss of the entire system. More details on PET systems based on NaI(Tl) detectors can be found in PET System Design (p. 107).

Other scintillation detectors

Position-sensitive and multi-channel photomultiplier tubes

Several research PET systems have been designed around multi-channel PMTs (MC-PMTs) and position-sensitive PMTs (PS-PMTs). MC-PMTs consist of an array of small, separate PMT channels within a single package. Position-sensitive PMTs have a segmented X and Y readout and are designed such that

the output signals are approximately linearly related to the position of the scintillation light that is incident on the photocathode. Both MC-PMTs and PS-PMTs could be used to replace the four single-channel PMTs in a block detector or to decode individual scintillator detector elements arranged into an array. Because of their compact size and the ability to provide positional information, these devices are often used to decode arrays with relatively large numbers of very small scintillator elements for high-resolution PET applications.^{19–21} Both MC-PMTs and PS-PMTs often have significant amounts of deadspace around their periphery, and, hence, they have sometimes been used with fiber optic coupling between the scintillator and the PMT to allow tight packing of detectors in a ring configuration.²² As the cost of PS and MC-PMTs are quite high relative to single-channel PMTs, their use has been largely limited to more specialized applications such as breast and animal imaging where a smaller number of detectors are required.


Depth-encoding detectors

The detectors described so far have all focused on determining the X, Y location of an interaction (the interaction location projected onto the front surface of the detector). This is fine for thin detectors; however, PET detectors typically require a 2- to 3-cm thickness of scintillator to achieve adequate efficiency. The detectors cannot be considered to be thin. The detectors discussed so far do not provide any information on the depth of the interaction of the annihilation photons inside the scintillator. This uncertainty in *depth of interaction* leads to a loss of spatial resolution in PET images as demonstrated in Resolution: Coincidence Response Functions (p. 38). If the PET detector can determine the Z or depth coordinate of the interaction, this resolution degradation would be removed. This is an active area of research and many possible approaches have been proposed. Two methods have emerged that promise a certain degree of success and are now finding their way into PET scanner design (Figure 17). The first approach uses two layers of scintillator materials (known as a *phoswich*) to provide a two-level (top half, bottom half) depth encoding capability.^{23,24} The scintillator materials are differentiated by their different decay times. The layer in which the interaction occurs can be simply determined by looking at the decay time of the pulses. The second approach places photodetectors at both ends of a scintillator array and uses the ratio of the signals between the two photodetectors to provide a measure of the depth of interaction. The photodetector at the far end must be thin and compact, both from a geometric standpoint, and also to minimize attenuation of the annihilation photons that must pass through this detector before reaching the scintillator. This approach has been studied extensively by Moses and colleagues²⁵ using a PIN photodiode array at the far end of the scintillator array to identify the crystal of interaction. A single-channel PMT at the back of the scintillator array provides information about the energy of the event and the timing signal. The ratio of the photodiode and PMT signal provides the depth of interaction information. The surface treatment of the scintillator elements and the use of reflectors along the sides of the crystal are critical in determining the distribution of scintillation light to the two ends of the array and hence the success of this approach.²⁶

Avalanche photodiodes

Extensive research has been invested in the development and application of APDs (p. 21) for PET. Both single-channel APDs (up to about 15 mm in size) and arrays of smaller APDs are now available. An 8×8 array of APD elements

FIGURE 17. PET detectors with depth-encoding capability. A: Detector is similar to the standard block detector but is made of two layers of scintillators that have different decay times. An interaction can be assigned to the top or bottom layer, depending on the decay time of the pulse that it generates. This provides one-level (top or bottom) depth of interaction information. (Reproduced from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, 3rd ed, W.B. Saunders, New York, 2003, with permission from Elsevier). B: An array of scintillator elements has photodetectors at both ends. A silicon PIN photodiode is used to determine which crystal the interaction took place in, and the single-channel PMT at the back of the array is used to generate the fast timing signal necessary for PET and a high signal-to-noise measure of the deposited energy. The ratio of the signal in the photodiode and PMT gives an indication of the depth of interaction within the detector element.²⁵ This provides continuous depth of interaction information but requires careful calibration of the depth information. Abbreviations: LSO, lutetium oxyorthosilicate.



with a 1-mm center-to-center spacing is shown in Figure 18. The thin profile of these photodetectors encourages a number of new design possibilities, for example, multiple concentric rings of detectors,²⁷ or the use of APD arrays on both front and back surfaces of scintillator arrays.²⁸ Both these designs also provide depth of interaction information, in addition to identifying the detector element in which the interaction occurred. APDs have been used successfully on a small scale in PET scanners for animals.^{29,30} The stability and longterm reliability of APDs and the need for large numbers of channels of electronics have so far limited their widespread use, although an increasing role is expected in the future.

Other gamma ray detectors

Scintillation detectors have been the dominant detector technology in PET, largely due to their high efficiency, robustness, and reasonable cost, particularly with the block and gamma camera approaches that require a relatively small number of




FIGURE 18. Photograph of an 8×8 avalanche photodiode photodetector array. Each pixel measures $1 \times 1 \text{ mm}^2$. (Photograph courtesy of Kanai Shah, Radiation Monitoring Devices Inc., Watertown, MA.) (See color insert.)

PMTs to read a large surface area of scintillator. However, other technologies have been, and are continuing to be, explored for possible applications in PET.

Multiwire proportional chambers (MWPCs) have long been used in high-energy physics as very cost effective detectors for covering large areas at high spatial resolution. These detectors consist of a chamber of gas with a set of finely spaced anode wires at high positive potential. Above and below the anode wire plane are cathode wires or strips held at ground potential that run in orthogonal directions. When the gas is ionized by a charged particle, the resulting electrons are attracted to the nearest anode wire and, because of the very high electric field close to the wire, an avalanche effect occurs, resulting in further ionization and a large signal. This in turn induces a charge on the nearest cathode strips, which provide information on the x and y position of the event. The fine spacing of the wires and the cathode strips allows very high spatial resolution to be achieved. For applications in PET, the incoming annihilation photons must first be converted into charged particles (electrons). This conversion can be achieved by making the cathode strips from thin layers of lead³¹ or by using some form of converter such as a stack of thin lead sheets interlaced with insulating sheets that are then drilled with a fine matrix of holes.³² Incoming annihilation photons can interact in the lead, ejecting electrons which are then drifted to the anode wires for amplification and detection. Figure 19 illustrates the principles of such a detector.

The problem for the application of MWPCs in PET has largely been achieving sufficient efficiency in conjunction with difficulties in matching the counting-rate performance, timing resolution, and energy resolution of scintillation detectors. To improve efficiency, multiple MWPC units have been stacked on top of each other, but even so, the efficiency of these detectors is typically on the order of 1% to 2%, compared to the 30% to 90% efficiency typical of scintillation detectors at 511 keV. Some of this efficiency loss can however be compensated by the good solid angle coverage of these large area detectors when placed in a scanner configuration. A second approach to improving the efficiency of these detectors has been to replace the thin lead converters with a sheet

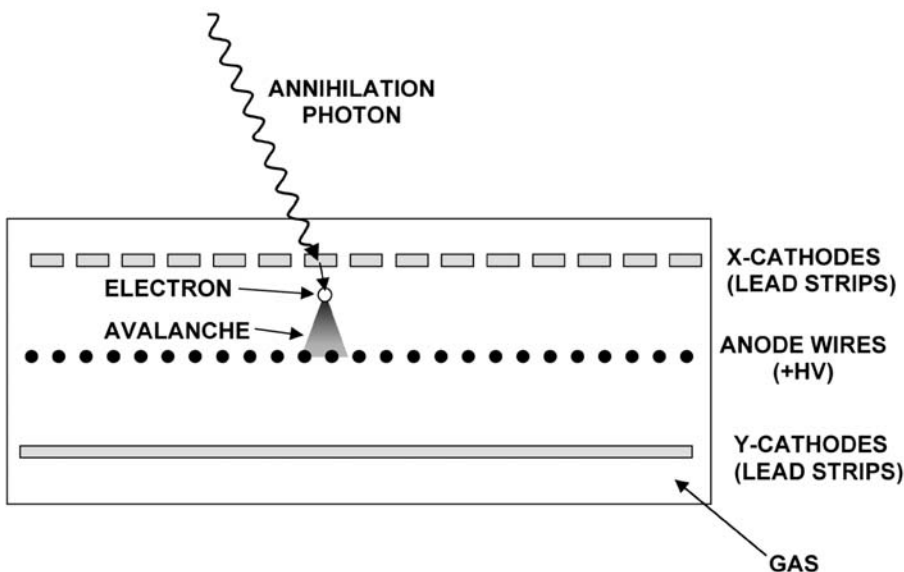


FIGURE 19. Diagram showing cross-section through a multiwire proportional chamber (MWPC) detector for PET. The annihilation photon interacts in thin strips of lead, ejecting an electron into the gas which is accelerated by a high potential difference to the anode wires, creating an avalanche of electrons at the wire. This, in turn, induces a signal on the nearby cathode strips that are alternately arranged in the x and y directions to provide the x and y coordinates of the event. Multiple units can be stacked on top of each other to improve efficiency and provide depth of interaction information.

or crystals of BaF_2 scintillator material. BaF_2 is one of the few scintillators that produces light as ultraviolet radiation (UV) that has sufficient energy to photoionize the gas tetrakis-dimethylamino-ethylene (TMAE). Incoming annihilation photons interact in the BaF_2 , and the subsequent scintillation light ionizes TMAE gas in the MWPC, with the position signal determined by the induced signal on cathode strips or wires as described previously.³³ While this improves efficiency, it still does not match the efficiency of BGO or lutetium oxyorthosilicate (LSO) scintillation detectors; the energy resolution and timing resolution remain poor. Working PET systems based on MWPCs or MWPC/ BaF_2 combinations have been developed for clinical studies,^{34–36} although the most successful application to date has probably been in small animal imaging.^{32,37}

Direct detection using semiconductor materials

The approach of direct annihilation photon detection using semiconductor materials has been relatively neglected but is likely to gain increasing attention in the future. The concept is to use the semiconductor material itself to directly detect the annihilation photons, thus eliminating the need for a scintillator. The detector would work like a standard silicon photodiode (Figure 11); however, in this case the annihilation photons directly create the electron-holes pairs. Silicon, although the most well-developed semiconductor material, has very poor efficiency at 511 keV and would not be the material of choice. Other semiconductor materials such as cadmium telluride (CdTe) or cadmium zinc telluride (CZT) have a stopping power that is similar to NaI(Tl) at 511 keV

and might be viable detector materials for PET. These materials are difficult to manufacture in bulk and are costly at this time. Achieving good energy and timing resolution from the relatively thick pieces of material needed to provide reasonable efficiency is a challenge. The approach is attractive in that it eliminates the conversion stage represented by the scintillator. The signal produced is very robust, as each 511-keV photon interaction will produce a large number of charge carriers ($\sim 10^5$). Other, even more dense (and, therefore, better stopping power) semiconductors such as PbI and TlBr exist, but these are in a fairly primitive stage of development and are mainly used as thin films at the present time. Should such materials become available in bulk at a reasonable price, they could be promising alternatives to scintillator-based detectors currently in use for PET.

DATA COLLECTION AND PET SYSTEM CONFIGURATIONS

Coincidence detection

In contrast to other nuclear imaging techniques, PET does not rely on absorptive collimation to determine the direction and location of the emitted photons. Instead, a technique referred to as *coincidence detection* is used. A simple coincidence detection system is illustrated in Figure 20, which consists of a pair of radiation detectors with associated electronics (amplifiers, pulse height analyzers, high voltage) and a coincidence circuit. If an annihilation occurs somewhere between two high-efficiency detectors, and the direction of the two 511-keV photons is such that each will have a chance to interact with one of the two detectors, it is very likely that a coincidence event will be recorded. Because all annihilation photons are emitted approximately 180° apart, a recorded coincidence indicates an annihilation occurred somewhere along the line (or more accurately, the volume) connecting the two detectors. This line or volume from which the detector pair can detect coincidences usually is referred to as a *line of response* or LOR. To reconstruct a complete cross-sectional image of the object, data from a large number of these LORs are collected at different angles and radial offsets that cover the field of view of the system (Image Reconstruction, p. 70).

The two detectors and associated circuitry should, under ideal circumstances, simultaneously generate the logic pulses necessary to generate a coincidence. However, due to stochastic processes in the emission of light in the scintillation detectors, a random time delay occurs in exactly when the detectors respond following the absorption of the annihilation photons in the detectors. This uncertainty in response or time resolution depends on the characteristics of the detector, primarily scintillation decay time constant and light output (Scintillators, p. 17). Furthermore, small differences are noted in the arrival times of the two photons depending on the difference in the distance of the annihilation site to each detector (Annihilation, p. 5). To avoid missing coincidence events, the logic pulses must have a certain finite width to ensure that the pulses overlap despite the finite time resolution. Typically, the width of the logic pulses, τ , should be at least as wide as the timing resolution of a pair of detectors (measured in FWHM). A typical timing resolution for a BGO- or NaI(Tl)-based PET detector is approximately 5 to 6 nanoseconds FWHM, while for LSO it is approximately 2 to 3 nanoseconds FWHM. It is important to keep the pulses as narrow




FIGURE 20. Diagram of a basic coincidence circuit. The two scintillation detectors with are connected to individual amplifiers (Amp) and pulse height analyzers (PHA). When a photon interacts in either of the detectors, the signals are amplified and analyzed to determine if the energy is above a certain threshold. If the energy criterion is satisfied, a logic pulse is generated by the PHA. These pulses are fed into a coincidence module (Coinc), which determines if there is an overlap of two pulses from the individual channels. An overlap occurs if both pulses occur within a time period of 2τ (e.g., they differ from each other in time by $\leq \tau$), where τ is the width of the pulse. If this is the case, a coincidence has been detected and the coincidence circuit generates a logic pulse that is fed into a counter for registration of the event. In a PET imaging system, the memory location corresponding to the two detectors in which the interaction occurred is incremented by one.

as possible to minimize the detection of events from unrelated decays that happen to strike the detectors within the time window determined by the overlap of the two logic pulses (Types of Events, p. 35).

PET camera: general concepts

The two small detectors shown in Figure 20 would not make a very effective PET system. They would only detect annihilation photons from decays occurring in the volume between the detectors and the tiny fraction of those decays in which the annihilation photons are directed towards the detectors. A complete PET system (Figure 21) consists of a large number of detectors (e.g., block detectors) placed around the object to be imaged. The most common detector configuration of a PET system is the ring geometry. When referring to directions within the plane of the detector ring, the terms *transverse* or *transaxial* are used. When referring to directions perpendicular to this plane (along the direction of the patient bed), the term *axial* is used. It is also possible to use a smaller number of large-area position sensitive detectors in a polygonal arrangement (e.g., continuous detector panels). Both these geometries allow many different LORs to be measured simultaneously and can sample an entire slice or cross-section of the object that is being imaged with little or no detector motion. This feature is of particular importance if the system is to be used for rapid dynamic imaging of the distribution of radiolabeled tracers.




FIGURE 21. Schematics of four common PET scanner configurations. A: Stationary block ring system. B: Rotating block ring system. C: Stationary NaI(Tl) system using six flat detectors. D: Stationary NaI(Tl) system using curved continuous panels.

Some lower cost PET systems are comprised of a partial ring of detectors, often two opposing detectors or detector plates. However, to acquire sufficient data to reconstruct a tomographic image, these detectors must be rotated around the object.

To improve the overall detection efficiency in modern PET scanners, the detectors usually extend 15 cm or more in the axial direction. This can be accomplished by stacking several rings of detectors next to each other or by having two-dimensional continuous detectors with large axial dimensions. Many slices of data can then be acquired simultaneously, ultimately producing a set of image slices that can be stacked into a 3-D image volume.

In PET scanners based on large continuous detectors, each detector will be in coincidence with detector heads on the opposing side of the scanner (Figure

21). In PET systems constructed from block detectors, the number of possible coincidence combinations is proportional to the square of the number of detector elements. It is, therefore, not practical to have a dedicated coincidence circuit for each possible detector pair. Instead, a large number of detectors are grouped together into detector banks or buckets^{38,39} where each detector group will look for coincident events in one or more opposite detector groups. Following the detection of a coincidence event, the electronics will then identify the detector elements (block detector) or detector locations (continuous detector) that produced the coincidence. The electronics will also check that the energy deposited in each detector is in the appropriate range for a 511-keV event. The energy window used is related to the energy resolution of the detector. A typical PET scanner uses an energy window of 350 keV to 650 keV to be sure of including all 511-keV photons, while rejecting photons that have lost a substantial fraction of their initial energy by scattering in the body.

Events that meet both the energy and timing criteria are then conveyed to the sorting hardware that writes the raw data in one of two ways: In *list mode*, each event is individually written to a file, with information about the two locations at which the annihilation photons interacted and the time at which the event occurred. In *histogram mode*, a memory location is assigned to each possible LOR, and each time a valid event is detected in that LOR, that memory location is incremented by 1. This provides the integrated number of events detected in each LOR and is frequently the most efficient manner to store the data, except for short duration acquisitions on cameras with very large numbers of LORs, where the average number of events per LOR < 1 . List mode data are advantageous for dynamic studies, as the events can be sorted into time “bins” after the completion of the study. In histogram mode, the events are integrated over a time interval that must be specified prior to data acquisition.

Types of events

Under ideal circumstances, only true coincidences would be recorded, that is, only events where the two detected annihilation photons originate from the same radioactive decay and have not changed direction or lost any energy before being detected. However, due to limitations of the detectors used in PET and the possible interaction of the 511-keV photons in the body before they reach the detector, the coincidences measured are contaminated with undesirable events, which includes *random*, *scattered* and *multiple* coincidences (Figure 22). All these events have a degrading effect on the measurement and need to be corrected to produce an image that represents as closely as possible the true radioactivity concentration. Another point to consider is that the vast majority (typically 90% or more) of photons detected by the PET scanners are *single* events, in which only one of the two annihilation photons is registered. The partner photon may be on a trajectory such that it does not intersect a detector (most PET scanners provide relatively modest solid angle coverage around the object), or the photon may not deposit sufficient energy in a detector to be registered or may not interact at all. These single events are not accepted by the PET scanner, but they are responsible for random and multiple coincidence events (see below). Because they must still be processed by the electronics to see if they form part of a coincidence pair, they are the determining factor in issues related to detector dead time (Dead Time Correction, p. 67).




FIGURE 22. Illustration of the four main coincidence event types. A: True coincidence. Both annihilation photons escape the body and are recorded by a pair of detectors. B: Scattered coincidence. One or both of the two annihilation photons interacts in the body prior to detection. This results in a mispositioning of the event. C: Random coincidence: A coincidence is generated by two photons originating from two separate annihilations. These events form a background in the data that needs to be subtracted. D: Multiple coincidence: Three or more photons are detected simultaneously. Due to the ambiguity of where to position the events, these normally are discarded. (Reprinted from *Physics in Nuclear Medicine*, 2nd ed, Cherry SR, Sorenson JA, Phelps ME, W.B. Saunders, New York 1986, with permission from Elsevier.)

Accidental coincidences

When positron annihilation occurs, the two 511-keV photons are emitted simultaneously. Therefore, the detectors should ideally respond simultaneously. Because of the finite time resolution of the detectors, as discussed earlier, signals must be accepted if they occur within a certain finite time interval or timing window. Because of the finite width of the timing window, it is possible that two unrelated single annihilation photons can be detected and registered as a valid coincidence. These unrelated events are referred to as *accidental* or *random* events. Because the random events are produced by photons emitted from unrelated isotope decays, they do not carry any spatial information about the ac-

tivity distribution and produce an undesired background in the final images. If the individual photon detection rates (counts per second) in a pair of detectors are given by N_1 and N_2 , then it can be shown that the rate of random coincidences, N_R (randoms per second) is given by:

$$N_R = 2\tau N_1 N_2 \quad (18)$$

where τ is the width of the logic pulses produced when a photon is absorbed in the detector. The term 2τ is often referred to as the *coincidence timing window*. Because the individual detection rates N_1 and N_2 are directly proportional to the activity in the field of view of the scanner, the rate of random coincidences is proportional to the square of the activity in the field of view. The randoms rate is directly proportional to the coincidence timing window, which is why it is important not to make this any wider than required by the timing uncertainties in true coincidence events.

EXAMPLE 8

Two detectors register a counting rate of 100,000 counts per second each when operating independently. What would the rate of random events be if they were placed in coincidence, and if the logic pulse width generated by each detector had a width of 6 nanoseconds?

ANSWER

From Equation 18, the randoms rate would be:

$$N_R = 2 \times (6 \times 10^{-9}) \times (100,000)^2 = 120 \text{ random events per second}$$

Scattered coincidences

Scattered coincidences are another type of background event in need of correction. These events are in essence true coincidences, but one or both of the two annihilation photons has undergone a Compton scatter interaction and changed direction before they reach the detector pair. Using the coincidence detection technique, it is assumed that all detected coincidence events originate from an annihilation which, in turn, originates from a position anywhere on a line connecting the detector pairs. Because of the change in direction of the photon(s) in a scattered event, this is not true and the event is assigned to the incorrect LOR. If not corrected, the scattered events produce a low spatial frequency background that reduces contrast. The distribution of scattered events depends on the distribution of the radioactivity and the shape of the scattering medium (i.e., the patient). As will be discussed later, this is probably the most difficult correction to perform in PET. The fraction of scattered events detected can range from 15% to well over 50% in typical PET studies, depending on the size of the object and the geometry and energy resolution of the PET scanner.

Multiple coincidences

Although only two detectors are required to be activated within the coincidence time window to register a valid coincidence, at high count-rates it is possible that three or more detectors are involved. In this case, it becomes ambiguous where the event should be positioned. Because of this ambiguity, these multiple coincidences normally are discarded. However, they can contain in-

formation about the quantity and spatial location of positron emissions because these events are often composed of a true coincidence together with a single photon from an unrelated decay. In this situation, up to three possible LORs can intersect the field of view, only one of which will be correct. In some circumstances, it may be better to randomly select one of the possible LORs rather than completely discarding the event.

Prompt coincidences

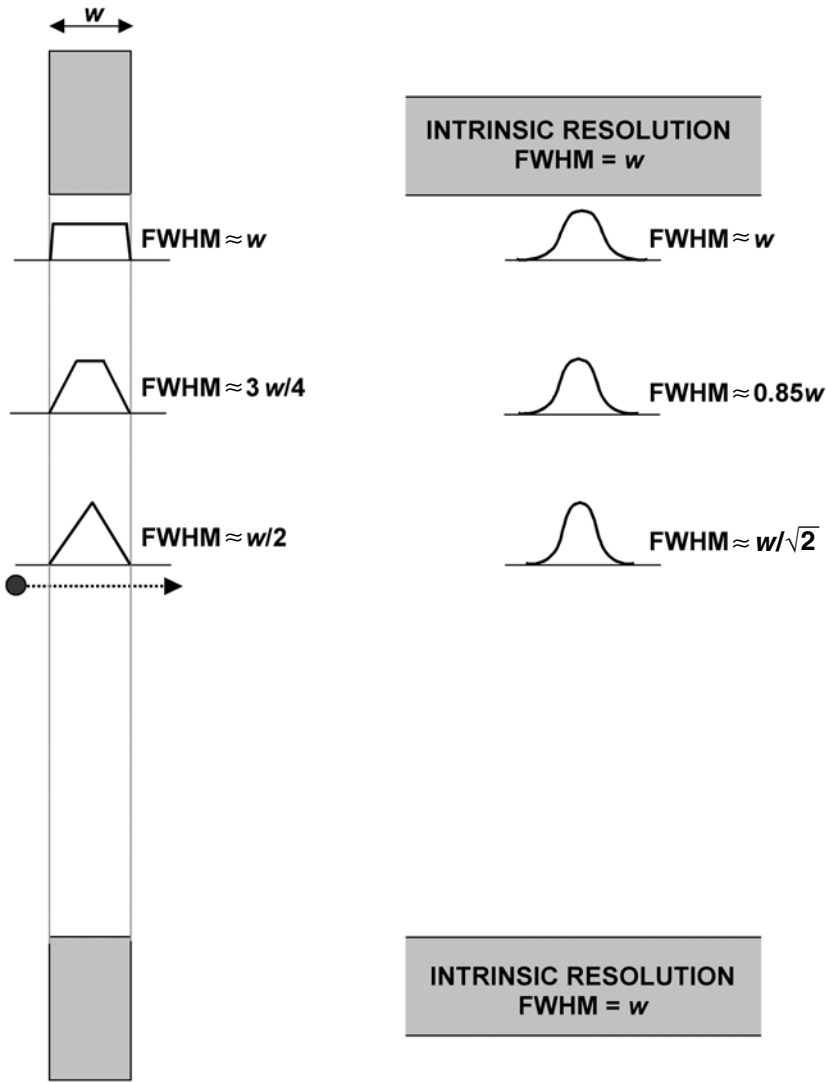
The total number of events detected by the coincidence circuit in a PET scanner are referred to as *prompt* coincidences. These events consist of true, scattered, and accidental coincidences where the true coincidences are the only ones that carry spatial information regarding the distribution of the radiotracer. It is, therefore, necessary to estimate what fraction of the measured prompt coincidences arise from scattered and accidental coincidences for each of the LORs. The contribution of scattered and accidental coincidences is then subtracted from the prompt coincidences to yield the net true coincidence rate for each measured LOR. Because both the scattered and accidental events are, in general, estimates, the accuracy of these estimates will affect the accuracy of the net calculated true coincidence rate. Any statistical or systematic noise in these estimates will also propagate into the net true coincidence rate.

Resolution: coincidence response functions

As discussed in Positron Range and Noncolinearity (p. 9) an ultimate resolution limit can be achieved in PET due to the physics of the positron decay. In addition to this limit, the design and properties of the detector used in the PET scanner, and the system geometry, will also contribute to the final image resolution. The intrinsic detector resolution can be divided into two main components: geometric and physical. The geometric component can be seen as the best possible resolution that can be attained for a particular design using a scintillation material with ideal detection properties (e.g., 100% detection efficiency). The physical component is caused by nonideal properties of the detectors (e.g., detector scatter, light sharing, cross-talk, and so on).

For a pair of discrete detectors, such as the detector elements in a block detector, the geometric resolution at the mid-point between the detector pair, can be described by a triangular shaped coincidence response function (Figure 23A) where the FWHM equals one half the detector width. This response function can be obtained by considering how many coincidence events would be detected as a small source is moved across the detector face. At any position close to either of the two detectors the response function changes and becomes trapezoidal in shape and eventually becomes a square function at the front surface of the detector. The intrinsic resolution in these types of detectors is therefore strongly influenced by the width of the detector elements.

In continuous detector systems, the intrinsic spatial resolution of the detector is largely determined by the number of scintillation photons available for determining the position of the event, not by geometric factors. The same types of physical components listed above, also contribute to the intrinsic spatial resolution. The intrinsic spatial resolution of a continuous detector can typically be approximated by a Gaussian with a particular FWHM. If this is the case, then it can be shown that the coincidence response function at the midpoint between



a) DISCRETE DETECTORS

b) CONTINUOUS DETECTORS

FIGURE 23. A: Geometric spatial resolution as a point source is moved between two discrete detectors. At the center of the field of view, the coincidence response function has a triangular shape with a FWHM equal to half the detector width w . As the source is moved towards one of the detectors, the coincidence response function becomes trapezoidal in shape and the FWHM increases linearly with the distance from the center of the field of view. B: Coincidence response function for a pair of continuous detectors, each with an intrinsic spatial resolution described by a Gaussian with a $FWHM = w$. The coincidence response function at the center of the field of view is $w/\sqrt{2}$, increasing to w at the detector face. (Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. Physics in Nuclear Medicine, W.B. Saunders, New York, 2003.)

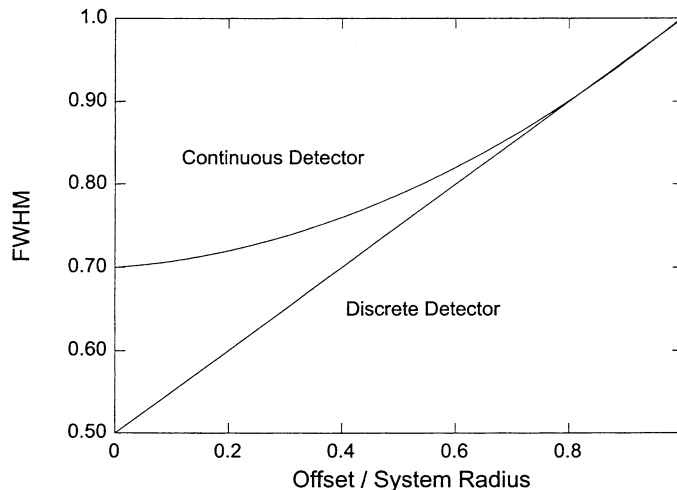



FIGURE 24. Comparison of the FWHM of the coincidence response function for discrete and continuous detectors with $w = 1$ (as defined in Figure 23) as a function of source position (offset) relative to the two detectors. A source offset of 0 corresponds to a source located exactly halfway between the detectors.

the detector pair is equal to the individual detector resolution $\text{FWHM}/\sqrt{2}$. For locations closer to one detector or another, the FWHM resolution will increase and at one detector face will eventually become equal to the intrinsic resolution of the individual detectors (Figure 23B). Figure 24 shows the change in the FWHM of the coincidence response function for discrete detectors and continuous detectors as a function of the source location.

The thick scintillation detectors used for PET imaging (typically 2–3 cm) lead to another geometric effect that degrades spatial resolution. This effect, which is referred to as *detector parallax* or the *depth of interaction* effect, is caused by the fact that the annihilation photons can interact at any depth in the scintillator material. Consider a ring geometry scanner consisting of either discrete or continuous detectors (Figure 21, top left or bottom right). At the center of the field of view, all emitted photons will enter the detectors perpendicular to the detector face. However, when a source is located with a radial offset, the detectors are angled with respect to the line of response and the annihilation photons may penetrate through the first detector they encounter and be detected in an adjacent detector as shown in Figure 25. There are two consequences of this effect. The coincidence response function becomes broader (because the detectors are at an angle and present a larger area to the line of response) and the event is also mispositioned towards the center of the scanner with respect to the line joining the two detectors of interaction. The amount of broadening depends on the width and thickness of the scintillator elements, the absorption characteristics of the scintillator material, and the separation of the detectors. It results in a worsening in the radial component of the spatial resolution of PET images as you move away from the center of the field of view. This can be a significant effect in small-diameter PET scanners that use thick scintillation detectors. This effect could be reduced if the depth of the interaction in the detector could be measured, which would allow a correct placement of the event (Other Scintilla-

FIGURE 25. In a ring geometry scanner, the point spread function becomes asymmetrical with increasing radial offsets due to detector penetration and the lack of information regarding the depth of interaction within the crystal. The result is a widening of the point spread function (degrading spatial resolution) and mispositioning of events towards the center of the field of view. The severity of these effects depends on detector ring diameter, detector depth, and the detector material.



tion Detectors, p. 27). PET scanners based on a polygonal geometry (Figure 21, bottom left) also suffer from these detector parallax effects, although the degradation in resolution is spread fairly uniformly across the entire field of view rather than being concentrated towards the peripheral field of view. These effects occur because for all source locations (even at the center of the field of view), annihilation photons enter the detectors with a range of angles with respect to the detector face.

The final system resolution for a particular system design is a convolution of all the resolution response functions, including the positron range, photon noncolinearity, geometric factors, intrinsic spatial resolution (for continuous detectors), and physical factors. In addition, insufficient sampling of lines of response through the object can degrade the resolution in the final reconstructed image (Limitations of Filtered Backprojection, p. 80).

EXAMPLE 9

Assuming that detector resolution, positron range, and photon noncolinearity can be approximated by Gaussian functions, calculate the system resolution at the center of the field of view for a clinical PET scanner with an 80-cm diameter and 6-mm discrete detector elements when imaging a ^{18}F -labeled radiopharmaceutical.

ANSWER

From Figure 23 the intrinsic detector resolution, R_{int} , at the center of the field of view of the scanner is:

$$R_{\text{int}} = 6 \text{ mm}/2 = 3 \text{ mm}$$

From Equation 10, the blurring due to photon noncolinearity, Δ_{nc} , is:

$$\Delta_{\text{nc}} = 0.0022 \times 800 = 1.76 \text{ mm}$$

From Figure 5, the blurring due to the positron range of ^{18}F , Δ_{pos} , has a FWHM of 0.102 mm.

The convolution of multiple Gaussian functions is another Gaussian with a FWHM given by adding the individual component FWHMs in quadrature. The system resolution, R_{sys} , is therefore:

$$R_{\text{sys}} \approx \sqrt{R_{\text{int}}^2 + \Delta_{\text{nc}}^2 + \Delta_{\text{pos}}^2} = \sqrt{3^2 + 1.76^2 + 0.102^2} = 3.48 \text{ mm}$$

Sensitivity: detector and geometric efficiencies

One of the most important factors in designing a PET system is to maximize the *system sensitivity*, since this will be a major determinant of final image quality. The more coincidence events that can be detected and used to form the image, the better. The number of events collected is dictated by the amount of radioactivity injected, the fraction of the injected activity that reaches the tissues of interest, the imaging time, and the sensitivity of the PET system. Practical limits are set for the amount of radioactivity that can be administered and on the imaging time. Thus, system sensitivity is a key factor in obtaining high-quality images. The system sensitivity is defined as the number of events (in counts per second) detected per unit of radioactive concentration (cps/Bq/ml) in a specific phantom. It is sometimes also expressed as the fraction of radioactive decays that produce a valid coincidence event (cps/Bq). The system sensitivity is a product of several factors, which include the efficiency of the detectors at 511 keV, the solid angle coverage of the detectors, the location of the radioactivity with respect to the detectors, and the timing and energy windows applied to the data.

The *detection efficiency* ε , of an individual detector is given by the product of the detection probability of the incoming photon in the detector volume and the fraction of these events, Φ , that fall within the selected energy window (typically set at 350–650 keV). The energy window helps to reduce the influence of scattered events by only accepting events that deposit energy close to 511 keV (e.g., photopeak events). The efficiency is:

$$\varepsilon = (1 - e^{-\mu d}) \times \Phi \quad (19)$$

where μ is the attenuation coefficient of the detector material (Table 4), and d is the thickness of the detector. A valid event requires that both photons be detected in opposing detectors and be within the appropriate energy range. The *coincidence detection efficiency* is, therefore, given by the square of Equation 19:

$$\varepsilon^2 = (1 - e^{-\mu d})^2 \times \Phi^2 \quad (20)$$

The *geometric efficiency* of the system is determined by the overall solid angle (Ω) coverage of the detectors with respect to the source location and the packing fraction. The solid angle subtended by the detectors of a circular system for a point source placed at the center is given by:

$$\Omega = 4 \pi \sin [\tan^{-1} (A/D)] \quad (21)$$

where D is the diameter of the detector ring. For a PET scanner consisting of a single ring of detector elements, A is just the height of the detector in the axial direction. For scanners consisting of multiple detector rings (e.g., a ring or rings of block detectors) or continuous detectors, A depends on the maximum acceptance angle over which data will be collected. This is discussed further in the next two sections.

In the manufacturing of the detectors in a PET system using discrete detector elements or block detectors, a small gap is always between the detector elements due to the need for reflective material on the detector walls and/or detector encapsulation. This dead space will produce a reduction in the overall efficiency and is referred to as the packing fraction (φ). The packing fraction is the ratio of the detector element area (width of detector element by axial height of detector element) to the total surface area, including the dead space:

$$\varphi = \frac{\text{width} \times \text{height}}{(\text{width} + \text{deadspace}) \times (\text{height} + \text{deadspace})} \quad (22)$$

The overall system sensitivity, η , for a point source placed at the center of a ring scanner is the product of the square of the detection efficiency ε and the geometric efficiency $\varphi \times \Omega$. Expressed as a percentage, it is given by:

$$\eta = 100 \times \frac{\varepsilon^2 \varphi \Omega}{4\pi} \quad (23)$$

Notice that ε is squared as a result of the coincidence detection. Because of this, a small reduction in ε , due to either a reduction in the thickness of the scintillator or a tighter energy window, will produce a significant loss in the overall sensitivity. φ appears as a linear term because of the angular correlation of the two photons; however, this is only an approximation and the true contribution of packing fraction losses is position and geometry dependent. For a distributed source, the expression becomes more complex due to variations in both the geometric and detection efficiencies across the field of view (FOV).

EXAMPLE 10

Compute the overall system sensitivity for a central point source in a PET scanner consisting of a single ring of BGO crystals if the ring diameter is 80 cm and the detector elements measure 4.9 mm in width (transaxial) by 6 mm in height (axial) by 30 mm deep. Assume the window fraction is 80%. Assume further that the dimensions for each detector element include a 0.25-mm thick layer of reflector all around the crystal, such that the actual size of the BGO elements is 4.4 mm \times 5.5 mm in cross-section.

ANSWER

From Equation 19, and using the linear attenuation coefficient for BGO in Table 4, the detection efficiency is:

$$\varepsilon = (1 - e^{-(0.96 \times 3)}) \times 0.8 = 0.755$$

From Equation 21, the geometric efficiency is:

$$\Omega = 4\pi \sin[\tan^{-1}(0.6/80)] = 0.094$$

The packing fraction is given by Equation 22 as:

$$\varphi = (4.4 \times 5.5) \div (4.9 \times 6) = 0.673$$

The overall system sensitivity, from Equation 23, is:

$$\eta = 100 \times 0.755^2 \times 0.094 \times 0.673 / 4\pi = 0.287\%$$

This example demonstrates the low sensitivity of a single-slice PET scanner, even when a complete ring of detectors surrounds the patient and high-efficiency detectors are used. In this situation, the geometric efficiency is the limiting factor.

Data representation—the sinogram

Consider a simple PET system consisting of 32 individual detectors in a ring, scanning an object with a 2-D distribution of radioactivity denoted by $a(x,y)$ (Figure 26). The raw data, which consists of the detection of annihilation photon pairs, usually is histogrammed into a 2-D matrix, where each element in the matrix corresponds to the number of events recorded by a particular pair of detectors (or along a specific line of response). The matrix is arranged such that each row represents parallel line integrals or a *projection* of the activity at a particular angle ϕ . Each column represents the radial offset from the center of the scanner, r . The relationship that relates which elements in this matrix (r,ϕ) record data from radioactivity in the object at location (x,y) is given by:

$$r = x \cos\phi + y \sin\phi \quad (24)$$

This 2-D matrix $s(r,\phi)$ is known as a *sinogram* because a point source at a location (x,y) traces a sinusoidal path in the matrix as given by Equation 24. The mapping of detector pairs into a sinogram is shown in Figure 27. In practice, as Figure 27 demonstrates, the data in each row do not come from a single angular view but rather from two adjacent angles that are interleaved together. The samples from the second angle fall exactly half way between those from the first angle. For typical PET scanner geometry, this leads to projection views that




FIGURE 26. Imaging geometry for a single-slice PET scanner consisting of 32 detectors. The relationship between the (x,y) coordinate system of the object and the (r,ϕ) coordinate system in which PET data are commonly stored is also shown.

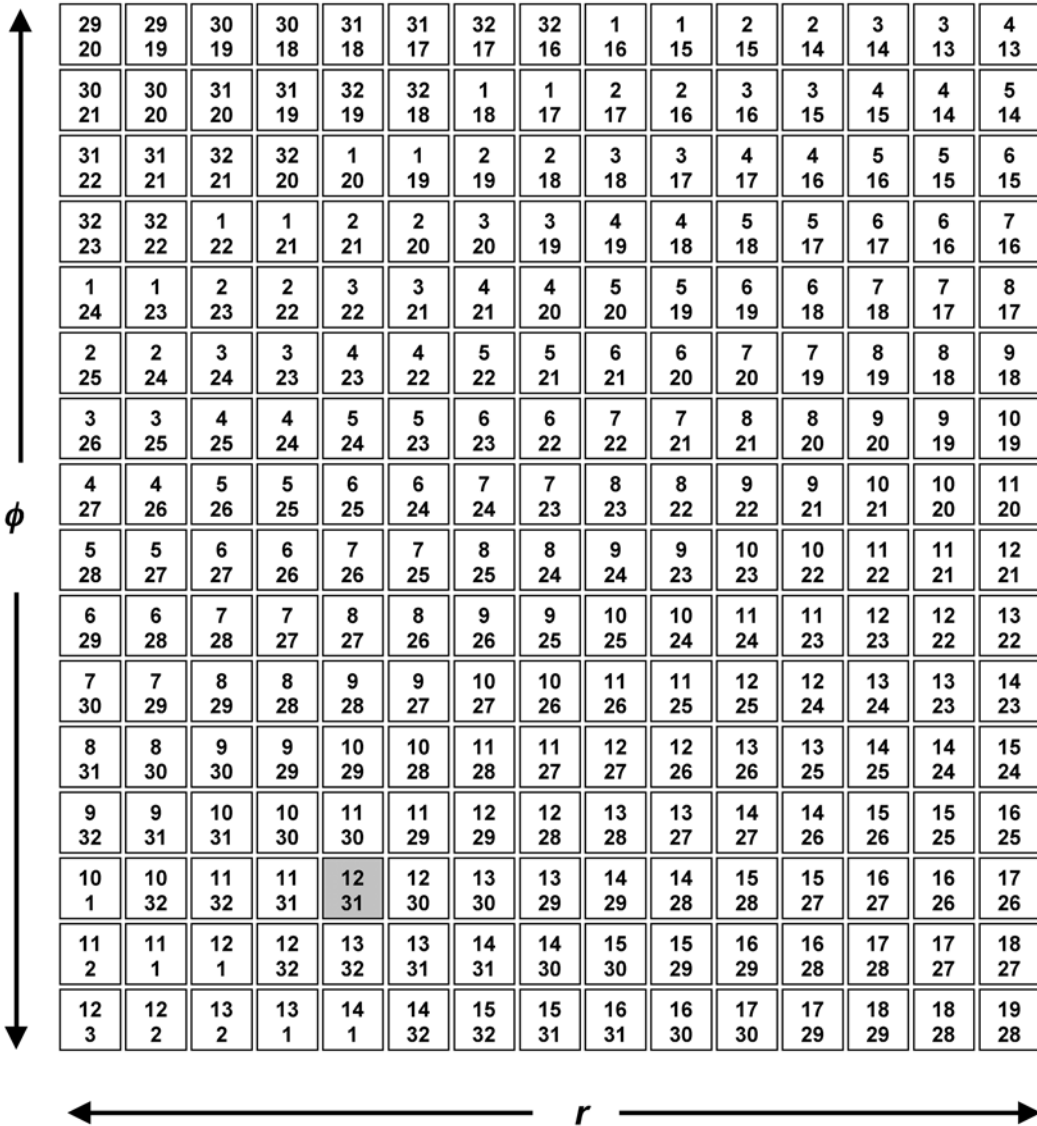


FIGURE 27. Formation of a sinogram based on the 32 detector PET system in Figure 26. The annihilation event shown in Figure 26 would be binned into the highlighted element in the sinogram. Each row in the sinogram corresponds to a projection of the radioactivity within the object at a specific angle ϕ . The angles in a sinogram extend over 180° , the other 180° is redundant as the data are based on pairs of detectors.

closely approximate parallel lines, particularly at the center of the scanner where the patient is located. This small rearrangement of the data from the strict definition results in twice the sampling near the center of the scanner, at the expense of reducing the number of angular samples by a factor of 2. This rearrangement helps produce data that are more appropriately sampled for image reconstruction (Image Reconstruction, p. 70). Figure 28 shows the sinogram that would be obtained from a simple cylindrical object containing two smaller re-




FIGURE 28. A simple object and the sinogram (simulation) that would result from taking projection views over 180° around this object. (Data courtesy of Dr. Andrew Goertzen. Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)

gions with different radioactivity concentrations. Notice how these regions with different radioactivity uptake trace a sinusoidal path in the sinogram.

EXAMPLE 11

What locations in object space will contribute data to a sinogram element at angle 30° and which is offset at a distance 6 cm from the center of the scanner?

ANSWER

From Equation 24, we get:

$$\begin{aligned} 6 &= x \cos 30^\circ + y \sin 30^\circ \\ 0.866 x + 0.5 y &= 6 \\ x &= 6.93 - 0.58y \end{aligned}$$

This is the equation for a line, and all (x,y) locations that satisfy the equation and lie along the line can contribute data to this particular sinogram element.

Two-dimensional data acquisition

In the first generation of multiring PET systems, coincidences were only recorded in *direct* and *cross planes*, where a direct plane is defined as coincidences between detector elements within the same detector ring and a cross plane are the average of coincidences recorded between detectors in two adjacent detector rings (Figure 29A). Collecting coincidences this way allows an improvement in axial sampling because the events collected by the cross planes originate primarily from the volume between the direct planes. Thus, in a system built up of N detector rings, N direct planes and $N - 1$ cross planes can be defined, resulting in a total of $2N - 1$ coincidence planes. Thin tungsten shields, known as *septa*, were used between detector rings to absorb annihilation photons incident at larger angles, thus reducing the overall count rates on the detectors to decrease the




FIGURE 29. Axial section through a multiring PET scanner showing 2-D and 3-D sinogram definitions: A) standard 2-D direct and cross-plane definitions, B) high sensitivity 2-D direct and cross-plane definitions used in many scanners, and C) full 3-D data acquisition.

likelihood of random coincidences and also to help absorb photons that had scattered in the body. This mode of operation is commonly referred to as 2-D *data acquisition* because the data collection is restricted to a set of almost parallel 2-D planes.

Newer, higher resolution PET systems are using smaller detector elements. In these scanners, direct and cross-plane definitions are still used, but because of the low sensitivity of each pure direct and cross plane, additional coincidence planes are accepted in which the detector pairs are separated by up to 5 to 6 detector rings (Figure 29B). This increases the azimuthal angle over which events will be allowed, also known as the *acceptance angle*. In terms of Equation 21, summing these coincidence planes increases the effective value of A , leading to higher sensitivity than would be achieved with the original definition of direct and cross planes. The drawback of this method is that the axial resolution at the edge of the FOV is significantly degraded due to the geometric divergence of the lines of response that are contributing to a particular image plane. There are also shadowing effects that come from the septa.^{40,41} In addition, the axial sensitivity drops off rapidly at the edge of the axial FOV because no additional cross-plane combinations can be added to the planes at the axial extremes.


In 2-D data acquisition, the data from the selected coincidence planes are averaged to produce a set of $2N - 1$ parallel sinograms, each of which can be used to reconstruct a cross-sectional image. Methods for reconstructing these sinograms will be discussed in Image Reconstruction (p. 70).

Three-dimensional data acquisition

The sensitivity of the PET system can be further improved by defining additional coincidence plane combinations, where the ring difference extends well beyond that used in 2-D data acquisition (Figure 29C). This requires the removal of the tungsten septa that would otherwise block these oblique lines of response. To avoid unacceptable resolution losses at off-center positions, these oblique coincidence planes are now stored in separate sinograms, with an associated azimuthal angle. This leads to N^2 sinograms for an N -ring PET scanner. Because the coincidence planes are no longer only limited to parallel planes, this acquisition mode is referred to as *3-D data acquisition*. Since the data are acquired in a 3-D manner, an appropriate 3-D algorithm has to be used to reconstruct the images (Three-Dimensional Analytic Reconstruction, p. 82).

The 3-D acquisition mode provides a dramatic improvement in sensitivity (typically a factor of $\times 5$ to $\times 7$) compared to 2-D acquisition,^{42,43} as A in Equation 21 now corresponds to the entire axial length of the scanner. This additional sensitivity can be used to improve signal-to-noise in PET images, reduce imaging time, or to reduce the amount of radioactivity that is injected. The sensitivity increase is not uniform across the field of view. The sensitivity profile for 3-D data acquisition is triangular in shape in the axial direction as shown in Figure 30. At the axial extremes of the field of view, the sensitivity is equivalent to that of 2-D data acquisition. The wide-open geometry of 3-D acquisition and lack of interplane septa results in a three- to fourfold increase in the fraction of scattered events detected. Furthermore, noise due to the subtraction of random coincidences becomes a problem at lower activity concentrations than in 2-D because each detector sees more of the radioactivity in the body and, therefore, has a higher singles event rate. This leads to further challenges in accurately cor-

FIGURE 30. Measured sensitivity profiles along the axis of a multiring PET scanner in 2-D mode (corresponding to B in Figure 29) and 3-D mode (corresponding to C in Figure 29). In 3-D mode, sensitivity peaks at the center of the axial field of view. At the extreme edges of the axial field of view, 2-D and 3-D sensitivity are equivalent.



recting data for random and scattered coincidences and can result in relatively poor image quality when large amounts of activity are present in the field of view, for example, when imaging over the bladder region. Many radiotracers are excreted through the kidneys and a significant fraction of the injected activity can end up in the bladder.

As a consequence of the improved sensitivity of the 3-D acquisition mode, the size of the raw data sets also is increased. This presents an additional challenge in handling the data from the initial collection of the data to the final archiving. As an example, the complete, uncompressed, 3-D data set from a high-

resolution clinical PET system can be more than 100 Mbytes. With the constant improvement in computing power and data storage, this size may not be a limitation in the near future, but with present technology a data set of this size poses significant challenges. Due to the relatively high-noise levels in the raw data, conventional loss less compression techniques are relatively inefficient in producing significant compression ratios. To produce a more significant reduction in the data sizes, a technique sometimes referred to as angular mashing is used. In the method, adjacent angles (projection angles within each sinogram and/or the oblique azimuthal angle) are simply added together to produce the compression. This is similar to the averaging used in high-sensitivity 2-D data acquisition shown in Figure 29B. Using this technique, the 3-D data sets can be reduced to a more manageable size of about 20 Mbytes. Because projection angles are added together, this method does produce a loss in both in-plane and axial resolution at off-center positions.

EXAMPLE 12

For the same scanner described in Problem 10, estimate the system sensitivity at the center of the field of view if there are 16 rings of BGO detectors and the scanner is operated with 3-D data acquisition.

ANSWER

The detector efficiency and packing fraction are the same as in Problem 10. The geometric efficiency is now given by Equation 21:

$$\Omega = 4 \pi \sin (\tan^{-1} (A/D)) = 4 \pi \sin (\tan^{-1} (6 \times 16 / 800)) = 1.50$$

The sensitivity Equation 23 is therefore:

$$\eta = 100 \times ((0.755)^2 \times 0.673 \times 1.50) / 4 \pi = 4.57\%$$

The result for 2-D acquisition in Problem 10 was 0.29%. This problem shows the dramatic increase in sensitivity at the center of the scanner for 3-D data acquisition compared with 2-D data acquisition.

Data acquisition protocols

The end point in most PET studies is to produce an image, from which diagnostic or quantitative parameters can be derived. These parameters can be as simple as a qualitative comparison of activity concentration in different tissue regions or more complex biologic parameters such as metabolic rate, receptor density, or levels of gene expression. The information that is to be extracted from the image will dictate how the PET data are collected (i.e., static or dynamic sequence).

The most basic data acquisition protocol in PET is the collection of a single data set or static frame over a fixed length of time. The image reconstructed from such a data set represents the average tissue activity concentration during the acquisition. This is the typical acquisition mode used in studies where the tissue activity distribution remains relatively static during the collection of the data. An example where this acquisition mode is commonly used is for 2-deoxy-2-[F-18]fluro-D-glucose (FDG) studies, where the tracer concentration remains fairly stable following an initial uptake period of 30 to 40 minutes. In these types

of studies, the biologic parameter of interest (in this case the metabolic rate for glucose) is then assumed to be directly proportional to the measured activity concentration.

For some radiotracers, it is necessary to follow the dynamic changes in concentration to extract a particular parameter of interest. In these studies, the data are collected as a sequence of dynamic time frames, where the PET images provide information about the changes in activity concentration distribution over time. This information represents the tissue response to the time course of the radiotracer in the plasma following intravenous injection. The tissue time-activity curve can then be processed with a compartmental model to determine the parameters of interest. These types of studies typically also require additional data such as the plasma radioactivity concentration and the plasma concentration of labeled metabolites, which can be determined from blood samples. An example of a dynamic study is shown in Figure 31.

The dynamic acquisition mode also is used in studies where the tissue radiotracer concentration remains constant (e.g., FDG brain scans), but the patient has to remain motionless in the scanner for an extended period of time. By collecting the data in a dynamic sequence (e.g., multiple 5-min frames), it is possible to determine if the subject moved during the acquisition and allow removal, or realignment of frames, when the subject moved.⁴⁴

Most PET systems have a relatively narrow axial field of view that limits the coverage to a single organ (e.g., the brain, heart, kidneys, and so on). To cover a larger extent of the body, the acquisition has to be performed in a series of steps in which the patient is moved through the scanner (Figure 32). To minimize the amount of patient motion, these scans are typically limited to 30 to 40 minutes. The main challenge in these whole-body scans is to collect enough counts in both the emission and transmission (see Attenuation Correction, p. 56) scans in this time frame to produce images of diagnostic quality, without excessive noise levels. The problem of noise contamination from the attenuation correction using short transmission scans has been greatly reduced through the development of fast and accurate image segmentation algorithms.^{45–47} Also, initial results from clinical PET scanners that are based on fast LSO detectors demonstrate the possibility of acquiring whole-body images in as little as 5 to 10 minutes. In many cases, the relatively high statistical noise originating from the emission data requires that the images be reconstructed using iterative algorithms (see Iterative Reconstruction Methods, p. 86). Although these algorithms tend to be computationally expensive, improvements in both acceleration techniques and a constant improvement in computing hardware, now allow routine reconstruction of high-quality, whole-body PET images in clinically acceptable times.^{48,49}

DATA CORRECTION

To produce an image volume in which each voxel value represents the true tissue activity concentration, a number of corrections need to be applied to the raw sinogram data. These corrections are typically applied to the sinograms as a series of multiplicative factors prior to image reconstruction.




FIGURE 31. Dynamic sequence of PET images showing a single transaxial slice through the human brain at different times following the injection of [F18]fluoro-L-DOPA, a radiotracer that reflects dopamine synthesis. Times are in minutes. Notice how the distribution pattern changes over time. The radiotracer is delivered to all brain tissue during the first few minutes. There is a gradual accumulation of the radiotracer in the caudate and putamen (the site to which dopaminergic neurons project) and clearance from other tissues. By 60 minutes, there is clear contrast between the specifically bound radiotracer in the caudate and putamen and the rest of the brain tissues. (Adapted with permission from Cherry SR, Phelps ME. Positron Emission Tomography: Methods and Instrumentation, in *Diagnostic Nuclear Medicine*, 4th Edition, Eds: Sandler MP, Coleman RE, Wackers FIT, Patton JA, Gottschalk A, Hoffer PB, Williams & Wilkins, Baltimore, 2002.)





FIGURE 32. To obtain a whole-body PET image, the body has to be imaged in segments. To achieve this, the patient bed is moved in steps through the scanner. At each discrete bed position, a scan is acquired. The data collected at each bed position are reconstructed and assembled into a whole-body volume which can be reoriented into coronal, sagittal, and transaxial images as shown on the right.

Normalization

Nonuniformities in individual detector efficiencies (physical dimensions), geometrical variations, and detector electronics (e.g., energy thresholds) all contribute to variations in coincidence detection efficiency between different LORs (i.e., pairs of detector elements) in the system. The normalization corrects each individual LOR, with a multiplication factor that compensates for these nonuniformities. Figure 33 illustrates the effect of normalization on the images from a scan of a uniform cylinder, acquired in 2-D mode. This figure illustrates that the normalization is not only a correction on the individual sinogram but is a volumetric correction that adjusts for any sensitivity variation between the different coincidence planes.

To generate a normalization correction, the individual detector efficiencies as well as any geometrical efficiency variations are measured. In addition, any variation in plane-to-plane efficiency is measured. The most straightforward method to determine the normalization correction factors is to collect data from a uniform plane source of activity, positioned at 6 to 8 equally spaced projection angles. This method will directly measure the relative variation in coincidence detection efficiencies between all the LORs in the system.⁵⁰ To avoid dead time and pile-up effects, a source of relatively low activity has to be used. Therefore, the main challenge using this method is to acquire enough counts per LOR (within a reasonable time frame) to provide a good estimate of the efficiencies with a minimum of statistical noise that would propagate into the final image.

FIGURE 33. Effect of normalization on the image of a cylinder containing a uniform concentration of radionuclide. Row A shows a transaxial cross-section through the reconstructed image of the cylinder, row B an axial cross-section. Without the normalization, the nonuniform axial sensitivity is revealed. The effect of the normalization is less visible in the transaxial image, but the difference image reveals ring artefacts as well as a “hole” in the center.



Furthermore, any nonuniformities in the plane source will propagate into the final image.

EXAMPLE 13

Assuming Poisson counting statistics, calculate the number of events that would be needed to normalize a PET scanner with a total of 10^6 lines of response (LORs) if a normalization accuracy of an average of 1% is desired. If the normalization source leads to a coincidence counting rate of 50,000 cps, how long would this normalization scan take?

ANSWER

For Poisson statistics, the standard deviation, σ , is equal to \sqrt{N} , where N is the number of events collected. For each LOR, we require that:

$$\sqrt{N}/N = 0.01$$

therefore, $N = 10,000$

For 10^6 LORs, the total number of counts needed in the normalization scan is:

$$10,000 \times 10^6 = 10^{10}$$

For a counting rate of 50,000 cps, the time for the normalization scan would be:

$$10^{10} \div 50,000 = 200,000 \text{ s} = 55.6 \text{ h}$$

This demonstrates the difficulty of making direct measurements of normalization factors for each individual LOR in the system.

An alternative method to determine the normalization matrix is the component-based method, in which the coincidence detection efficiency of a pair of detectors i and j is assumed to be composed of the product of the individual detector efficiencies, ϵ , and geometrical factors, $g_{i,j}$. The normalization correction factors $n_{i,j}$ are therefore given as:

$$n_{i,j} = \frac{1}{\epsilon_i \times \epsilon_j \times g_{i,j}} \quad (25)$$

These geometrical factors include a correction for the angle of incidence of the annihilation photons, systematic variations in crystal efficiency dependent on the position of the crystal in a detector block modules, and relative plane efficiency. The individual detector efficiencies for all detector elements in the system can be determined from a scan of a uniform cylinder or any circularly symmetric source, centered in the scanner FOV. For each detector element in the system, the sum of the coincidences between the detector of interest and all of its opposing detectors (corrected for random and scattered coincidences) is determined. The assumption is made that by averaging over a large number of opposing detector elements, this sum will be directly proportional to the individual detector efficiency. The geometrical factors are typically determined once for a particular system using very high counting statistical acquisitions at the factory and can be assumed to remain constant. The coincidence detection effi-

ciency for a given LOR in the system is the product of the measured individual detector efficiencies adjusted for any geometrical efficiency variations,⁵¹⁻⁵³ and the normalization factors that need to be applied to the emission data are the reciprocal of this quantity (Equation 25). This method provides an estimate of the normalization factor with virtually no statistical noise because the measurement of detection efficiencies is averaged over a large number of detector elements; however, the method dose depends heavily on a number of measured and empirical geometrical factors that could introduce systematic errors as discussed by Badawi et al.⁵⁴ Figure 34 shows the projection data before and after correction for detector normalization.

Attenuation correction

At 511 keV, a relatively high probability exists that one or both annihilation photons will interact in the subject, predominantly through Compton interactions.

FIGURE 34. A: Unnormalized sinogram. Each detector in the system traces a diagonal line in the sinogram. Detectors that have a high or a low efficiency will lead to bright and dark diagonal lines, respectively. B: Measured normalization factors. C: After multiplying the unnormalized sinogram by the normalization factors, differences in detector efficiencies have been effectively removed. The images on the left show an individual sinogram; the images on the right show normalization effects in the axial direction.





FIGURE 35. Effect of attenuation correction on a uniform cylinder.



As discussed earlier, the results of these interactions are the removal or attenuation of primary photons from a given LOR and the potential detection of scattered photons in a different LOR. Thus, attenuation and scatter are manifestations of the same physical process. Correction involves removing scattered events from the LORs and then subsequently correcting each LOR for the fraction of events that were scattered, or attenuated from, that LOR.

Attenuation of the signal from a given LOR can be corrected either by a direct measurement or using a mathematical model or a combination of the two. Figure 35 illustrates the effect of photon attenuation on the image of a cylinder containing a uniform radioactivity concentration. Without the correction, the central portion of the cylinder appears to have lower activity than the outer edge because photons coming from the center of the cylinder, on the average, must pass through more material to reach the detectors than photons at the edge of the cylinder.

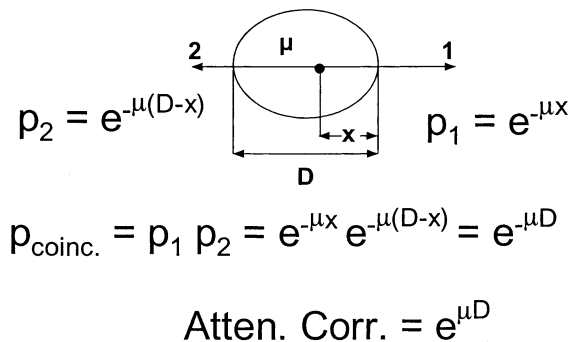
Consider a point source located at an unknown depth x in a uniformly attenuating medium with an attenuation coefficient μ . If the thickness of the object is D along the LOR (Figure 36), then the probability that annihilation photon 1 will escape the object is the result of Equation 15:

$$p_1 = \frac{I(x)}{I(0)} = \exp(-\mu x) \tag{26}$$

The probability that annihilation photon 2 will escape is:

$$p_2 = \frac{I(D-x)}{I(0)} = \exp(-\mu(D-x)) \tag{27}$$

FIGURE 36. Derivation of the equations for attenuation correction in PET. p_1 and p_2 are the probabilities that each of the two annihilation photons from a source located at an unknown depth x in a uniform attenuator of thickness D will escape the object. The product of p_1 and p_2 is the probability that both annihilation photons will escape the object and be available for detection. The attenuation correction for this line of response is simply $(p_1 \times p_2)^{-1}$, which is independent of the location of the source.



The probability that both annihilation photons will escape the object is the product of the individual probabilities:

$$p_1 \times p_2 = \exp(-\mu x) \times \exp(-\mu(D - x)) = \exp(-\mu D) \quad (28)$$

As can be seen from this equation, the reduction in photon flux is independent of the location of the source and only dependent on the total thickness of the object along the LOR and the attenuation coefficient of the object. This is unique to annihilation coincidence detection and makes attenuation correction straightforward. The attenuation correction factors, $a_{i,j}$ that need to be applied to the emission data for the LOR joining detector i and j , are given simply by the reciprocal of Equation 28:

$$a_{i,j} = \exp(\mu D_{i,j}) = I(0)/I(D_{i,j}) \quad (29)$$

where $D_{i,j}$ is the tissue thickness for the LOR between detector i and detector j . This equation forms the basis for the various methods for correcting photon attenuation in PET.

EXAMPLE 14

A line of response passes through 30 cm of soft tissue and 2 cm of bone. Compute the attenuation correction factor that would need to be applied to this line of response.

ANSWER

From Table 3, $\mu(\text{bone}) = 0.17 \text{ cm}^{-1}$ and $\mu(\text{tissue}) = 0.096 \text{ cm}^{-1}$ at 511 keV.

The total attenuation correction factor is the product of the two components:

$$\exp(0.17 \times 2) \times \exp(0.096 \times 30) = 25$$

This demonstrates how to compute attenuation correction factors for a mixture of tissues and the very large correction factors that need to be applied when the annihilation photons have to pass through large amounts of tissue.

Calculated attenuation correction

In its simplest form, the calculated attenuation correction assumes that the outline of the object that being imaged can be approximated with a geometrical shape such as an ellipse. Furthermore, the attenuation coefficient within this object is also assumed to be constant. To generate the attenuation correction, the chord length (representing $D_{i,j}$) is determined for each individual LOR intersecting the ellipse and $a_{i,j}$ is calculated from Equation 29.

This method is primarily useful in imaging phantoms where the attenuation coefficient is typically uniform and the shape of the phantom can most of the time be approximated with a geometrical shape such as a circle or ellipse. This method was also used in early PET systems as the attenuation correction in brain studies. To apply this correction, the emission data were initially reconstructed without correction for attenuation. The users then fitted an ellipse to the outline of the skull on each individual cross-sectional slice. From these ellipses, the attenuation correction was calculated and applied to the raw data, which was reconstructed

again. This method produces images that are largely free of attenuation artifacts, however, due to the approximations of the shape of the object and the assumption of a uniform attenuation coefficient, the method tends to underestimate the attenuation. This method is also very prone to artifacts, such as asymmetries, depending on the positioning of the ellipse. Furthermore, the method can be very labor-intensive on a system generating a large number of image slices.

An improvement of this method was introduced by Bergström et al⁵⁵ and later refined by Siegel and Dahlbom.⁵⁶ The outline of the head was determined directly from the sinogram data. This eliminated the need of fitting the ellipse to each individual slice and more accurately models the outline of the head. Furthermore, in this refined method, the attenuation of the skull can be modeled. This method improves the original one but still has a tendency to create quantitative errors and small artifacts, especially at the base of the skull due to the presence of air cavities and thicker bones.

Measured attenuation correction

The most accurate method to determine the attenuation correction is through direct measurements. As was shown in Equation 28, the amount of attenuation is independent of the location of the source. This means that if a source is placed outside the object along the LOR of interest, the amount of attenuation would be the same as for a source inside the object. Therefore, by placing a positron-emitting source outside the object, the amount of attenuation can be measured directly. In measured attenuation correction, either a set of ring sources, or a set of rotating rod sources, is placed just inside the detectors, enabling the attenuation factors for all LORs in the scanner to be measured in a single scan. Initially, a reference or a *blank scan* is measured, in which data from these external sources are acquired with no object in the scanner. This corresponds to measurements of $I(0)$ in Equation 29 for each LOR. Then the object is placed in the scanner, and a *transmission scan* is acquired. This provides $I(D)$ from Equation 29. The attenuation correction factors for each LOR are simply given by taking the ratio between the blank sinograms and the transmission sinograms as shown in Figure 37. The normalized emission sinogram is multiplied by these factors to obtain attenuation corrected sinograms. Figure 38 shows measured sinogram data before and after applying attenuation correction. The sources used in blank and transmission scans usually are made from ⁶⁸Ge which has a 273-day half-life. These sources generally need replacing every 12 to 18 months.

The main advantage of this method is that the attenuation is directly measured, and no assumptions are made in regards to the shape of the object nor the distribution of attenuation coefficients. The main difficulty in these transmission scans is to collect an adequate number of counts along each attenuated LOR. It is not unusual to have attenuation correction factors of over 50 in the region of the abdomen, with less than 10 measured transmission counts measured in an LOR. Thus, the statistical quality of the transmission scan is typically very poor and this noise will propagate into the emission data if left unprocessed. The most common method to process the transmission data is to apply a spatial smoothing filter to the blank and transmission scans prior to computing the attenuation correction (i.e., the blank/transmission ratio). This method is typically adequate if the transmission scan is fairly long (> 20 min).

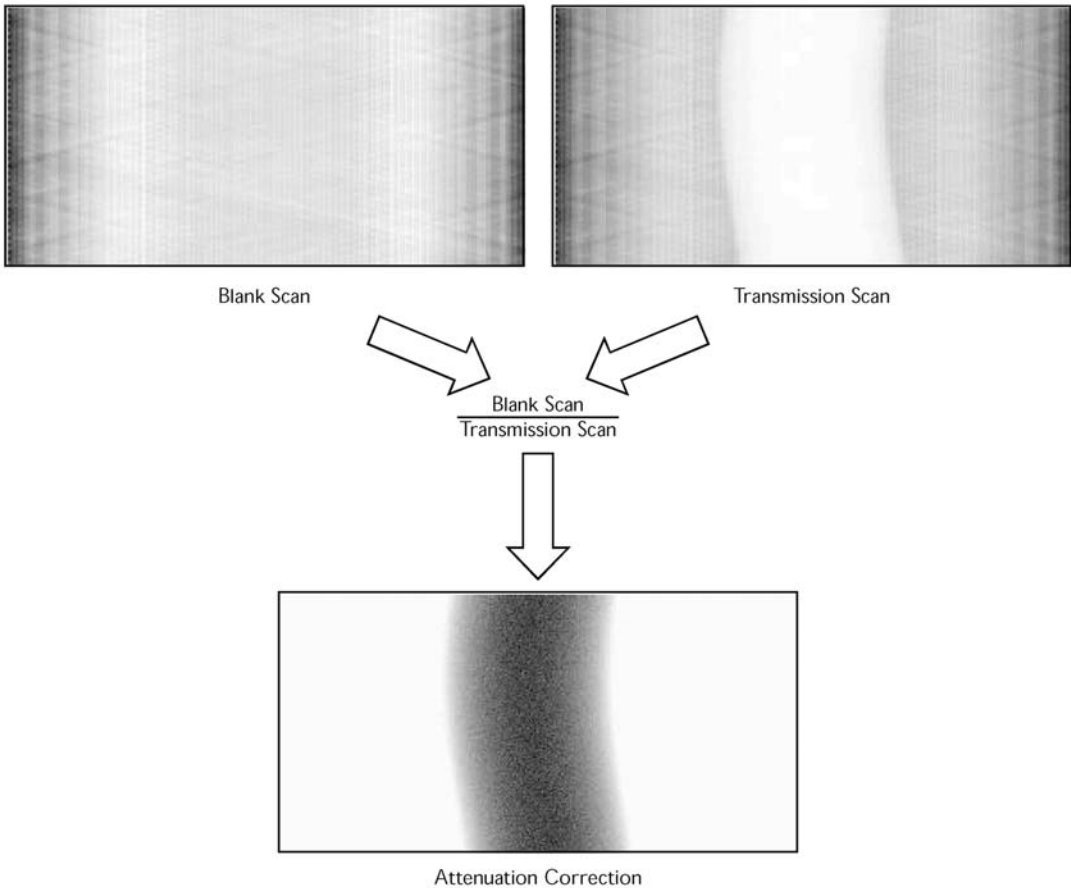


FIGURE 37. The measured attenuation correction matrix is created by dividing the blank scan sinogram (acquired without the subject in the scanner) by the transmission scan sinogram (acquired with the subject in scanner). This operation is performed on every element (i.e., line of response) in the sinogram. These blank and transmission scan measurements typically are taken with external ring or rotating rod sources containing positron-emitting radionuclides.

Hybrid techniques of the calculated and measured attenuation correction methods also have been developed, where a virtually noise-free attenuation correction can be created. In these methods, an image of the attenuation coefficients is first reconstructed by taking the natural log of the attenuation correction sinograms and reconstructing these sinograms using computed tomography (CT) techniques. This reconstruction yields a noisy CT-like image, but the information content of these images is typically enough to allow an image segmentation, where the attenuation coefficients in the images are classified or segmented into a fixed number of attenuation coefficients (e.g., soft tissue, lung tissue, and bone). The segmented image can then be traced along each LOR to calculate an almost noiseless estimate of the attenuation correction.^{45–47} A drawback of the method is that assumptions about the actual attenuation coefficients must be made and population averages are used.

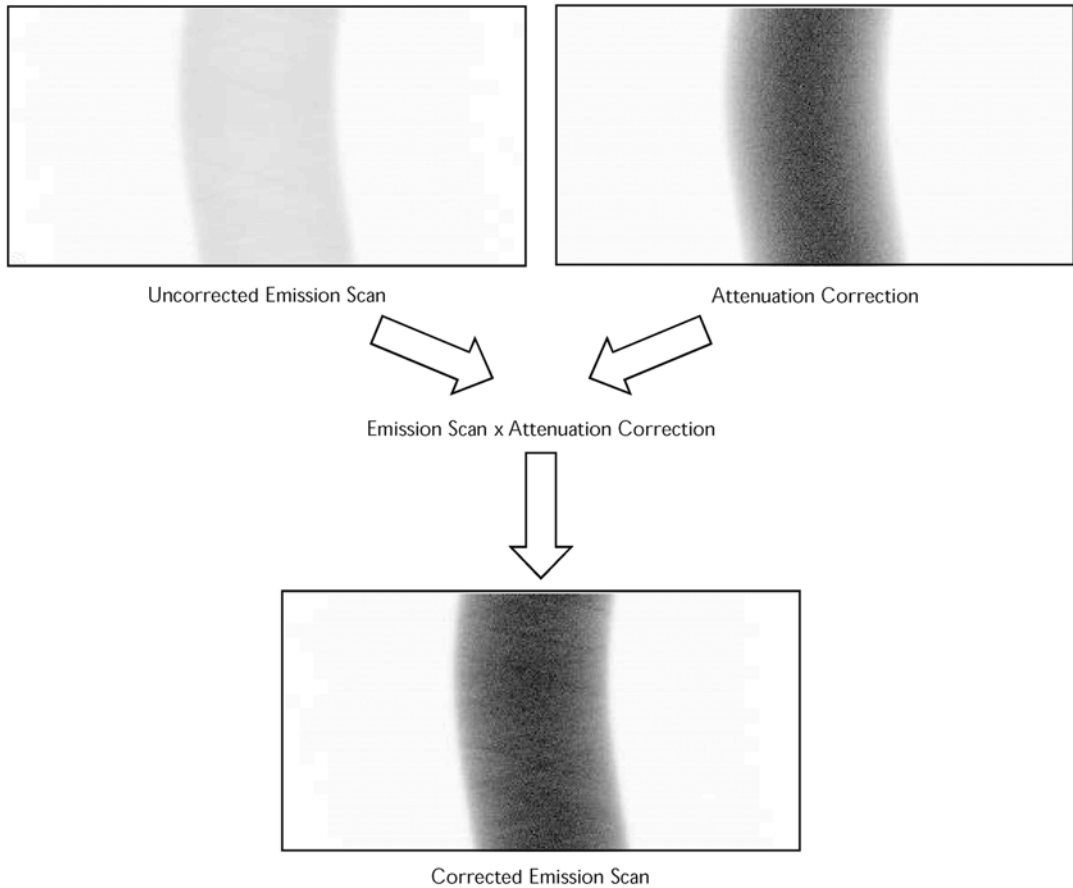


FIGURE 38. The attenuation correction is performed by multiplying the normalized emission sinogram with the attenuation correction matrix. This operation is performed on every element in the sinogram.

The introduction of rotating rod sources instead of ring sources has allowed a more efficient use of the scanner with the acquisition of transmission scans after the injection of the tracer.^{57–59} This approach is known as *post-injection transmission scanning*. In clinical imaging protocols (Data Acquisition Protocols, p. 50), this eliminates the need for the patient to remain in the scanner during the required time for the radiotracer to distribute in the body. Because most transmission sources use a positron emitter, a coincidence originating from the transmission source will be indistinguishable from an emission coincidence. However, when using a rod source, the position of the rod is monitored during the acquisition. If the line of response for a detected event aligns with the location of the rod source, a high likelihood exists that the event originated from an annihilation in the rod source. This assumption can be made if the total activity in the source is much higher than the total emission activity along any given LOR. This is typically the case if the activity in the rod sources is approximately 2 mCi. In areas of high accumulation of radioactivity, such as in the bladder in FDG scans, this assumption may not be true. This can lead to an undercorrec-

tion for attenuation in these regions, due to the substantial number of emission coincidences that are erroneously assigned as transmission events.

Using rotating rod sources, the emission and transmission scans can be acquired simultaneously.^{60,61} However, several practical difficulties occur in performing simultaneous emission and transmission scans. The relatively high activity in the rod sources produces an elevated counting rate of random coincidences (Types of Events, p. 33), which can, in some instances, exceed the emission counting rate. The subtraction of these random coincidences then results in excessive noise levels in the corrected sinograms (Correction for Random Coincidences, p. 65). To overcome the problems associated with high random counting rates in simultaneous emission/transmission scans, the activity in the rod sources has to be reduced by a factor 4–5. This reduction requires that the transmission data must be corrected for the contamination of emission events. Furthermore, the emission counts also have to be corrected for contamination of transmission events. The main advantage of this method is the perfect spatial registration of emission and transmission data. However, the need for the various cross-contamination factors may introduce systematic errors and noise in the estimation of the net true counting rate for each LOR.

The main reason for the poor statistical quality of many transmission scans is that only a limited amount of activity can be put in the transmission sources. Because the transmission source will be very close to the detector ring, the near detector is always exposed to a very high photon flux. Therefore, the amount of activity in the rod source is primarily dictated by the count rate capabilities of the detector system. One approach to overcome this problem is to only measure the photons that pass through the object instead of requiring coincidence detection.⁶² The line connecting the detector element and the position of the transmission source then gives the directional information of the detected photon. This method is generally referred to as *singles transmission scanning* because only one of the two photons is used in the detection process. The main advantage of this method is the improvement in count rate compared to coincidence measurements. A significant drawback of this method is that scattered radiation cannot be rejected by the requirement that the position of the rod source has to align with the coincidence LOR as described above. Therefore, the scatter fraction is typically very high in *singles transmission scans*, which, in turn, results in a significant underestimation of the attenuation correction. The use of rod sources that extend across the entire axial FOV also results in attenuation correction maps with very poor axial resolution that introduce inaccuracies in attenuation coefficients, especially at interfaces of tissues of significantly different attenuation coefficients. To alleviate these problems, single or multiple collimated point sources are typically used instead of a single rod source.^{63,64}

Another drawback of using the singles measurement with ⁶⁸Ge as the transmission source is that postinjection transmission scans cannot be performed because the emission singles cannot be distinguished from the transmission singles. To allow postinjection transmission scans, one can use a transmission source that emits photons of energy different from 511 keV. Several singles transmission systems use ¹³⁷Cs, which emits 662 keV photons.^{63,64} Using a higher photon energy reduces the amount of emission photon cross-contamination of the transmission data. The amount of cross-contamination depends on the energy resolution of the detector system but cannot be entirely eliminated even in

systems using NaI(Tl) detectors, where the energy resolution is as good as 10%. The drawback of using ^{137}Cs as the source is that the transmission measurement is performed at a different energy compared to the emission data (i.e., 511 keV). To overcome this problem, transmission images are typically segmented and the appropriate attenuation values are assigned at 511 keV.^{63–65}

Scatter correction

Correction for scatter is probably the most difficult correction that is required in PET, mainly because a scattered event is indistinguishable from a true event except on the basis of energy. Unfortunately, there is no simple way to measure the number of scattered events, and it is, therefore, important to minimize the detection of these events, accomplished to a certain degree by energy discrimination, collimation, and geometrical considerations. When an annihilation photon undergoes a Compton interaction in the body and scatters, it will lose some of its energy in the process. The direction of the scattered photon follows the Klein–Nishina probability function (Figure 8) and the amount of energy lost in the scattering interaction is given by Equation 12. At 511 keV, forward scatter, in which only a small amount of energy is lost in the interaction, is favored. In fact, 50% of all Compton interactions produce photons with a scatter angle of 60° or less. If PET detectors only accepted events with an energy of 511 keV, all scattered events could be eliminated. But this would require a detector with extremely good energy resolution. Most BGO scintillation detectors used in PET have an energy resolution of about 20% to 25%, which makes it very hard to separate small angle, high-energy scattered photons from the primary 511-keV photons. Even PET systems that use high light output NaI(Tl) scintillation detectors have an energy resolution of at best 10%. An additional difficulty in separating scattered events from primary events is caused by the fact that a significant fraction of the primary 511-keV photons will only deposit a portion of the energy within the detector volume. Although these events are “good” events, they are detected in the same energy range as scattered events. Thus, if the system would only accept events within a narrow energy window at approximately 511 keV, the overall detection efficiency of the system would be very poor (see Sensitivity: Detector and Geometric Efficiencies, p. 42). Therefore, to maintain a reasonable detection efficiency, most PET systems operate with a relatively large energy window between 350 keV to 650 keV, which also results in the detection of a certain amount of scattered photons. Energy discrimination is most efficient in rejecting low energy, large angle scatter. The presence of tungsten interplane septa (Figure 29) in a PET system helps reduce detected scatter to a relatively low level. For brain imaging, the fraction of the detected events that have undergone Compton scatter is in the range of 0.1 to 0.15 for a 2-D scan with the interplane septa in place. If the septa are removed, and the system is operated in 3-D mode, the scatter fraction for brain imaging increases to approximately 0.3 to 0.4.

Although scattered events are spread across the field of view and have a fairly low spatial frequency distribution, their contribution, particularly in 3-D studies, needs to be corrected to produce images with high contrast and acceptable quantification. Figure 39 shows the effects of scatter in images of a phantom. It is important to notice that scatter and attenuation are really one and the same phenomenon. When one of the annihilation photons scatters, the event is removed from its original LOR (attenuation), but because the photon is rarely ac-





FIGURE 39. A: Profiles through the reconstructed images of a cylinder containing a uniform radioactivity concentration showing the contribution of scattered events. The corrected profile is flat, as expected. B: Images corresponding to the profiles; from left to right: Uncorrected image (trues + scatter), corrected image (trues), and scatter image.

tually absorbed in tissue, it may still be detected in a different LOR (scatter). For this reason, the order in which the corrections are applied is important, and scatter correction should be performed prior to attenuation correction. The scatter correction methods can be divided into three main categories: *analytical*, *dual energy window*, and *simulation methods*.

Analytical methods

In the simplest analytical scatter correction methods, the amount of scatter was estimated by fitting a smoothly varying function to the counts appearing outside the object in the sinograms.^{66,67} This method assumes that the scatter distribution varies slowly across the FOV and is relatively independent of both the source distribution and the scattering medium. One of the most widely used analytical methods for scatter correction in 2-D PET is the Bergström et al. method.⁶⁸ This method is based on phantom measurements of the scatter dis-

tribution from line sources placed at different locations inside scattering media of different dimensions. This information is then used to essentially deconvolve the measured data which is blurred by these scatter distributions. The advantage of this method compared with the fitting method is that it considers the source distribution. However, the shape of the scattering distribution is highly dependent on the shape and size of the scattering medium; thus, the phantom measurements are at best an approximation.

Dual energy method

The dual-energy window technique is based on scatter correction techniques developed originally for SPECT. In this method, coincidences are acquired in two energy windows, for example, 400 keV to 600 keV and 250 keV to 400 keV. The idea is that the upper window contains scattered and unscattered photons and the lower window contains only scattered photons. To correct for scatter, some fraction of the low-energy window counts are subtracted from the high-energy window counts. A difficulty in using this method is that in reality both energy windows contain a mixture of both unscattered and scattered events due to the limited energy resolution of the detectors, and that the low-energy window is more likely to contain multiply scattered photons that in general will have a different spatial distribution than singly scattered photons. This method also requires a set of measured calibration constants to account for efficiency differences in the two windows and to determine what fraction of the low-energy window counts should be subtracted. These constants, in turn, depend on object size. This method is typically only used in studies in which the object geometry is well-defined and remains fairly constant between studies (i.e., brain scans).⁶⁹

Simulation methods

The most accurate scatter correction methods developed to date are probably the simulation-based methods. In these methods, the scatter is estimated by first reconstructing the emission data without scatter correction. Using these images as the initial estimate of the source distribution together with an attenuation map (i.e., reconstruction of the transmission images), the scatter is then simulated, either using a simple single scatter model,⁷⁰ an approximate analytical model,⁷¹ or using a Monte Carlo simulation.^{72,73} The advantage of this method is that it takes into account the 3-D distribution of radioactivity and attenuation coefficients in the object that is being imaged. The obvious drawback is that it is more time consuming because it involves additional image reconstructions and computationally expensive simulations, especially if a Monte Carlo simulation is performed.

Correction for random coincidences

Random coincidences, like scatter, result in additional events being recorded in LORs.⁷⁴ These events, because they contain no spatial information (the two annihilation photons come from separate decays), are distributed quite uniformly across the field of view. Without correction, they lead to a loss in image contrast, adversely affect quantification, and can lead to significant image artifacts. Random correction often is performed in real time in modern PET systems and, therefore, is transparent to the user.

Two main approaches correct random or accidental coincidences in PET. As was shown in Equation 18, the random counting rate can be estimated from

the singles counting rate for a given detector pair and coincidence time window. In theory, the number of random events for every detector pair in the scanner can be estimated and subtracted. To implement this singles method, one would need a data acquisition system that can, in addition to recording coincidences, also accurately monitor the singles rate for each detector element. In addition, the coincidence time window needs to be accurately known for each detector pair.

A different approach is to directly measure the accidental coincidences, which can be achieved by adding a parallel coincidence circuit to the one measuring the prompt coincidences. In this second coincidence circuit, the logic pulse from one of the two detectors is delayed in time such that the detector pair cannot produce any true coincidences. Therefore, any coincidences seen in this coincidence circuit can only be caused by accidental coincidences, which is also an estimate of the number of accidental coincidences in the prompt circuit.⁷⁵

It should be noted that in both methods, the correction for randoms is not a correction on an event-by-event basis, because a random event is indistinguishable from a true event for the coincidence circuit. Instead, the correction method provides a statistically separate measurement of the number of random events detected by each detector pair which is then subtracted from the prompt (trues + randoms) coincidence measurement. This subtraction of two measurements leads to an increase in the statistical uncertainty of the true coincidence rate. Assuming that Poisson statistics applies, the net number of true counts (N_{true}) after correction are:

$$N_{true} = N_{prompt} - N_{random} \quad (30)$$

where N_{prompt} and N_{random} are the number of prompt and random counts, respectively. If the delayed coincidence technique is used for correction, the statistical error in the true counts (ΔN_{true}) is given as:

$$\begin{aligned} \Delta N_{true} &= \sqrt{\Delta N_{prompt}^2 + \Delta N_{random}^2} \\ &= \sqrt{N_{prompt} + N_{random}} = \sqrt{N_{true} + 2 \times N_{random}} \end{aligned} \quad (31)$$

As can be seen from Equation 31, the statistical error in the estimate of N_{true} increases with increasing numbers of random events. If the singles method is used for estimating the accidental rates, then it can be shown using Equation 18, and error propagation analysis that the error in the net true coincidence counts is:

$$\Delta N_{true} = \sqrt{N_{true} + N_{random} + 2 \times 4\tau^2 \times N_{single}^3} \quad (32)$$

where N_{single} is the number of single events (no coincidence requirement) recorded by the detectors (assumed to be equal in the two detectors). The third term in Equation 32 is typically much smaller than N_{random} and can be neglected in terms of its contribution to the statistical error in the estimate of N_{true} . Equation 32 does not consider any errors in estimating the coincidence window 2τ , which typically has to be measured.

Comparing Equations 31 and 32, the singles correction method produces a statistically superior estimate of the number of random events, due to the negligible noise contribution from the term containing N_{single} . On the other hand, there will most likely be systematic, additional measurement errors in the de-

termination of the coincidence time window, which can vary due to differences in cable and trace lengths in the electronics, variations in PMT transit times and other factors that are difficult to measure or account for. Although the delayed coincidence method produces a noisier estimate of the number of random events, this method is virtually free of systematic errors because the delayed coincidences are measured in the same circuitry as the prompt coincidences.

EXAMPLE 15

You are performing a 5-minute, whole-body PET study in a patient. The scanner records a prompt coincidence counting rate of 50,000 cps and a randoms rate of 10,000 cps. If there are 10^6 LORs in the scanner, compare the % uncertainty in the number of true coincidences per LOR for the two different randoms correction methods if you assume that the prompt and random coincidences are equally spread among all LORs in the system.

ANSWER

The number of prompt coincidences per LOR is given as:

$$50,000 \text{ cps} \times (5 \times 60) \text{ secs} / 10^6 = 15 \text{ prompts per LOR}$$

The number of random coincidences per LOR is:

$$20,000 \text{ cps} \times (5 \times 60) \text{ secs} / 10^6 = 6 \text{ randoms per LOR}$$

The number of true coincidences per LOR is:

$$15 - 6 = 9 \text{ trues per LOR}$$

From Equation 31, the uncertainty in the number of trues for the delayed method is:

$$\Delta N_{true} = \sqrt{N_{true} + 2 \times N_{random}} = \sqrt{9 + 12} = 4.58$$

The % uncertainty in N_{true} is

$$100 \times \frac{\Delta N_{true}}{N_{true}} = \frac{4.58}{9} = 50.9\%$$

From Equation 32, the uncertainty in the number of trues for the singles method, assuming that the third term is negligible, is:

$$\Delta N_{true} = \sqrt{N_{true} + N_{random}} = \sqrt{9 + 6} = 3.87$$

The % uncertainty in N_{true} is:

$$100 \times \frac{\Delta N_{true}}{N_{true}} = \frac{3.87}{9} = 43\%$$

This demonstrates the advantage of the singles method, particularly when N_{true} is not much greater than, or even less than, N_{random} .

Dead time correction

In an ideal system, the net true count rate of the system should increase linearly with increasing activity in the field of view. However, there are a number of com-




FIGURE 40. Illustration of the variation of true counting rate as a function of activity concentration in a uniform cylinder (solid line). The dashed line indicates the ideal linear response of the true counting rate. At low activity concentrations, the measured count rate follows the ideal linear response. As the activity concentration is increased, the deviation from the ideal response increases due to dead time in the processing electronics.

ponents in the detection chain that will experience some level of dead time as the activity increases. This is illustrated in Figure 40, where the measured true counting rate is shown as a function of the activity concentration in a 20-cm diameter, 20-cm tall cylinder. The ideal true counting rate is also shown (extrapolated from the measured true counting rate at low activity concentrations). The main source of dead time in most PET systems is the processing of each event in the detector front-end electronics. This dead time is mainly dictated by the extent of signal integration necessary for accurate energy discrimination and event positioning in the detector module. The integration time is, in turn, dictated by the scintillation decay time constant of the crystal material. The integration time is typically set at 3 to 4 times the decay time constant (i.e., 900–1200 ns for BGO and 120–160 ns for LSO). The most common way to characterize dead time is to fit the scanner response using either a paralyzable or nonparalyzable dead time model. Mathematically, these models are described by:

$$\text{Paralyzable model} \quad R_{Meas} = R_{True} \times \exp(-R_{True}\tau) \quad (33)$$

$$\text{Nonparalyzable model} \quad R_{Meas} = R_{True} / (1 + R_{True}\tau) \quad (34)$$

where R_{Meas} and R_{True} are the measured (or observed) and true counting rates, respectively, and τ is the dead time constant, which in the case of the detector dead time would be the integration time. Thus, if the integration time can be reduced, for instance, by using a faster scintillator, the amount of detector




FIGURE 41. Measured true counting rate (closed circles) as a function of ideal true counting rate. The solid line shows the fit of the data to a paralyzable dead time model with a dead time constant of $1 \mu\text{s}$. The dashed line shows the corresponding fit for a nonparalyzable, dead time model.

dead time can be reduced. Figure 41 shows the measured counting rate as a function of true counting rate, using the data from Figure 40. Figure 41 also shows the best fit of the two dead time models, described by Equations 33 and 34, to the data. In this particular case, the paralyzable model best characterizes the overall dead time for the system.

Other contributions to dead time in a PET system can come from the coincidence event processing, real-time sorting of data into sinograms, and data transfer.⁷⁶ Correction for dead time typically involves a model of the dead time behavior of the system at different count rate levels and will generally be some combination of paralyzable and nonparalyzable dead time factors contributed by the different processing stages in the system. Input to the overall dead time determination is usually the measured average detector singles rates and coincidence rates.⁷⁷

Another problem at high counting rates is pulse pile-up which may result in resolution losses and artifacts. Pulse pile-up occurs when two photons are absorbed in the same detector module within the integration time of the electronics. Because of the broad energy windows used in PET and the high likelihood of Compton interactions in the detector, it is quite possible for the total energy deposited by the two photons to fall within the energy window of the system. In this case, the two photons are accepted as a single event and the detector electronics will assign the event to a location between the two interaction positions. This leads to both loss of events (the two events are recorded as one) and mispositioning of the event. At very high counting rates, pile-up can become a limiting factor and may degrade resolution and introduce artifacts.⁷⁸

EXAMPLE 16

Determine the true count rate at which 20% of all counts are lost for a BGO ($\tau = 1.2 \mu\text{s}$) and LSO ($\tau = 160 \text{ ns}$) detector. Assume a paralyzable dead time model.

ANSWER

The 20% count rate loss occurs when:

$$R_{Meas} = 0.8 R_{True}$$

Using Equation 33:

$$\begin{aligned} 0.8 R_{True} &= R_{True} \exp(-R_{True}\tau) \\ R_{True} &= -\ln 0.8/\tau \end{aligned}$$

For BGO:

$$R_{True} = -\ln 0.8/1.2\mu\text{s} = 186 \text{ kcps}$$

For LSO:

$$R_{True} = -\ln 0.8/160\text{ns} = 1.39 \text{ Mcps}$$

IMAGE RECONSTRUCTION

The goal of *image reconstruction* is to provide quantitatively accurate *cross-sectional* images of the distribution of positron-emitting radiopharmaceuticals in the object that is being scanned, using the externally detected radiation along with the mathematical algorithms of *computed tomography*. This essentially allows us to see “inside” the body in a completely noninvasive fashion. The reconstruction step is necessary because the raw PET data only defines the location of the emitting atom to within a line across the object (Figure 1). To reconstruct tomographic images also requires that data from the object be adequately sampled. For this reason, PET scanners generally consist of rings of detectors that fully encompass the object to be imaged or sets of opposing detectors that can be rotated about the object. Both geometries allow data to be collected from many different angles around the object.

Initially, we will consider a highly simplified PET scanner that consists of a single ring of individual detectors that can localize incident annihilation photons (Figure 26). A PET scan consists of the detection of a large number of pairs of annihilation photons (typically 10^6 – 10^8). During the course of the PET scan, the total number of counts measured by a particular detector pair will be proportional to the integrated radioactivity along the line joining the two detectors. This data are commonly referred to as *line integral* data. The role of image reconstruction is to convert the line integrals measured at many different angles around the object into a 2-D image that quantitatively reflects the distribution of positron-emitting atoms (and, therefore, the molecule to which it is attached) in a slice through the object parallel to the detector plane.

There are two basic approaches to image reconstruction. One approach is analytic in nature and utilizes the mathematics of computed tomography that relates line integral measurements to the activity distribution in the object. These

algorithms have a variety of names, including filtered backprojection and Fourier reconstruction. The second approach is to use iterative methods that model the data collection process in a PET scanner and attempt, in a series of successive iterations, to find the image that is most consistent (using appropriate criteria) with the measured data. This section provides a basic introduction to the methods of image reconstruction. A discrete formulation is used, as all real PET data acquired are ultimately sampled in discrete bins rather than on a continuous basis. For more details on the mathematics or details on specific reconstruction algorithms, the reader is referred to the list of further reading at the end of the chapter, as well as the individual references made in the text. An excellent summary review on image reconstruction is given by Leahy and Clackdoyle.⁷⁹ Details on analytic reconstructions can be found in the textbook by Kak and Slaney.⁸⁰

Backprojection

A basic algorithm used as part of many reconstruction methods, and an intuitively appealing way to approach image reconstruction, is *linear superposition of backprojections*, often known simply as *backprojection*. First, an image matrix is defined (typically, 128×128 pixels for PET). For a valid line of response (e.g., coincidences between detectors 12 and 31 in Figure 26), a line is drawn between the detectors and through the image matrix. The value added to each pixel that is intersected by the line is given by $N \times w$, where N is the number of counts detected by the detector pair (after all corrections described in Data Correction, p. 51, have been applied) and w is a weighting factor proportional to the pathlength of the line through the pixel. The value is therefore larger if the line passes across the center of the pixel and smaller if the line passes through the corner of the pixel. This is illustrated in Figure 42. In essence, the counts from a detector pair are being projected back along the line from which they originated. This process is repeated for all valid detector pairs in the PET system, the counts from each subsequent detector pair being added to the counts that have been backprojected for all preceding detector pairs, hence the name linear superposition of backprojections.

There are two different methods to implement simple backprojection on a computer. The first approach, *ray-driven backprojection*, is the method described above. The lines of response, or rays, are traced through the 2-D image matrix, and the pathlength through each pixel is calculated. A more common (and efficient) way to perform backprojection, when the projection data are stored as sinograms (Data Representation: The Sinogram, p. 44) is to use the *pixel-driven backprojection* algorithm. For each image pixel (x,y) , and each projection angle ϕ , we calculate the sinogram coordinate r (from Equation 24) that will contribute data to that pixel. In general, the calculated value of r will not fall exactly on one of the sinogram elements, and linear interpolation is used between adjacent elements to calculate the contribution to add to pixel (x,y) . In this algorithm, each image pixel is handled one at a time, and the calculation proceeds as a loop over all the projection angles ϕ for each image pixel. Mathematically, pixel-driven backprojection can be written as:

$$a'(x,y) = \frac{1}{N} \sum_{n=1}^N s(r, \phi_n) = \frac{1}{N} \sum_{n=1}^N s(x \cos \phi_n + y \sin \phi_n, \phi_n) \quad (35)$$

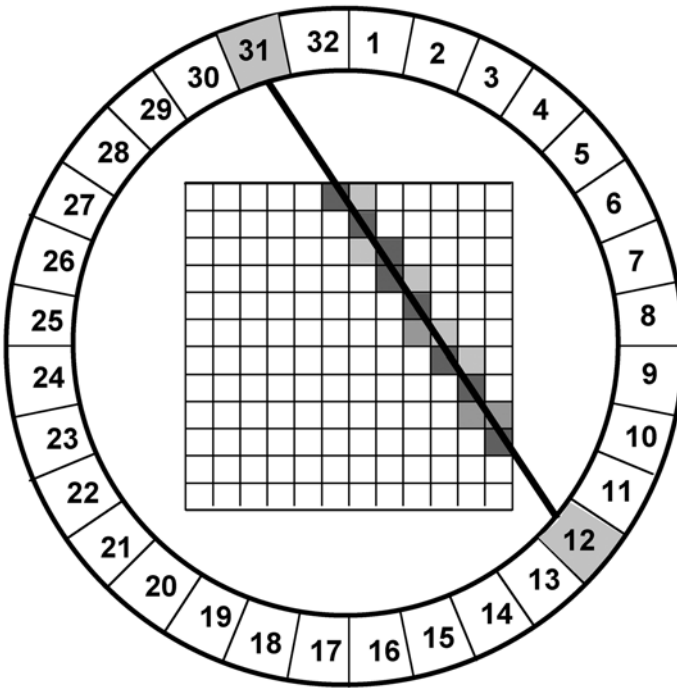


FIGURE 42. Illustration of backprojection. An image matrix consisting of an array of square-image pixels is defined. Events detected by a given detector pair (in this case detectors 12 and 31) are placed in pixels intersected by the line joining the two detectors. A weighting factor is applied to account for the pathlength of the line of response through the pixel. This is repeated for all valid detector pairs to build a backprojected image (see Figure 43).

where N is the number of different equally spaced projection angles over which data have been obtained, $s(r, \phi_n)$ is the number of counts in the sinogram element at angle ϕ_n and radial offset r , and $a'(x, y)$ is the backprojected image. Both approaches to backprojection are mathematically identical, although results can differ slightly depending on the details of how the weighting factors and interpolations are carried out.

Simple backprojection of the data does result in an image that resembles the true distribution of radioactivity in the object, but it is only an approximation, hence the designation $a'(x, y)$. Backprojection places counts outside the boundaries of the object which is clearly incorrect, and for any complex object, it is readily apparent that backprojection will result in a blurred representation of the object because counts are distributed equally along the line from which they originated. The result of reconstruction of the computer phantom shown in Figure 28 by backprojection is shown in Figure 43. The blurring in backprojection is proportional to $1/r$, where r is the distance from the source. Mathematically, it can be shown that the relationship between the backprojected image $a'(x, y)$ and the true activity distribution $a(x, y)$ is given as:


$$a'(x, y) = a(x, y) \otimes \frac{1}{r} \tag{36}$$

where \otimes denotes the operation of convolution.

Analytic reconstruction: projection slice theorem and direct Fourier reconstruction

To develop a reconstruction algorithm that eliminates the $1/r$ blurring factor, we turn to an important theorem known as the *Fourier or projection slice theo-*

FIGURE 43. Simple backprojection reconstruction of the sinogram in Figure 28. The reconstructed image bears some resemblance to the object in Figure 28 from which the sinogram was measured, but the image is badly blurred. Counts are placed outside of the object where there is no radioactivity. (Data courtesy of Dr. Andrew Goertzen. Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, W.B. Saunders, New York, 2003.)



rem (sometimes also known as the *central section theorem*). This theorem states that the measured projection data $s(r, \phi)$ contained within the sinogram can be related to the activity distribution $a(x, y)$ in the object through the use of a widely used mathematical tool known as the *Fourier transform*.

The Fourier transform expresses a function $f(x)$ in terms of its component *spatial frequencies* (expressed as a weighted sum of sine and cosine terms of different frequencies) rather than in terms of its magnitude as a function of position. The component spatial frequencies, ν , and the magnitude with which each frequency contributes to the signal, $F(\nu)$, is determined by the manner with which the signal $f(x)$ changes with position. For example, a signal that is fairly uniform over space can be represented using lower spatial frequencies, while one with sharp discontinuities and regions of rapid change will require higher frequencies for an accurate representation. This is perhaps best illustrated with some simple examples as shown in Figure 44. The two extreme cases would be a uniform signal (i.e., $f(x) = \text{constant}$) and a delta function ($f(x) = 0$, except at $x = 0$, where $f(x) = 1$). In the former case, there is no frequency information and the Fourier transform of the function only contains information at the zero frequency. In the latter case, an infinite number of frequencies are required to represent the sharp delta function, and the Fourier transform contains equal contributions at all frequencies. The last function illustrated is a more general and discrete function, such as might be measured as the projection data from a particular angle in a sinogram. Here again, there is a distribution of frequencies, but high-frequency components are limited by the resolution of the detectors that are being used and the spatial sampling of the data. Discrete Fourier transforms of data can be calculated very quickly on computers; many data analysis packages have built-in functions for computing them. In this text, we will not provide details of the exact formulation and calculation of Fourier transforms. This involves the use of complex exponentials and is beyond the scope of this book. The interested reader is referred to the textbook by Bracewell⁸¹ for further details. We will simply denote the Fourier transform $F(\nu)$ of a function $f(x)$ as:

$$F(\nu) = \text{FT} [f(x)] \quad (37)$$




FIGURE 44. Illustration of the Fourier transform which expresses a spatially varying signal $f(x)$ in terms of its frequency components $F(v)$. Three cases are shown: A (top row): Signal is uniform over space and the only nonzero component of the Fourier transform is at zero frequency. B (middle row): A delta function. All frequencies are required to represent this sharp spike. C (bottom row): A more realistic situation that could represent a profile through the radioactive distribution in a patient. A range of frequencies is required to represent this distribution, with higher amplitudes at lower frequencies. The high-frequency cut-off in real data will be determined by the resolution of the detector and the sampling of the data.

The inverse Fourier transform takes the frequency representation of a function $F(v)$ and computes the spatial representation of the function $f(x)$ and is denoted by:

$$f(x) = \text{FT}^{-1} [F(v)] \tag{38}$$

Fourier transforms can also be calculated for 2-D functions $f(x,y)$, where the transform is also a 2-D function $F(v_x, v_y)$ representing spatial frequency components in the x and y directions.

The projection slice theorem states that the one-dimensional (1-D) Fourier transform of a projection at angle ϕ (one row in the sinogram) is equal to the 2-D Fourier transform of the image evaluated along a radial profile at angle ϕ with respect to the x -axis (Figure 45), which can be written as:

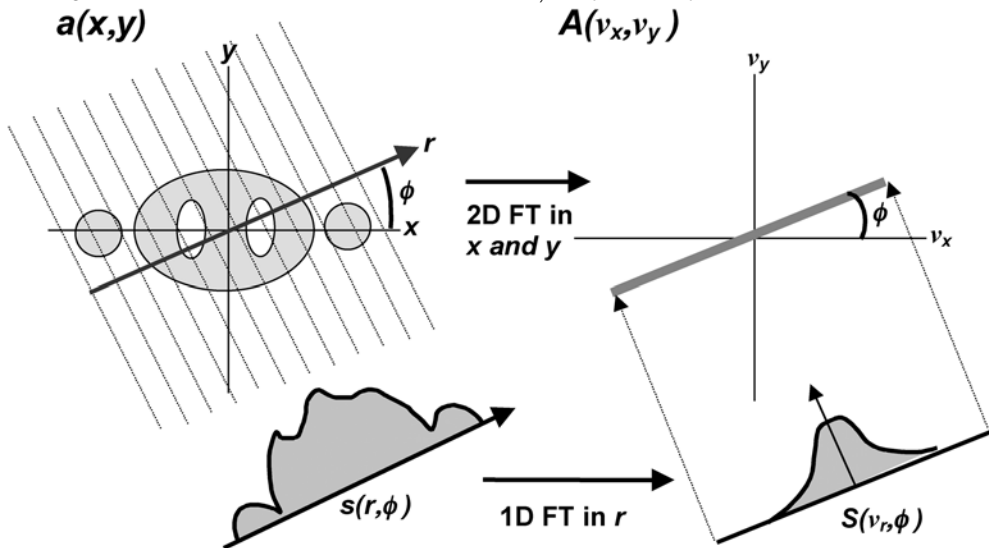
$$S(v_r, \phi) = A(v_x, v_y)|_{v_x = v_r \cos \phi, v_y = v_r \sin \phi} \tag{39}$$

where $S(v_r, \phi)$ is the 1-D Fourier transform of $s(r, \phi)$ with respect to r , and $A(v_x, v_y)$ is the 2-D Fourier transform of the activity distribution $a(x, y)$ which is the quantity desired to reconstruct in the image. This relationship immediately leads to the direct Fourier reconstruction of a cross-sectional image $a(x, y)$ by the following approach:

1. Take the 1-D Fourier transform of the first row in the sinogram.
2. Interpolate and add onto a 2-D rectangular grid $A(v_x, v_y)$ according to Equation 39.
3. Repeat for all subsequent rows in the sinogram.
4. Take the inverse 2-D Fourier transform of $A(v_x, v_y)$ to find the image $a(x, y)$.

The practical difficulty with this approach is that the Fourier-transformed projection data $S(v_r, \phi)$ are sampled along radial lines and must then be interpolated onto a rectangular grid to form $A(v_x, v_y)$ before taking the inverse Fourier transform. Accurate interpolation is computationally intensive, and the resultant image is very sensitive to interpolation errors. However, when properly implemented, this leads to a reconstructed cross-sectional image $a(x, y)$ that will match the distribution of radioactivity in the original object within the limits

FIGURE 45. The Projection slice theorem. The object with activity $a(x, y)$ is scanned to produce projection data $s(r, \phi)$ at different angles ϕ (sinogram data). The 1-D Fourier transform (FT) of the projection data at angle ϕ , $S(v_r, \phi)$ is equal to the 2-D Fourier transform of the image $A(v_x, v_y)$ evaluated at angle ϕ . To reconstruct an image, the Fourier transforms of the projection data at each angle are inserted onto a rectangular grid according to this theorem. The reconstructed image can then be obtained by taking the inverse 2-D Fourier transform of $A(v_x, v_y)$ to yield $a(x, y)$.



imposed by statistical noise in the measured projection data, errors in quantification of the line integrals (due to any of the factors discussed in Data Correction, p. 51), and resolution and sampling limitations of the data.

Two-dimensional analytic reconstruction: filtered backprojection

A more elegant approach to reconstruction can be achieved by reformulating Equation 39 in the spatial rather than frequency domain. The result is:

$$a(x,y) = \frac{1}{N} \sum_{n=1}^N s^*(r,\phi_n) \quad (40)$$

where s^* is the original sinogram data that have been modified by a filter function in the spatial frequency domain with form $H(v) = |v|$:

$$s^*(r,\phi) = \frac{1}{2\pi} \text{FT}^{-1}[S(v_r,\phi) \times |v_r|] \quad (41)$$

$S(v_r,\phi)$ is the 1-D Fourier transform of the original projection data $s(r,\phi)$ with respect to r . Comparing Equation 40 with Equation 35, it can be seen to be the same equation as for backprojection, the only difference being that the projections have been modified by a filter according to Equation 41. Hence, the name *filtered backprojection*.

The reconstruction filter, $H(v)$, is known as the ramp filter because of its shape in the frequency domain (Figure 46) which results in a larger weighting factor for higher spatial frequencies. The blurring that occurs in images reconstructed with backprojection can be thought of as suppressing high-spatial frequency information (high-spatial frequencies give “sharpness” to an image). The shape of the ramp filter, therefore, makes sense intuitively, as it amplifies the high-spatial frequencies with respect to the low-spatial frequencies, reversing the effects of the $1/r$ blurring.




FIGURE 46. The reconstruction filter that is applied to the projections in frequency space to achieve accurate reconstructions is a simple ramp function, where the amplitude of the filter is proportional to the frequency, up to a maximum frequency v_{\max} which is defined by the sampling of the data.

The implementation of filtered backprojection (FBP) is as follows:

1. Take the 1-D Fourier transform of the first projection angle in the sinogram.
2. Multiply this by the filter function.
3. Take the inverse Fourier transform.
4. Backproject the modified (filtered) projection.
5. Repeat for all angles around the object.

This algorithm is very fast and easy to implement on a computer. It involves only 1-D Fourier transforms and simple linear interpolation (in general, the selection of x , y , and ϕ will result in values of r that do not fall exactly on the line of response defined by a detector pair). It has become the method of choice for analytic reconstruction.


In the discrete formulation of image reconstruction appropriate for PET (the recorded projection data are discrete, not continuous functions), the reconstruction filter $H(\nu)$ should be cut-off at a maximum frequency ν_{max} (Figure 46) which is given as:

$$\nu_{max} = \frac{1}{2 \times \Delta r} \quad (42)$$

where Δr is the distance between samples in the sinogram and ν_{max} represents the highest frequency that can be faithfully recorded in the discrete, sampled data.

The ramp filter often is modified in practice to improve signal-to-noise in reconstructed PET images. A typical projection view contains a range of frequencies, with a tendency for higher amplitudes at low frequencies and lower amplitudes at higher frequencies. Statistical noise (noise related to the finite number of annihilation photon pairs contributing to each element in the profile) has a uniform spectral appearance and contributes equally at all frequencies (Figure 47). Thus, if a reconstructed image is too noisy due to limited sta-

FIGURE 47. Fourier transform of the signal (no-noise) (gray line) and the statistical noise (black line) in a projection measurement. Notice how the signal drops off quickly with increasing frequency, while the noise remains quite constant across all frequencies. The individual data points have been smoothed to produce the solid curves to better visualize the trends in the two components. (Data courtesy of Dr. Andrew Goertzen.)



tistics, one option is to attenuate or remove higher frequencies, thus improving signal-to-noise. This improvement, however, is at the expense of degrading image resolution, as the signal contained within these same high frequencies is responsible for the finer detail (rapidly changing activity, sharp edges, and so on) in an image. Thus, by modifying the reconstruction filter, it is possible to trade-off signal-to-noise and spatial resolution in the reconstructed image. A simple way to do this is to cut off the ramp filter at a frequency $\nu_{\text{cutoff}} < \nu_{\text{max}}$. More commonly, apodizing filters are used, which are based on the ramp filter at low frequencies but have a reduced magnitude at high frequencies. These also avoid a sharp change in the filter at the cut-off frequency which reduces the chances of introducing artifacts into the reconstructed images. The functional forms for some of the more common reconstruction filters are:


$$\text{Ramp:} \quad H(\nu) = |\nu| \quad (43A)$$

$$\text{Hann:} \quad H(\nu) = 0.5|\nu| \left(1 + \cos \frac{\pi\nu}{\nu_{\text{cut-off}}} \right) \quad (43B)$$

$$\text{Shepp-Logan:} \quad H(\nu) = \frac{2\nu_{\text{cut-off}}}{\pi} \sin \frac{|\nu|\pi}{2\nu_{\text{cut-off}}} \quad (43C)$$

In all cases, $H(\nu)$ is set to zero for $|\nu| > \nu_{\text{cut-off}}$; $\nu_{\text{cut-off}}$ can have a maximum value of ν_{max} . These filters are shown in Figure 48. Figure 49 shows the reconstruction of a dataset using several different reconstruction filters, demonstrating the trade-off between signal-to-noise and resolution as the filter is changed and as $\nu_{\text{cut-off}}$ is varied. In practice, it is rare that sufficient annihilation photon pairs are collected in a human PET study to reconstruct PET images with a ramp fil-

FIGURE 48. Three reconstruction filters that are commonly used in filtered backprojection. The Shepp-Logan and Hann filters reduce the amplitude at high frequencies, improving signal-to-noise but reducing spatial resolution. They also avoid “ringing” artifacts from the very sharp cut-off of the ramp filter at $\nu = \nu_{\text{cut-off}}$.






FIGURE 49. Effects of filters and cut-off frequencies on the same reconstructed image. Notice trade-off between signal-to-noise and spatial resolution as the filter is changed and its cut-off value is varied. (Data courtesy of Dr. Andrew Goertzen.)

ter cut-off at ν_{max} ; therefore, these modified filters are almost always used to yield acceptable signal-to-noise images.

So far, this chapter has discussed producing a single 2-D image slice from a single ring of detectors. Most PET scanners, however, consist of multiple rings of detectors or area detectors that can provide information from many different locations along the axis of the patient. These scanners can acquire and reconstruct many parallel and contiguous 2-D image slices simultaneously and can be stacked to form a 3-D volume of image data. The thickness of each slice is determined by the size of the detector elements in the axial direction in the case of block-style detectors—the slice profile is a triangular shape, with the width of the base of the triangle equal to the width of the detector element. For gamma-camera style detectors, the slice thickness is determined by the resolution of the detector in the axial direction. Modern PET cameras produce 3-D image volumes that can be resliced from the transaxial orientation in which they were acquired, into coronal or sagittal views, or any arbitrary orientation that is desired. It is important to notice that even though 3-D image volumes are assembled in this manner, they are produced from a stack of independent 2-D reconstructions. This approach only uses annihilation photons that are emitted parallel, or close to parallel, to the desired image section (see Two-Dimensional Data Acquisition, p. 46). It is also possible to use annihilation photons that are emitted obliquely to the desired image slices, but this requires approximations or a fully 3-D reconstruction algorithm (Three-Dimensional Data Acqui-

sition, p. 48). The reconstruction of 3-D PET data is covered in *Three-Dimensional Analytic Reconstruction*, p. 82.

Limitations of filtered backprojection

Although 2-D filtered backprojection (FBP) is a fast method for reconstructing cross-sectional images, it has requirements that must be met to successfully reconstruct images. FBP also has assumptions about the data that are approximations in the real world. One major requirement for successful reconstruction using FBP is adequate sampling.⁸² Two types of sampling are of concern. The center-to-center spacing between samples in a row in the sinogram (sometimes referred to as linear sampling or projection sampling) and the angular sampling that is defined by the number of rows in the sinogram. Based on the sampling theorem, the data should be sampled with a projection sampling interval Δr that is at least one half the highest anticipated spatial resolution of the reconstructed image. The highest resolution that can be achieved is dictated by the intrinsic resolution of the detectors combined with additional factors such as the blurring effects of positron range and noncolinearity (R_{sys} from Example 9). The projection sampling criterion can therefore be expressed approximately as:

$$\Delta r \leq 0.5 \times R_{\text{sys}}(\text{FWHM}) \quad (44)$$

The required angular sampling depends on the diameter of the object being imaged because the sampling density decreases with distance from the center of the scanner. Reconstructing PET data requires projection data that are acquired over 180° . To achieve a sampling interval equal or better than Δr along the circumference of an object of diameter D over 180° requires that:

$$N \geq \pi D / 2\Delta r \geq \pi D / R_{\text{sys}}(\text{FWHM}) \quad (45)$$

where N is the number of angular samples. Inadequate linear or angular sampling leads to a reduction in spatial resolution and can cause artifacts in the reconstructed images (Figure 50).

In PET systems with continuous detectors, the choice of linear and angular sampling intervals is made by the choice of bin size into which the measured interaction coordinates are histogrammed into a sinogram. It is, therefore, quite easy to meet the sampling requirements. In a stationary ring PET system composed of individual detectors, the sampling distance and number of angular samples is fixed by the width of the detector elements and the number of detectors in the ring. In these systems the data are undersampled in the linear direction and, therefore, the projection data usually are rearranged by interleaving adjacent angular views to form a sinogram that has double the linear sampling and half the angular sampling (see *Data Representation—The Sinogram*, p. 50). This leads to a better balance between linear and angular sampling, although the data are still undersampled based on strict sampling criteria. To address this, some PET systems have a small built-in detector motion (often known as “wobble”) that allows Δr and N to be increased so that they meet the sampling criteria.⁸² Some PET scanners have inactive or “dead” regions between detector modules or panels, leading to gaps in the sinogram data. In these regions the data are not sampled at all, which can lead to major reconstruction artifacts. This is often solved by rotating the detectors such that these gaps can be filled in or by using extrapolation or other techniques to estimate what the missing data should be.




FIGURE 50. Effects of linear (A) and angular (B) sampling on reconstructed image quality. Data at top left have sufficient sampling. (Data courtesy of Dr. Andrew Goertzen. Reproduced with permission from Cherry SR, Sorenson JA, Phelps ME. Physics in Nuclear Medicine, W.B. Saunders, New York, 2003.)

EXAMPLE 17

A fixed ring PET scanner is made up of 512 detector elements, each 4.9-mm wide, that are tightly packed in a 79.8-cm diameter detector ring. The spatial resolution of each detector is 3.8 mm. Calculate the required projection and angular sampling for this system for imaging an object that is 40 cm in diameter and compare this with the actual sampling intervals.

ANSWER

From Equation 44, the required projection sampling for 3.8-mm resolution detectors would be:

$$\Delta r \leq 0.5 \times R_{\text{sys}}(\text{FWHM}) \leq 1.9 \text{ mm}$$

From Equation 45, the required number of angular samples over 180° would be:

$$N \geq \pi D / 2\Delta r \geq \pi D / R_{\text{sys}}(\text{FWHM}) \geq \pi \times 400 \text{ mm} / 3.8 \text{ mm} \geq 330$$

Based on how the sinograms for PET are defined (Figure 27), Δr is equal to one half the detector spacing; in this case $\Delta r = 4.90/2 = 2.45$. The number of angles in the sinogram equals half the number of detectors, and so in this case $N = 256$. This demonstrates that PET data are slightly undersampled, although in practice the effects of this level of undersampling on the reconstructed images are relatively small.

The major drawbacks of FBP relate to assumptions that this analytic approach makes about the detection system and the data that are collected. First of all, FBP is based on line integrals and assumes that the PET detectors are point detectors. In practice, PET detectors have finite dimensions, thickness, and resolution; they measure data along a volume joining two detectors. The solid angle subtended by the detectors to points within the volume varies with source position, giving nonuniform sensitivity. Furthermore, FBP cannot model any of the other degrading factors that occur in a PET scanner, such as intercrystal scatter, positron range, and noncolinearity. Lastly, FBP takes no account of the statistical properties of the data. It assumes noise-free data and weights all lines of response equally, independent of signal-to-noise. Because different angular views can never be completely consistent with each other due to the variation of statistical noise from one sinogram element to the next, streak artifacts are common unless very high count data are available for reconstruction. As this is rarely the case in PET, compensation is usually made by using a reconstruction filter that improves signal-to-noise, at the expense of degrading image resolution as explained in *Two-Dimensional Analytic Reconstruction: Filtered Backprojection* p. 76.


Three-dimensional analytic reconstruction

PET scanners operating in 2-D mode only process data that come from detectors lying in or close to the plane of the desired transverse image. Obliquely angled lines of response are rejected (either physically by the interplane septa or by not enabling these lines of response in the data acquisition system) as dis-

cussed in Two-Dimensional Data Acquisition p. 46. This does not make good use of the emitted radiation from a patient, in which annihilation photon pairs are randomly oriented and, therefore, very few of which are aligned parallel to the imaging planes. 3-D PET data acquisition (Figure 29C) involves removing the interplane septa and acquiring data between all possible axial locations in the scanner. This makes better use of the emitted radiation, leading typically to the detection of somewhere between $\times 5$ and $\times 10$ more events for a given radiation dose and imaging time.^{42,43} 3-D PET data also are usually stored in sinograms, with each sinogram being characterized by an average axial location \bar{z} , and an azimuthal angle θ or ring difference δ (Figure 51). A scanner where the data are binned into 32 different intervals in the axial direction would, for example, produce 32^2 possible sinograms, with ring differences up to ± 31 . To incorporate the oblique lines of response into the reconstructed image requires that the approach described in Two-Dimensional Analytical Reconstruction: Filtered Backprojection (p. 76) be modified.

Several approximate reconstruction methods seek to convert the collected 3-D data into a set of parallel transverse sinograms ($\theta = 0^\circ$) so they can be reconstructed using conventional 2-D filtered backprojection methods. The simplest of these methods, often referred to as *single-slice rebinning*,⁸³ takes the average axial location, \bar{z} , of a detected event and places the event in the sinogram(s) that most closely corresponds to that axial position as illustrated in Figure 52. Along the central axis of the scanner, this approximation works perfectly. However, it steadily becomes worse with increasing radial distance. This approach only yields reasonably acceptable images when the object being imaged takes up a small fraction of the field of view and when the axial acceptance angle (the maximum oblique angle accepted) is small. Otherwise, significant blurring of data occurs in the axial direction that becomes apparent when the data are resliced into sagittal or coronal views. It is a very quick way to reconstruct 3-D data, as it only requires that the data are resorted into conventional 2-D sinograms and reconstructed using standard 2-D FBP.

FIGURE 51. Definition of sinograms in a 3-D PET system. A sinogram is produced for each possible combination of axial positions in the scanner, which are defined by the azimuthal angle θ (or ring difference δ) and the average location, \bar{z} , from which the data come. The sinogram data are therefore denoted as $s(r, \phi, \theta, \bar{z})$ where a separate sinogram is produced for each value of (θ, \bar{z}) viewed by the scanner.






FIGURE 52. Illustration of single-slice rebinning in which the 3-D dataset is reduced to a set of parallel 2-D sinograms that can be reconstructed with 2-D FBP. All four lines of response illustrated would be binned into the same sinogram, as they all have the same value for z . These sinogram data are then assumed to come from radioactivity in the gray slice perpendicular to the scanner axis. At the center of the field of view, data are positioned correctly, but it is clear that the axial resolution will degrade significantly for sources away from the scanner axis, as data are averaged from the divergent lines of response.


More recently, a technique known as *Fourier rebinning* (FORE) has been introduced.⁸⁴ The details of FORE are beyond the scope of this text, but it is based on a principle that relates the 2-D Fourier transform of the oblique sinograms ($\theta > 0^\circ$) to the 2-D Fourier transform of the transverse ($\theta = 0^\circ$) sinograms. Although this is still an approximate method, it yields substantially better results than single-slice rebinning, even for large objects and large acceptance angles,⁸⁵ while retaining much of the advantage in terms of short reconstruction times. This has become the algorithm of choice for very large 3-D datasets, for example, those from dynamic PET studies involving perhaps 20 to 30 frames of 3-D data.

The gold standard for analytic 3-D image reconstruction is the *3-D reprojection algorithm* (3DRP).⁸⁶ It is based on an extension of filtered backprojection methods to 3-D. However, two major differences exist between 2-D and 3-D PET datasets, one of which creates a problem, and the second of which provides a solution to that problem. The first difference is that the 3-D dataset is in some senses incomplete. In 2-D a complete set of angular data is obtained by having a ring of detectors or by rotating detectors around the object. The projection dataset is complete because projections are available for ϕ ranging from 0 to 180° . In 3-D PET, an analogous situation would occur if the scanner had the geometry of a sphere. In this case, projections would be available for both ϕ and θ ranging from 0 to 180° , and all possible projection angles would be measured. However, PET scanners generally have a cylindrical geometry and the limited axial length of the scanner results in projections that are truncated in the axial direction (Figure 53). As the azimuthal angle θ is increased, the truncation becomes more severe. An accurate reconstruction requires that this truncation be removed for the azimuthal angles that are to be included in the reconstruction. As it stands, the only angle for which the data are not truncated are the sinograms with $\theta = 0^\circ$ which correspond to the standard 2-D dataset.

The solution for dealing with this comes from the second feature of 3-D PET datasets—data redundancy. We have already discussed how the stack of parallel 2-D sinograms can be reconstructed and stacked into a 3-D image volume. Therefore, sufficient information is contained within the set of transverse sinograms alone to reconstruct the image volume. If the measured 2-D sinogram data were noise-free, there would be no need for 3-D data acquisition and reconstruction. However, statistical noise dominates in PET; hence, there is good reason for wanting to incorporate the data from the oblique sinograms to improve the signal-to-noise of the reconstructed image. Data redundancy provides the solution to the data truncation problem as follows. The conventional 2-D sinograms ($\theta = 0^\circ$) are first extracted from the 3-D dataset. Each is reconstructed with 2-D filtered backprojection and stacked to form a 3-D image volume. This represents a low statistical estimate of the image. This image volume can now be used to estimate the missing data and remove the truncation by adding the activity along oblique lines of response that were not actually measured by the scanner. This process is known as *reprojection* or *forward-projection* and is the inverse process of back-projection. In this way, the missing data are estimated, and the dataset now fulfills the requirements for reconstruction by 3-D filtered backprojection.

Reconstruction then proceeds along the lines of 2-D filtered backprojection, except the dimensionality of the data is increased by one. Each projection, or

FIGURE 53. Truncation of projections in the axial direction caused by the limited axial extent of the scanner. For $\theta \neq 0$ A (top), parts of the object (shaded area) are not sampled. These missing lines of response are indicated by the dashed lines. As θ increases, the amount of truncation becomes worse. However, the dashed lines of response can be estimated by reconstructing an initial image volume (B, bottom) using 2-D FBP of parallel sinograms ($\theta = 0$) and then summing the activity along these lines to estimate what would have been measured in the missing sinograms.



angular view, is now 2-D and represents the projection of activity in (r,z) for a particular value of ϕ and θ . The reconstruction filter is also 2-D and takes on a form dependent on the maximum azimuthal angle, θ_{\max} , of the data used in the reconstruction.⁸⁷ Finally, the filtered data are backprojected through a 3-D image volume. The steps involved in 3-D filtered backprojection using the 3DRP can be summarized as:

1. Extract 2-D sinograms ($\theta = 0^\circ$).
2. Reconstruct each of these with 2-D FBP and stack to form 3-D image volume.
3. Forward project through 3-D image volume to calculate missing LORs to remove truncation.
4. Extract 2-D projection data for angle ϕ and θ .
5. Take 2-D Fourier transform of the projection.
6. Multiply by 2-D reconstruction filter.
7. Take inverse 2-D Fourier transform.
8. Backproject data through 3-D image matrix.
9. Repeat for all angles $0 \leq \phi < 180^\circ$ and $-\theta_{\max} \leq \theta \leq \theta_{\max}$.

Because of the need to compute the missing data, and the fact that backprojection occurs through a 3-D volume rather than across a 2-D plane, the computational complexity of 3DRP is approximately an order of magnitude higher than 2-D FBP. Nonetheless, it enables the oblique sinogram data to be accurately incorporated into the reconstructed image leading to significant improvements in signal-to-noise. In many cases, this signal-to-noise improvement allows a higher cut-off frequency to be used in the reconstruction filter, leading to improved spatial resolution if desired. Figure 54 shows a comparison of 2-D and 3-D data acquisition and reconstruction demonstrating the signal-to-noise gain in 3-D PET. A detailed review of 3-D reconstruction methods for PET, which also presents a more detailed mathematical treatment, can be found in Bendriem and Townsend.⁸⁸

Iterative reconstruction methods

The analytic techniques described above have historically been the most commonly used reconstruction methods for PET. Another class of reconstruction techniques, known collectively as *iterative reconstruction* methods, offer an alternative approach. These methods are computationally more intensive than FBP; for this reason, they have been found to have less clinical use to date. However, as computer speed continues to improve, and with acceleration techniques, these approaches are beginning to be of more general use.

The basic idea behind iterative reconstruction approaches is summarized in Figure 55. An initial guess is made of the image distribution $a^*(x,y)$ (often a blank or uniform grayscale image). The next step is to calculate what projection data would be measured for the radioactivity distribution in the initial guess. The simplest way to do this is to use a process known as *forward-projection*. This is exactly the inverse of backprojection, and involves summing up all the activity in pixels that are intersected by the line of response that corresponds to the measured sinogram element. Once this process is complete, we have a set of estimated projection data based on our initial guess that can be compared with the actual measured projection data. Obviously, they will not agree because it is




FIGURE 54. Comparison of PET images of the brain acquired and reconstructed in A) 2-D (reconstructed with standard 2-D FBP) and B) 3-D (reconstructed with 3DRP). Injected dose and imaging time were the same for both studies. Notice the improvement in signal-to-noise in the 3-D study. (Reproduced from *Brain Mapping: The Methods*, 2nd ed, Eds: Toga AW, Mazziotta JC, Academic Press, San Diego, 2002, with permission from Elsevier.)

very unlikely that the initial guess of $a^*(x,y)$ is anything like the true radioactivity distribution $a(x,y)$. Based on the differences between observed and measured projections, the initial guess is then adjusted, and the whole process is repeated. If the method by which the image estimate is updated is properly formulated, then with each successive iteration through this process, the image estimate will start to converge towards the true image. After a while, the estimated image should closely match (within the limits imposed by the statistical quality of the data and the resolution and sampling of the detector system) the true distribution of radioactivity in the object.

There are many different types of iterative algorithms each differing in some aspect of their formulation and implementation. One of the factors that distinguish these algorithms is the *cost functions* they use. The cost function is a function that gives a measure of the difference (or similarity) between the estimated and measured projections and is the function that we seek to minimize (maximize) during the reconstruction. The second important component of the algorithm is the search or update function, that is, how the image estimate is up-

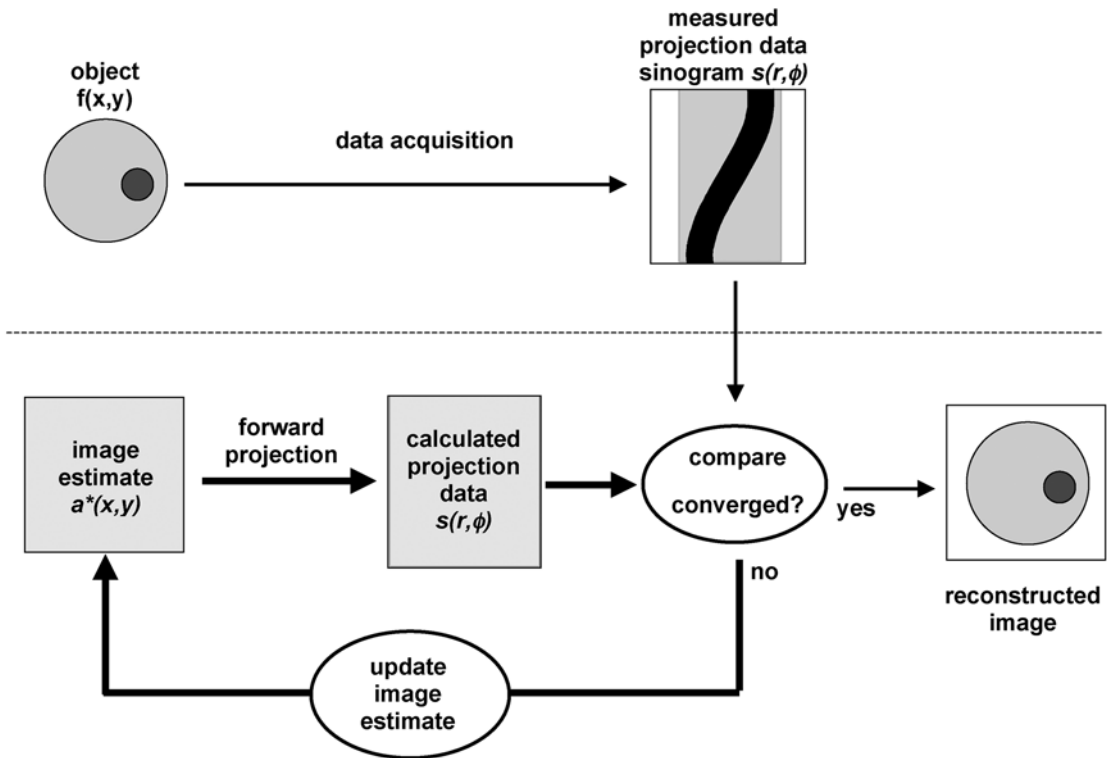


FIGURE 55. Basic flowchart for iterative image reconstruction. Image estimate is forward-projected to calculate a sinogram, which is then compared with the measured sinogram. Based on the difference between calculated and measured sinograms, the image is updated. (Reproduced from Cherry SR, Sorenson JA, Phelps ME. *Physics in Nuclear Medicine*, 3rd ed. W.B. Saunders, New York, 2003, with permission from Elsevier.)

dated at each iteration step. The goal is to use methods that reliably converge to the minimum or maximum of the cost function as quickly as possible. Other differences between algorithms include whether and how the statistical nature of the data is implicitly modeled and whether any other prior information (e.g., the expected smoothness of the image) is considered. Some algorithms also force the reconstructed image to be non-negative. For a concise history and review of different iterative reconstruction methods, refer to the review by Leahy and Clackdoyle.⁷⁹

While simple forward-projection is the most straightforward way of calculating the projection data from an image estimate, most iterative algorithms use a more sophisticated approach that models the probability that a gamma ray pair emitted at point $a(x,y)$ in the object/image space is detected in projection element $s(r,\phi)$. By avoiding the simplistic line integral model that is one of the limitations of filtered backprojection, many of the different factors (e.g., system geometry, object and septa scatter, detector characteristics, positron range, and noncolinearity) that determine whether or not a gamma ray is detected, and where it is detected, can be included in the reconstruction process.⁸⁹ The forward-projection step essentially becomes a simulation of the entire imaging system. This can lead to more quantitatively accurate images and improvements in signal-to-noise or spatial resolution.

Two factors make iterative reconstruction approaches much more computationally intensive than their filtered backprojection counterparts. First, each iteration is essentially equivalent in time to a full-filtered backprojection reconstruction. Backprojection is the most time-consuming part of the filtered backprojection algorithm, and backprojection and forward-projection are computationally very similar processes. Most iterative algorithms require multiple iterations (anywhere from 2 up to several hundred, depending on the algorithm and the data) to reach an acceptable image. Second, as soon as the line integral model of simple forward-projection is replaced by a more accurate model of how gamma rays are detected, the time required to compute the projection data increases, as many more image pixels may now contribute to a particular projection element. A number of approximations have been developed to speed up these algorithms. One of the most popular is called *ordered subsets*, also known as OSEM, in which only a subset of the projection angles are used in any one iteration.⁹⁰ This speeds up the algorithm, as the time per iteration is directly proportional to the number of angles that need to be forward-projected.

The most widely used iterative reconstruction approaches are based on *maximum likelihood* (ML) methods. *Likelihood*, is a general statistical measure that is maximized when the difference between the measured and estimated projections is minimized. The *expectation-maximization* (EM) algorithm is an iterative algorithm that maximizes likelihood under a Poisson data model. It implicitly treats the projection data as having a Poisson distribution determined by the counting statistics in each projection bin and thus takes into account the statistical noise in the data. The derivation of the form of the algorithm that is used in PET is lengthy and beyond the scope of this book, but the interested reader is referred to articles by Shepp and Vardi⁹¹ and Lange and Carson.⁹² It is instructive to examine the implementation of the algorithm and to see how the image is updated on each iteration.

As a starting point, the reconstruction problem can be written as follows:

$$s_j = \sum_i M_{i,j} a_i \quad (46)$$

where a_i is the activity in an image pixel i and s_j is the number of counts in projection element j . For projection data taken with a PET camera at 128 angles around the object, with each projection measuring 256 elements, the index j would run from 1 to 32768 (128×256). If this is to be reconstructed onto a 128×128 grid, the index i would run from 1 to 16384 (128×128). $M_{i,j}$ is a large matrix (32768×16384 in the example just given) which provides the probability that gamma rays emitted in pixel i will be detected in projection element j . This matrix provides the model of the imaging system and is a much more sophisticated approach than simple forward-projection. This matrix can be determined by calculation, simulations, measurements, or a combination of these approaches. For example, a point source could be placed at a location corresponding to a pixel in the image to determine the counts detected in every projection element for that image location. This would then be repeated for all image pixels. This would be extremely tedious even though symmetry arguments can be used to reduce the number of measurements considerably. In practice, many of the geometric effects can be calculated, and other factors that it may be desirable to include in the matrix (for example detector scatter), can often be simulated.

The ML-EM algorithm for PET can be written as:

$$a_i^{k+1} = \frac{a_i^k}{\sum_j M_{i,j}} \sum_j \frac{M_{i,j} s_j}{\sum_i M_{i,j} a_i^k} \quad (47)$$

This equation shows how the image pixel intensity a_i at iteration $k + 1$ is calculated based on the estimated image pixel intensity a_i at iteration k and the measured projection counts p_j . Notice that when the estimated projection data exactly equal the measured projection data s_j that (substituting from Equation 46):



$$a_i^{k+1} = a_i^k \quad (48)$$

and the image does not change any more. This never actually occurs in practice because of noise in the data and inevitable errors and approximations in $M_{i,j}$. Figure 56A shows the reconstruction of a phantom study as a function of iteration number using the ML-EM algorithm and also shows the reconstructed images obtained using an accelerated OSEM approach.


The iterative reconstruction approach has several key advantages. First, it replaces a simple line integral approximation of what the imaging system is measuring with a model that, if correctly implemented, accurately reflects the probability that a gamma ray emitted at a certain location in the object is detected in a given projection element. This can lead to improved spatial resolution as demonstrated by Figure 56B which shows a comparison of FDG brain images reconstructed with filtered backprojection with a maximum a posteriori (MAP) algorithm (155) that contains an accurate model of the PET system used to acquire the data. Second, the sampling criteria are relaxed, in that the effect of missing data tends to produce lower spatial resolution, rather than streak artifacts that would commonly be encountered with FBP. Third, the statistics of each line of response can be taken into account, with more weight being given to measurements that have better statistical quality.

Generally, iterative methods can produce better signal-to-noise at a given spatial resolution than filtered backprojection methods, partly for the reasons above, and partly because they do not directly convolve the data by a filter that amplifies high-spatial frequency components such as noise. In some instances, factors of two improvements in signal-to-noise over FBP have been demonstrated with iterative algorithms, which is equivalent to a fourfold increase in the effective sensitivity of the scanner. Another perspective is that the same quality image could be obtained with iterative methods using just one quarter of the injected activity or in one quarter of the imaging time. The disadvantages of iterative reconstruction relate to the computational issues already discussed and to the fact that these methods are nonlinear, which can make their behavior quite difficult to predict. This un-

FIGURE 56. A: Illustration of ML-EM reconstruction of a phantom study as a function of iteration number. Also shown are images reconstructed using an accelerated OSEM algorithm with 4 iterations (using 16 subsets) and 8 iterations (using 16 subsets). The OSEM images are placed under the ML-EM images that are qualitatively of similar quality. In this case, the OSEM required 8 to 16 times fewer iterations compared to the conventional ML-EM algorithm to produce images of similar quality.



A.



B.
FIGURE 56. *Continued.* B: A comparison of filtered backprojection (FBP) and maximum a posteriori (MAP) reconstructions of FDG brain images. Notice the improvement in spatial resolution in the MAP images due in large part to the accurate system model incorporated in this particular algorithm. (Image courtesy of Dr. Richard Leahy, University of Southern California.)

predictability, together with the many different possible implementations of iterative approaches, have caused a certain amount of confusion and concern, which has limited more widespread use. However, the option to use ML-EM algorithms is now found on many commercial PET cameras and powerful multiprocessor desktop computers, and, combined with acceleration techniques such as OSEM, makes reconstruction times acceptable in many instances.

Iterative approaches also can be easily adapted to reconstructing 3-D PET data, although the computational complexity increases dramatically. The matrix $M_{i,j}$ becomes very large because of the added dimensionality of the problem and reconstruction times, even with accelerated approaches or specialized hardware, can take many tens of minutes or longer depending on the specific geometry of the problem.

IMAGE ANALYSIS

Image display

The end result of the PET acquisition and image reconstruction is a 3-D image volume where each individual voxel (volume element) represents the regional tissue radioactivity concentration. The most common way to visualize the data is to display the volume as a series of transaxial cross-sectional images on a gray scale (or pseudo color) display, where each gray level (or color level) represents a particular activity concentration, or if the data are processed through a tracer kinetic model, a specific biologic parameter. In most modern PET systems, the axial sampling or plane separation is fine enough to allow the reorientation of the data into coronal planes, sagittal planes, or any other arbitrary orientation. The most common display format for clinical studies is to show coronal sections for whole-body imaging and transaxial sections for brain imaging. It also is common to simultaneously display transaxial, coronal, and sagittal sections in which the cursor is linked between the three image windows, allowing for easy navigation through a large 3-D dataset.

Calibration and region of interest analysis

One of the advantages of PET over some other imaging modalities is that it can accurately determine the activity concentration of the radiotracer within a volume. This can provide information that would allow, for instance, the classification of a lesion in terms of its metabolic rate. To do this, it is necessary to accurately calibrate the system such that it is possible to convert the image count density into an activity concentration, which is most commonly done by scanning a uniform cylindrical phantom with a known activity concentration, A (Bq/ml). The phantom data are then corrected for attenuation, scatter, and so on, and reconstructed with the same parameters that are to be used in the clinical studies. From the reconstructed images, the average count density, C (counts per voxel per second), within the central portion of the phantom image is determined. Because the activity concentration in the phantom is known, the calibration factor (CF) between image counts and activity concentration is then:

$$CF = A(\text{Bq/ml}) \times B.F. / C(\text{counts/voxel/sec}) \quad (49)$$

B.F. is the branching fraction, which is the fraction of decays that occur via positron emission for the radionuclide of interest. This term is necessary because the radionuclides used in PET do not necessarily decay by 100% positron emission (Table 2). The calibration factor allows the determination of the activity concentration within a region in the reconstructed images. Provided that all corrections have been applied, the accuracy of this activity concentration is typically within 5%.

EXAMPLE 18

A region of interest analysis of a noncalibrated ^{18}F FDG PET image resulted in an average image count of 3240 counts per voxel/sec. A uniform calibration cylinder (volume 6000 ml) filled with 40 MBq $^{68}\text{Ge}/^{68}\text{Ga}$ (branching fraction for ^{68}Ga is 0.89—see Table 2) was imaged earlier and resulted in an image count of 72300 counts per voxel/sec. Determine the activity concentration in the region in the ^{18}F FDG PET image. The branching fraction for ^{18}F is 0.97.

ANSWER

First determine the calibration factor CF from the ^{68}Ge cylinder data.

$$\begin{aligned} \text{CF} &= (40 \text{ MBq} / 6000 \text{ ml}) \times 0.89 / 72300 \text{ (counts per voxel/sec)} \\ &= 0.082 \text{ (Bq} \cdot \text{sec /ml} \cdot \text{counts per voxel)} \end{aligned}$$

To determine the activity concentration in the ^{18}F FDG PET image, multiply the counts from the image with the CF and correct for the branching fraction for ^{18}F :

$$\begin{aligned} \text{Act. Conc.} &= 3240 \text{ (counts per voxel/sec)} \\ &\quad \times 0.082 \text{ (Bq} \cdot \text{sec /ml} \cdot \text{counts per voxel)} / 0.97 \\ &= 274 \text{ (Bq/ml)} \end{aligned}$$

The most common way to determine what the local activity concentration is in a PET image (or volume) is to define a *region of interest* (ROI) on an image, using an image analysis software package. The average voxel value within the region is calculated. Using the calibration factor from Equation 49, this value is then converted to an activity concentration in becquerels per milliliter. Because PET data sets usually are made of a stack of images that form a volume, a *volume of interest* (VOI) can be defined by connecting several ROIs defined on multiple contiguous planes into a single VOI. By defining a VOI, the error in the average activity concentration due to counting statistics is reduced.

If a dynamic sequence of images has been acquired, the ROIs or VOI that have been defined can be applied to the same region (or volume) on all images to generate a *time activity curve* (TAC) that shows the radiotracer concentration in a specific region of the body over time. This time dependent data set can then be used with a compartmental model to determine biologically meaningful parameters and to construct parametric images.

Image segmentation

Image segmentation is an analysis tool used in image processing that classifies pixel elements into regions or classes that are homogenous with respect to one or more characteristics. In the analysis of a CT or MRI study, image segmentation is used to delineate different tissue types, such as the separation of gray and white matter areas in the brain. In the case of a PET study, image segmentation could aid in the determination of the extent of areas of radiotracer uptake. Image segmentation also can be used in combination with image registration techniques (see Image Registration, p. 95) where different tissue types or regions are identified on an image with structural information (such as an MRI or CT image). These regions are then transferred onto a spatially registered PET image to determine the level of activity uptake in these regions.

One of the primary uses of image segmentation in PET is the processing of measured attenuation correction scans as a method to reduce image noise.^{45–47} As discussed in Attenuation Correction p. 56, measured attenuation corrections are inherently noisy, and this noise propagates into the final emission image. Image segmentation can be used in the process of the measured attenuation data to reduce the amount of noise that propagates from the measured attenuation correction. Figure 57 illustrates the effect of image segmentation on a whole-body transmission scan. Figure 57A shows the reconstructed transmission data, in which the gray levels represent the distribution of linear attenuation coefficients. The relatively high level of noise in this image is due to poor counting statistics, which, in turn, are due to the relatively short acquisition time (~15 minutes). Using image segmentation, where pixel values within specific gray-scale ranges corresponding to air, lung, soft tissue, and bone are replaced with a fixed value or a narrow range of values, the noise in the transmission image is greatly reduced (Figure 57B). This image volume is then forward-projected to generate the appropriate attenuation correction factors, which will contain less noise compared to the original attenuation correction.

Image registration

Image registration refers to the process in which image volumes are realigned into a common anatomical coordinate space. The three main applications for image registration in PET are:

1. Correction for patient motion.
2. Correlation of PET images to other imaging modalities (e.g., MRI and CT).
3. Comparison of image data within and between different subjects.

A common complaint about PET imaging is that the entire procedure is lengthy, especially in quantitative research protocols with multiple isotope injections. Therefore, patient motion is potentially a problem. Motion during a scan not only introduces a loss in spatial resolution in the final image but may also make regional quantitation impossible in dynamically acquired studies or in studies involving multiple scans. A common clinical protocol for FDG brain imaging is to acquire data for 20 to 40 minutes following the injection and uptake period. For many patients, it is very difficult to remain motionless during the scan time, despite head restraints. Therefore, the data are collected as a series of short frames, which can be viewed, following the acquisition, as a dynamic sequence to detect patient motion. If the patient moved, then only the frames in which the patient remained stationary are added together, and the frames with motion are discarded. Although the amount of patient motion has been reduced, the drawback is that, by discarding data, the image noise is increased due to the lower number of counts that contributes to the final image. This problem could potentially be eliminated if all images are registered, prior to summation; thus, all the acquired information contributes to the final image.

There are many methods for image registration and numerous articles have been published on this subject. An excellent overview of medical image registration techniques is provided by Hill et al.⁹³ The most successful application of these techniques has been in the registration of brain images, where the registration process for within subject registration is typically limited to rigid translation and rotations of the two image volumes. The challenge in image registration is in determining the transformation that will produce the best possible




FIGURE 57. A (left): Whole-body transmission image. B (right): Segmentation of image on the left in which different gray levels have been segmented out and assigned a narrow range of attenuation values that correspond to the known values for air, soft tissue, and bone.

alignment. One method for aligning two volumes is to view them on a screen and manually identify common features or manually adjust registration parameters until the two volumes appear to match. Using manual alignment methods, a trained user can relatively quickly register two image volumes within a few millimeters, provided the images are free of spatial distortion. Because manual registration can be labor-intensive, several more or less automated algorithms have been developed over the years. One group of these is based on the assumption that the information content in the two images is similar. The main difference between the published implementations is in what criterion is used for determining image correspondence (e.g., stochastic sign difference or the sum of absolute differences).^{94,95} This approach is suitable for correction of patient motion within a scan, such as in the example given earlier, or to register two image volumes of the same subject acquired on two different occasions.

Many times it is desirable to register the image of the same subject obtained using different tracers or, more commonly, to register intramodality images (e.g., PET to CT or PET to MRI). Intramodality registration is of particular interest as it allows the biologically specific PET signal to be mapped onto the high-resolution anatomy provided by MRI or CT. In these situations, the information content in the two data sets are typically dramatically different. Using the similarity criterion generally does not provide a robust solution to the registration problem. Nonethe-

less, robust registration methods have been developed that seek to maximize the overlap of the volumes or surfaces⁹⁶ or minimize the standard deviation of the ratio of the pixel intensities.⁹⁷ If the PET radiotracer has very specific uptake in a small fraction of the volume (e.g., [F-18]fluoro-L-DOPA in the human brain) or the images have relatively high noise levels, these methods may fail. In these situations, it might be better to register the PET transmission images rather than the emission images; however, this assumes that the patient remained stationary between the transmission and emission scans.

The most sophisticated type of image registration involves registration of images of different subjects into a standardized atlas (e.g., Talairach space). This application has been extensively used in brain activation studies but also in studies of different types of diseases, such as dementia. Although, at a gross level, the anatomical and functional structure of the human brain is common among individuals, broad deviations are seen in size and shape, and there is significant variation in the appearance of the cortex at the gyral level. To register one person's brain to another, the registration process must involve different nonrigid transformations that deforms or reshapes the brain. The main challenge here is to not only introduce a deformation that makes the two brains similar in shape but also to ensure that functional areas are registered. The process of elastic deformation into a common space is also sometimes referred to as spatial normalization.^{98–100} The use of spatial normalization allows the comparison of regions of brain activation across several subjects that are given the same stimulus.¹⁰¹ Spatial normalization also allows the use of a standardized ROI atlas, from which uptake in specific anatomical or functional areas can rapidly be extracted, thus eliminating the need for manual drawing of regions.¹⁰² Several groups also have assembled databases of the normal uptake of a specific tracer, such as FDG, to which an individual's image is compared to determine areas of abnormal uptakes. This has been applied in the detection of Alzheimer's disease using ¹⁸FDG.¹⁰³ Figure 58 shows an example of intrasubject image registration.

Partial volume effects

One of the main difficulties in the ROI and VOI analysis is to accurately determine the activity concentration in regions or volumes that are small compared with the resolution of the PET scanner. This is of particular importance in the quantitative characterization of small lesions or structures. As discussed earlier, a PET system has limited spatial resolution, such that decays from an infinitely small point source of radiation will be smeared out and appear as a finite-sized blob (equal to the point spread function of the scanner) of lower activity concentration in the reconstructed PET image. This is known as the *partial volume effect*.¹⁰⁴

The result of the partial volume effect is that small objects appear to have lower activity concentration in comparison to larger objects of equal activity concentration (Figure 59). In this figure, a set of spherical sources of equal activity concentration are simulated (A, top row). In Figure 59B, a smoothing filter of 10-mm has been applied in all three dimensions to simulate a 10-mm resolution detector system. As can be seen from the profiles (Figure 59C) through the images, the activity concentration is accurately measured in the larger objects, although there is smearing at the edge of the sources due to the limited spatial resolution. As the sources become smaller, the plateau of the measured activity concentration in the center of the sources diminishes, and, eventually, a suppression of the peak activity concentration is seen. The degree of suppres-




FIGURE 58. Illustration of intrasubject registration of PET images. A (top row) shows single sections through the midbrain of 6 individuals from their FDG-PET scans prior to any image registration. B (second row) shows the images following a nonlinear registration to a “reference” brain shown in C (bottom left). The sum of the 6 images in B is shown in D (bottom right). Notice the high correlation in location and shape of structures after image registration (B) that allows the scans from these 6 different individuals to be added together while still maintaining resolution of the major brain structures.





FIGURE 59. Simulated data showing the partial volume effect on a set of spherical sources (diameter ranging from 0.5 cm to 4.5 cm) of equal activity concentration (A, top row). Due to the partial volume effect (10-mm Gaussian smearing), the activity concentration in the smaller spheres appears to be lower (B, second row). The sphere diameter in centimeters is given above each profile (C, bottom).

sion is both a function of object size and the system reconstructed image resolution; it is characterized by the recovery coefficient (RC):

$$RC = \frac{\text{Measured peak activity concentration}}{\text{True activity concentration}} \quad (50)$$

The recovery coefficient for a spherical object as a function of object diameter, normalized to the image resolution, is shown in Figure 60.




FIGURE 60. Calculated recovery coefficient for the system simulated in Figure 59. For a sphere with a diameter greater than 3 times the image resolution (FWHM), the activity concentration is accurately preserved.

To accurately estimate the activity concentration using ROI analysis, it is, therefore, important to know the size of the object and the reconstructed image resolution. If the size of the object is approximately three times greater than the image resolution, a small central region of interest placed over the object would accurately represent the activity concentration within the object. The size of the ROI has to be small enough to minimize partial volume effects but also large enough to reduce statistical noise (by averaging across voxels). The curves in Figure 61 show the effect of the region of interest size on the measured activity concentration as a function of object size, normalized to the image resolution. As can be seen from this figure, if the size of the region is identical to the object size, underestimation of activity concentration will occur. However, as the ROI size is reduced, the amount of underestimation is reduced.

If the recovery coefficient can be determined, then it is possible to correct for partial volume effects. The recovery coefficient depends on both the physical dimensions of the object and the image resolution. The image resolution of the PET system can be easily determined by measuring its point spread function. However, the dimensions of the structure or lesion of interest are in general much more difficult to measure, with one particular problem being that they cannot easily be determined from the PET images because the extent of the structure or lesion seen in the PET images is distorted by the partial volume effect.¹⁰⁴

The dimensions of the structure can be estimated if anatomical information from other high-resolution modalities such as either a registered MR or CT is provided with the PET image. Using this information together with the meas-




FIGURE 61. Effect of the region of interest size (relative to the image resolution R) on the estimated activity concentration for spheres of varying diameter. If the region size is large in comparison to the object size, the estimated activity concentration is underestimated. As the region of interest size approaches zero, the degree of underestimation approaches the recovery coefficient (RC).

ured PET image resolution, a recovery coefficient can be estimated.^{105–107} However, the problem remains a complex one, because not only is it necessary to correct for the apparent suppression of the activity concentration, but if the region of interest is surrounded by tissue with radiotracer uptake, it is also necessary to correct for cross contamination or activity spillover from these surrounding areas into the region of interest due to the limited spatial resolution. An example of this “spillover” effect would be the contribution to the signal in gray matter regions of the brain from activity in the adjacent white matter, or the contribution of signal from radioactivity in the blood to the determination of myocardial activity concentrations due to the large blood pool in the heart. Furthermore, the radiotracer may not be homogeneously distributed in the structure or lesions that are of interest. A more generalized model for partial volume correction that accounts for the cross contamination and heterogeneous tracer uptake has been proposed by Rousset et al.¹⁰⁸ Nonetheless, in most imaging situations, correction for partial volume errors are only estimates at best, as some assumptions are always needed.

PERFORMANCE EVALUATION OF PET SYSTEMS

To objectively compare the performance of different clinical PET systems, the National Electrical Manufacturers Association (NEMA) has developed guidelines on how certain performance parameters, such as spatial resolution and

sensitivity, should be evaluated and presented.^{109–111} The guidelines allow a user, in the process of selecting a PET system, to obtain a relatively unbiased comparison of system parameters. The ability to make these comparisons is probably of most importance in multicenter clinical trials where a variety of systems may be in use and a certain minimum performance standard is required. The original NEMA standard focused on the system performance using a relatively small phantom (20 cm diameter, 20 cm tall), which is appropriate to use to simulate imaging brain-sized objects. Since the publication of the initial standard,¹¹⁰ whole-body PET has become the predominant use of PET, at least in routine clinical practice. Therefore, the use of the small phantom does not adequately describe the count rate situation in whole-body PET studies. Furthermore, the short axial extent of the phantom used in the original standard makes comparisons of systems with large axial FOVs to narrow FOV systems difficult. The new NEMA standard, NEMA NU2-2001,¹⁰⁹ takes some of these shortcomings into account and adds an image quality measurement appropriate for whole-body imaging. Approaches for making a number of important performance measurements are outlined below.

Reconstructed spatial resolution

Spatial resolution measurements are made using an ^{18}F -point source (dimensions 1 mm or less). The FWHM and the FWTM are reported for several prescribed source positions in the FOV (Figure 62) so that variations in resolution can be determined. At least 100,000 counts are collected for each acquisition and the data are reconstructed with a ramp filter and, if possible, a zoom that results in pixels of dimension of at most $0.1 \times$ the anticipated FWHM. Three components of resolution are measured by taking orthogonal profiles through the reconstructed image of the point source: the radial and tangential components are in the transaxial plane defined as shown in Figure 62, and the axial component is along the axis of the scanner.

Scatter fraction

The *scatter fraction* (SF) is a measure of the contamination of the data from scattered photons, which depends on factors including the geometry of the scanner, the shielding (such as septa), and the energy window. The scatter fraction is defined as the ratio of scattered to total events measured at low counts rate to minimize accidental coincidences and dead time. In the initial NEMA standard from 1994,¹¹⁰ the scatter fraction is estimated from data acquired with a line source of activity placed at different radial offsets in a cylindrical phantom (20 cm diameter, by 19 cm tall, inner diameter) filled with (nonactive) water. A 24-cm FOV is defined for all scanners. Profiles through the sinogram are used to estimate the number of scattered events within the FOV and the number of true events within a 2-cm radius of the source (Figure 63). Scatter within the peak is estimated by assuming a constant background of scatter under the peak. The measurement is repeated at three radial positions: 0, 4.5 and 9 cm. At 4.5 cm and 9 cm, the sinogram profile must be analyzed as a function of angle. The average scatter fraction is calculated by weighting the totals counts and scattered counts measured at each position of the source by the relative area of the annulus at that radius (Figure 63):

$$SF = \frac{S(0) + 8 \times S(4.5) + 10.75 \times S(9)}{T_{tot}(0) + 8 \times T_{tot}(4.5) + 10.75 \times T_{tot}(9)} \quad (51)$$




FIGURE 62. Illustration of the positioning of the point sources to measure the transaxial and axial resolution in a PET system. Abbreviation: FOV, field of view.

where S is the number of the scattered counts per unit activity and T_{tot} is the total number of counts (true + scattered) per unit activity. This measurement is performed in both 2-D and 3-D. This measurement represents the scatter fraction in brain imaging. In NEMA NU-2 2001,¹⁰⁹ the scatter fraction is defined for whole-body imaging. The scatter fraction is determined in a similar way with the difference that a 20-cm diameter and 70-cm long cylinder is used. Furthermore, the scatter fraction is only determined for a single off-center position (4.5 cm).

Sensitivity

The sensitivity of a PET scanner is defined as the counting efficiency of the system for a known amount and distribution of activity. To measure the absolute sensitivity of the scanner, a 700-mm long, 5-mm diameter tubing is filled with a known amount of activity. The activity in the source should be

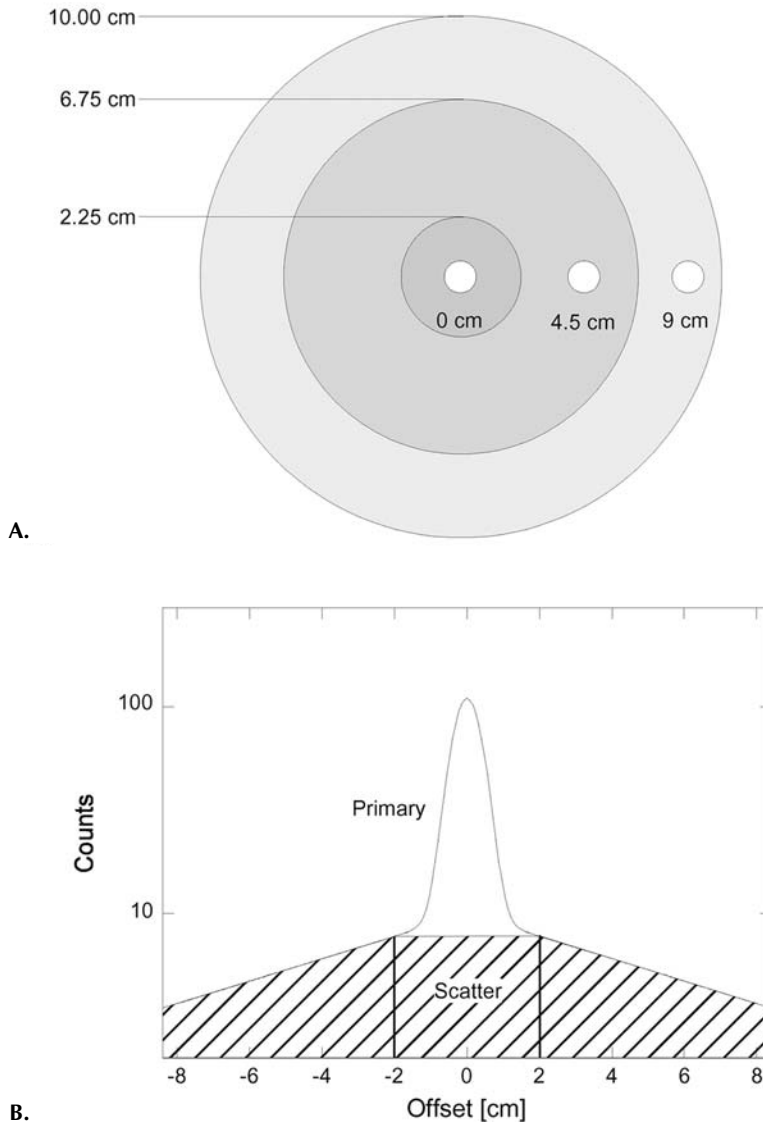


FIGURE 63. The scatter fraction is determined from measurements of a line source placed within a 20-cm diameter phantom (A). The sinogram profile (B) is used to estimate the number of scattered events within the FOV and the number of true events within a 2-cm radius of the source. Scatter within the peak is estimated by assuming a constant background of scatter under the peak. The measurement is repeated at three radial positions: 0, 4.5 cm, and 9 cm; at 4.5 cm and 9 cm, the sinogram profile must be analyzed as a function of angle. The average scatter fraction is calculated by weighting the totals and scatter measured at each position of the source by the relative area of the annulus at that radius (Equation 51).

low enough to ensure counting losses of less than 1% and a random coincidence fraction of less than 5% of the true coincidence counting rate. To ensure that the emitted positrons annihilate and produce a pair of 511-keV photons, the source has to be surrounded by an attenuating sleeve. Although the


sleeve will ensure positron annihilation, it will also attenuate a fraction of the emitted photons, which will prevent a direct measurement of the absolute sensitivity. Instead, by successive measurements of the count rate using different sleeves of known thickness, the attenuation-free sensitivity can be determined by extrapolation.¹¹² This measurement is performed in the center and at a 10-cm radial offset in the FOV.

Count-rate performance and dead time

Measurements of count rate as a function of activity concentration are performed with a 20-cm diameter, 70-cm long, uniform cylinder phantom initially filled with a high activity concentration (typically 25–50 kBq/ml) and allowed to decay. The recommended radionuclide is ^{11}C because the shorter half-life makes the total measurement time more manageable than with ^{18}F . Data should be acquired until the fraction of random coincidence events and the system dead time are negligible. The rate of prompts (true + scatter + randoms), random (from delayed window) and scattered (using scatter fraction, p. 102) coincidences are recorded at convenient time points (e.g., 4 points per half-life).

The deviation of the trues rate at high activity concentrations from an ideal, linear dependence is due to scanner (detector + electronics) dead time. The percentage dead time, as a function of increasing activity concentration, is defined as: $\%DT = 1 - T/T_{ex}$, where T is the actual trues rate and T_{ex} is the trues rate linearly extrapolated from low count rate data (Figure 40). The concentration at which the dead time reaches 50% also is a measure of scanner performance (Figure 64). A high sensitivity scanner (e.g., one operated in 3-D) typically saturates at relatively low activity concentration (compared to 2-D), even though the maximum trues rate is higher than that of a less sensitive

FIGURE 64. Example of a typical dead time curve as a function of activity concentration within a 20-cm diameter cylinder phantom. The concentration at which the dead time is 50% should be reported.



system which peaks at a higher activity concentration (e.g., the same scanner operated in 2-D).

Noise equivalent count rate

The count rate measurements do not directly indicate image signal-to-noise in the presence of relatively changing trues, randoms and scatter rates. A better measure of signal-to-noise is provided by the *noise equivalent count rate* (NEC),¹¹³ defined by:

$$NEC = \frac{T^2}{T + S + 2 fR} \tag{52}$$

where T , S and R are the true, scatter, and random coincidence counting rates, f is the fraction of the sinogram width subtended by the phantom, and the factor 2 comes from on-line randoms subtraction (see Correction for Random Coincidences, p. 65). The NEC provides only a global measure of the signal-to-noise ratio because it is not sensitive to regional variations of the source distribution. Figure 65 shows a set of typical count rate curves from a clinical scanner for prompt, true, and random coincidences as well as the resulting NEC count rate.

Image uniformity

Uniformity measures the deviations in the reconstructed image from a uniform response. A 20-cm diameter uniform cylinder is filled with a moderate activity concentration (fraction of random coincidences and dead time < 20%) and imaged when positioned ~2.5 cm off-axis transaxially. An average of 20 million




FIGURE 65. Example of prompt (squares), random (triangles), and true + scatter (circles) count rate curves. The dashed curve is the resulting NEC curve calculated from the measured count rate curves using Equation 52.

counts per slice are acquired and reconstructed with a ramp filter. A grid of 1-cm \times 1-cm regions of interest are inscribed in a circle 18 cm in diameter centered on the image of the cylinder in each slice. Positive and negative nonuniformity (NU) are defined by:

$$NU(+)=+(C_{max}-C_{ave})/C_{ave} \quad (53)$$

$$NU(-)=- (C_{ave}-C_{min})/C_{ave} \quad (54)$$

where C_{max} is the maximum number of counts in any square ROI within the slice, C_{min} is the minimum number of counts, and C_{ave} is the average number of counts of all regions in the slice. The necessity to collect such high statistics is to minimize the nonuniformity due to statistical effects and focus on those due to scanner imperfections, software corrections, and the reconstruction algorithm.

PET SYSTEM DESIGN

High-performance dedicated clinical PET scanners

Dedicated PET systems have undergone dramatic changes since the first PET systems were designed in the mid 1970s.¹¹² Although the first systems could produce images at a resolution of 1 cm to 2 cm, these were low-sensitivity, single-slice, small-diameter systems, using heavily collimated NaI(Tl) detectors. State-of-the-art, high-resolution whole-body systems now have an intrinsic resolution of 3 mm to 5 mm, orders of magnitude of higher sensitivity, and an axial coverage of 15 cm or greater, with a minimal amount of collimation.¹¹⁴⁻¹¹⁸ The design of the most widely distributed high-end PET systems is fundamentally the same. The detector system in these systems is based on the block detector concept (Block Detector, p. 22). Depending on the manufacturer and the particular model, the size of the individual detector elements vary. For example, they are $4.0 \times 4.4 \text{ mm}^2$ in cross-section on the ECAT HR+ (CTI/Siemens, Knoxville, TN) and $6 \times 8 \text{ mm}^2$ on the Advance (GE Medical Systems, Waukesha, WI). These systems generally cover an axial field of view of 15 cm to 16 cm, producing between 35 to 63 simultaneous cross-sectional image planes. The total number of detector elements in these systems can be as high as 18,000. Table 5 summarizes the design and performance of several commercial, dedicated PET systems. A photograph of a typical clinical PET scanner is shown in Figure 66 and Figure 67 shows a whole-body ^{18}F -FDG PET scan acquired on a similar system.

These systems all have the ability to collect data in both 2-D and 3-D mode, as discussed in Image Reconstruction (p. 86), where the lead septa located between each detector ring can be removed to allow the collection of oblique coincidence lines of response. These systems all have built-in rod sources that can automatically be extended and retracted for the acquisition of transmission data. These high-end dedicated systems provide the user with the maximum flexibility in the type of PET studies that can be acquired. The large number of detector channels allows high count rate studies to be performed with nominal dead time losses (especially in 2-D mode). The full ring geometry allows fast dynamic scans to be acquired.

TABLE 5. Comparison of Several Commercially Available Clinical PET Systems


	Scanner					
	ECAT EXACT 47	ECAT EXACT HR+	ECAT ACCEL	Advance NXi	C-PET plus	ALLEGRO
Detector Material	BGO	BGO	LSO	BGO	NaI (TI)	GSO
Diameter [cm]	82.4	82.4	82.4	92.7	90	90
Detector Dimensions [mm] (Transaxial × Axial × Depth)	6.75 × 6.75 × 20	4.05 × 4.39 × 30	6.75 × 6.75 × 20	3.9 × 8.2 × 30 mm	Curved Panel 500 × 300 × 25 mm	4 × 6 × 20 mm
Detectors per Block or Module	64	64	64	36	1	638
Number of Rings	47	63	47	35	64/128	90
Axial FOV [cm]	16.2	15.5	16.2	15.2	18	25.6
Transverse FOV [cm]	58.5	58.5	58.5	55	57.6	57.6
Septa	Yes	Yes	No	Yes	No	No
Spatial Resolution [mm]						
2-D Transaxial:						
1 cm	6.0	4.6	6.2	4.8	—	—
10 cm	6.7	5.4	6.7	5.4	—	—
2-D Axial:						
1 cm	4.5	4.2	4.3	4.0	—	—
10 cm	5.9	5.0	5.9	5.4	—	—
3-D Transaxial:						
1 cm	6.0	4.6	6.3	4.0	5.0	4.8
10 cm	6.7	5.4	6.8	5.4	6.4	5.9
3-D Axial:						
1 cm	4.6	3.5	4.7	6.0	5.5	5.4
10 cm	6.5	5.3	7.1	6.3	5.9	6.5
Sensitivity [kcps/mCi/cc]						
2D	180	200	200	200	—	—
3D	780	900	925	1100	450	>700
Peak NEC [kCps]						
NEMA NU-2 (2001)						
2D	44	54	>90	83	—	—
3D	25	38	61	28	16	48

Abbreviations: BGO, bismuth germanate; FOV, field of view; GSO, gadolinium oxyorthosilicate; LSO, lutetium oxyorthosilicate; NaI(TI), thallium-activated sodium iodide; NEC, noise equivalent counts; NEMA, National Electric Manufacturers Association; PET, positron emission tomography; 3-D, three-dimensional; 2-D, two-dimensional.



FIGURE 66. Photograph of ECAT EXAT clinical PET scanner. (Courtesy of CTI Inc., Knoxville, TN.)

FIGURE 67. Whole-body FDG image acquired on the GE Advance clinical PET scanner. (Courtesy of GE Medical Systems, Waukesha, WI.)



Since the early 1980s, the majority of dedicated PET systems have used BGO as the scintillation material.¹¹⁹ Although this material has excellent absorption properties, the relatively long scintillation decay time (Table 4) limits the count rate performance, especially in the high counting-rate environment of 3-D whole-body scans. Because of this limitation, the injected dose has to be reduced when operating a BGO in this 3-D mode, which tends to offset the sensitivity gain. Recently, the first commercial whole-body system was introduced using LSO as the scintillation material. The speed of LSO allows 3-D scans to be acquired at the full injected dose, which will provide an improvement in image quality and/or improved patient throughput.¹²⁰ For example, whole-body scan times can be reduced from about 1 hour to under 10 minutes.

Lower cost clinical PET scanners

The overall cost of a PET scanner is directly proportional to the number of detector modules and the amount of associated electronics. These components account for approximately 50% of the manufacturing cost of a PET system. One approach to reduce the overall cost of the system is to reduce the number of detector channels. This is an approach used in the ECAT ART scanner,¹²¹ which is a partial ring system (Figure 21, upper right) in which the detectors rotate to collect a complete data set. To compensate for the lost sensitivity caused by the reduced number of detector channels, this system has no interplane septa and is operated in 3-D mode only. The need for detector motion limits how fast dynamic frames can be acquired.

A different approach for reducing overall system cost is to use less expensive detector technology. In the Philips/ADAC C-PET systems, NaI(Tl) is used as the detector material, which is less expensive in comparison to BGO. This system uses continuous detectors based on curved plates of NaI(Tl) read by a matrix of relatively large PMTs (Continuous Gamma Camera Detector, p. 25).¹²² The total number of PMTs used in this design is roughly a factor of 3–4 less compared to PET systems based on the block design. The use of less expensive continuous NaI(Tl) detectors allows the construction of a system with a longer axial FOV, which to a certain degree compensates for the lower efficiency of NaI(Tl). Like the ECAT ART, the C-PET system is also a 3-D-only system. Recently, Philips/ADAC (ADAC Laboratories, Milpitas, CA) introduced the Allegro system, where individual $\sim 4 \times 4$ mm GSO crystals are mounted on an array of PM tubes. This is similar to the panel detector concept shown in Figure 14 (bottom). GSO has both better absorption characteristics and a shorter decay time in comparison to NaI(Tl), which should provide significant improvements in both sensitivity and count rate performance.

Coincidence imaging on gamma cameras

Shortly after the first scintillation camera was invented in the late 1950s, it was proposed to use a pair of these devices, operated in coincidence mode, for detection of annihilation radiation.¹²³ Since the algorithms for tomographic reconstruction were not yet developed, the device was restricted to planar imaging. With the increased availability of positron-emitting radiopharmaceuticals from distribution centers in the mid 1990s, the idea of using two conventional scintillation cameras in coincidence for tomographic

imaging was revived.¹²⁴ Because such a system is basically a dual-headed SPECT system, upgraded with coincidence circuitry, this system could then also be used with collimators for conventional nuclear medicine studies using single-photon emitting radionuclides.

One of the difficulties in using standard scintillation cameras for PET imaging is that they are highly optimized for low-energy gamma ray imaging (~ 140 keV). Because of the high energy of the annihilation radiation (511 keV), the detection efficiency is very poor using standard three-eighths-thick NaI(Tl) detectors. This can be compensated for by increasing the detector thickness; however, this reduces the intrinsic spatial resolution of the detector. The count rate capability of these coincidence systems also is limited, which is caused by the fact that all detected events have to be processed through only two detector channels. Because these coincidence systems are operated with a minimum of collimation and all events have to be processed by only two detector channels, dead-time and pile-up effects become a serious problem at relatively low activity levels. The end result is that the total number of counts that can be collected within a reasonable time frame is limited. This, in turn, requires the use of low-resolution filters in the reconstruction algorithm to keep statistical noise at acceptable levels for diagnostic quality images, which may limit the visualization of small lesions.

To overcome both the detection efficiency and count rate limitation of the coincidence systems, CTI/Siemens (Knoxville, TN) designed a hybrid PET/SPECT, using block detector technology as found in conventional PET systems and the phoswich concept (Figure 17, top). Using this technology, several parallel channels process the photon flux. To overcome the efficiency problem, this system uses two layers of scintillators, where a front layer of NaI(Tl) is primarily used for detection of the low-energy gammas from single-photon emitters. A second layer of LSO is used to improve the detection efficiency of the 511-keV photons.¹²⁵

High-performance brain imaging systems

Several companies and universities have developed or are in the process of developing high-performance research PET scanners, with a particular emphasis on high resolution and high sensitivity brain imaging. The *ECAT EXACT 3D* (CTI/Siemens, Knoxville, TN) is a 6-ring, 3-D only version of the clinical *ECAT EXACT HR+*, resulting in an axial field of view of 23.4 cm.¹²⁶ This is the highest sensitivity PET scanner built to date, with an absolute sensitivity at the center of the field of view of 10%, and it is largely directed towards low concentration receptor studies in the brain. Another system, the *HRRT (High Resolution Research Tomograph)*; (CTI/Siemens, Knoxville, TN), introduced a number of new features in the quest for higher spatial resolution in the brain.¹²⁷ It combines the quadrant-sharing approach (Figure 14) with a LSO/GSO phoswich detector block (Figure 17) that provides 1-bit depth of interaction information. This was the first human scanner to use LSO scintillator and the first human PET scanner to have depth of interaction capability. It is also the highest resolution commercial PET system built to date for human imaging, with a reconstructed image resolution as high as 2.5 mm. Figure 68 shows one of the first brain studies performed on this ma-




FIGURE 68. Brain images showing glucose metabolism acquired on the high-resolution HRRT scanner—the highest resolution commercial human PET scanner currently in existence. (Image courtesy of CTI Inc., Knoxville, TN.)

chine, showing glucose metabolism in the brain that demonstrates the astonishing detail that can now be visualized by PET. Other high performance brain imaging systems based on GSO detectors and depth encoding LSO detectors with photodiode and PMT readout are being developed at the University of Pennsylvania¹²⁸ and Lawrence Berkeley National Laboratory,¹²⁹ respectively.

Other human PET scanners

A number of other prototype systems based on novel approaches or technologies have been designed and built by research groups across the world. Most notable are scanners based on multiwire proportional chambers with lead or lead glass converters,^{35,36} a combination of BaF₂ scintillator with MWPC readout³⁴ and a variable field of view BGO camera,¹³⁰ which also included an early implementation of the quadrant-sharing scheme (Figure 14).

There has also been considerable interest in developing dedicated positron imaging systems for breast imaging. FDG-PET has been shown to have high sensitivity and specificity in the detection of breast lesions.¹³¹ By placing detectors around the breast, rather than around the whole cross-section of the patient, sensitivity can be dramatically increased, and it should be possible to achieve higher resolution images at a reasonable noise level. A number of designs have been developed, some of which provide simple projection images through the breast,^{132,133} others of which will be capable of some form of tomography. The projection-based systems are designed for incorporation into mammography or biopsy gantries so that coregistered mammograms can be acquired along with the PET data. At the present time, these systems are just entering clinical trials,¹³⁴ so the cost-effectiveness and diagnostic utility of these systems for breast imaging have yet to be determined.

Multimodality PET imaging

A major advance in the late 1990s has been the concept of combining molecular imaging by PET with anatomic information obtained from other modalities, using integrated imaging systems rather than software-based approaches.¹³⁵ In particular, combined PET and CT scanners have been developed that enable coregistered PET and CT images to be acquired in quick succession.¹³⁶ The value




FIGURE 69. Fused FDG-PET and anatomical CT images from the Reveal combined PET/CTscanner. (Image courtesy of CTI PET Systems, Knoxville, TN.) (See color insert.)

of fused FDG-PET and CT images has become apparent to physicians, particularly in oncology,¹³⁷ where the correlation of anatomy (CT) and metabolism (FDG-PET) has become a powerful new diagnostic tool. Furthermore, the CT scan can be used to provide the tissue density information with which to calculate the attenuation correction for the PET images. It is expected that combined PET/CT systems will have important applications in radiation therapy, surgical planning, and guided biopsy procedures.

A number of companies have unveiled products that combine a PET scanner and a CT scanner. One example is the Biograph/Reveal system (CTI/Siemens, Knoxville, TN) which consists of an ECAT HR+ PET scanner (with either BGO or LSO block detectors) integrated with a Siemens Somatom EMOTION spiral CT system. An example of a fused PET/CT image obtained from this system is shown in Figure 69. The Hawkeye system (GE Medical Systems, Waukesha, WI) uses a Millenium VG coincidence gamma camera system with a simple CT system based on a linear detector array.¹³⁸ A new combined PET/CT system based on the GE ADVANCE PET scanner and a high-end, multi-plane spiral CT has also been released (Figure 70). ADAC/Philips (ADAC Laboratories, Milpitas, CA) have a PET/CT system based on their Allegro GSO PET system combined with a spiral CT system. This is a rapidly developing area in clinical PET; systems in which there is more complete integration of the PET and CT systems, in terms of the detector hardware, gantry, and the software, are to be expected.

There have also been some early attempts at producing PET systems that are compatible with MRI scanners so that PET and MRI images can be acquired in



FIGURE 70. Photograph of GE Discovery LS combined PET-CT scanner. (Image courtesy of GE Medical Systems, Waukesha, WI.)

the same setting. This research, however, remains at a very preliminary stage and has been limited so far to phantom and animal studies.¹³⁹

Animal scanners


PET has recently increased its impact in basic biological research, leveraged largely by the development of very high-resolution, relatively low-cost PET systems designed specifically for imaging laboratory animals.¹⁴⁰ The ability to measure a range of relevant molecular and biological processes noninvasively in animal models by PET has opened many new possibilities in animal research, both in research laboratories and in the pharmaceutical industry. High-performance PET systems have been developed for brain imaging in nonhuman primates,¹⁴¹ but perhaps the biggest advances have been in the development of systems for imaging small rodents, particularly mice and rats. An early system based on the same detectors found in clinical scanners provided an initial demonstration of the utility of dedicated animal PET scanners for neuroreceptor studies in the rat brain.¹⁴² Since then, a number of groups have built functional prototype systems using a wide array of PET detector technology including avalanche photodiodes,^{28,29} position-sensitive PMTs,^{143–145} fiberoptically coupled multichannel PMTs with LSO scintillator,¹⁴⁶ and multiwire proportional chamber technology.^{32,37} Figure 71 shows the microPET[®] scanner^{146,147} developed for small-animal imaging. This was the first PET scanner to incorporate the new scintillator LSO. When combined with an accurate iterative reconstruction algorithm,¹⁴⁸ this system can produce animal images with a spatial resolution of approximately 1.5 mm. FDG images of the rat brain obtained with this scanner and reconstructed with the MAP algorithm described by Qi et al¹⁴⁸ are also shown in Figure 71.

Several companies now offer commercially available animal PET systems. The most widely distributed at the time of writing is the microPET[®] scanner

FIGURE 71. Photograph of microPET[®] small animal scanner (A, top). Coronal FDG images of the rat brain obtained using this system and reconstructed with an iterative algorithm containing an accurate system model (B, bottom). (See color insert.)



A.



B.





FIGURE 72. A: Photograph of the microPET[®] Focus animal PET scanner with covers removed to reveal ring of detector modules (courtesy of Concorde Microsystems Inc., Knoxville, TN). B: ¹⁸F-FDG whole-body imaging in a single mouse showing tumor growth (arrows) in mammary fat pads over a period of eight weeks. (Courtesy of Craig Abbey, UC Davis.) (See color insert.)

(Concorde Microsystems Inc., Knoxville, TN). This system is made up from detector modules that consist of an array of $2.2 \times 2.2 \times 10$ mm LSO crystals coupled via a short optical fiber bundle to a position-sensitive PMT.¹⁴⁹ Two different configurations exist, one with a 26 cm diameter bore that can accommodate small non-human primates¹⁴⁹ and a rodent-only version with a 14.8-cm bore.¹⁵⁰ Reconstructed image resolution is approximately 2 mm with filtered backprojection and the sensitivity is 2.2% and 3.4% for the 26 cm and 14.8 cm bore system respectively at the center of the field of view. Images showing the development of tumors in a mouse acquired using a MicroPET[®] system are shown in Figure 72.

A new system, called the MicroPET[®] Focus, has recently been developed by the same company and uses $1.5 \times 1.5 \times 10$ mm crystals to improve the spatial resolution to around 1.75 mm, with a sensitivity of 3.4%. With MAP reconstruction, this system achieves images with a spatial resolution approaching 1.25 mm. Oxford Positron Systems (Weston-on-the-Green, UK) offers a very high resolution animal PET system³² called the quad-HIDAC that is based on the multiwire proportional chamber technology. This system achieves 1 mm reconstructed resolution with iterative algorithms and has stacks of detector mod-



A.

B.

FIGURE 73. A: Photograph of the quad-HIDAC small-animal PET scanner. B: ^{18}F -fluoride ion bone scan in a mouse acquired with this system. (Courtesy of Oxford Positron Systems, Weston-on-the-Green, UK.)

ules to provide depth of interaction information, thus minimizing parallax errors. The sensitivity is around 1.8%. An image from this system is shown in Figure 73.

It is clearly still possible to make significant improvements in spatial resolution and sensitivity for small-animal imaging. A number of groups are developing detectors and systems that can realize approximately 1-mm reconstructed spatial resolution in all three dimensions, leading to a $1\text{-}\mu\text{l}$ volumetric resolution.^{151–153} The combination and integration of small-animal PET with other modalities such as CT¹⁵⁴ and MRI¹³⁹ also is being explored. Finally, the development of detector technology for small-animal PET applications has the added benefit of providing a convenient testing ground for technologies that may ultimately also be applicable for clinical PET systems used in humans.