slide1:

Welcome to Lecture 17 on SPECT and PET. Today, we're finishing the nuclear imaging section. Before I explain the content of this lecture, let's get started by recalling why these modalities matter in biomedical imaging. SPECT stands for Single Photon Emission Computed Tomography, and PET means Positron Emission Tomography.

Both enable molecular imaging and are vital for diagnostics, functional evaluation. They use specific tracers and detectors, which we'll explore in detail today

slide2:

As you can see from the course schedule, we are right on track. Today's lecture focuses on SPECT and PET. After this, we will proceed to MRI, ultrasound, and optical imaging, before concluding the course with the final exam.

So far so good, and everything is proceeding according to plan.

slide3:

Let's review a little bit and address some points that were confusing in the last lecture. Have you reviewed the slides and the book chapter? Have you reconsidered why the dynamic curve was shown on the previous slide?

We utilize a generator system to produce technetium-99m (Tc-99m) and molybdenum-99 (Mo-99), and we observe the dynamic range illustrated. I want to underline that the relationship between parent and daughter nuclides is constant—the decay rate of the daughter is governed by the half-life of the parent. This is a core principle in nuclear medicine generators, particularly in SPECT imaging.

There are three equations we looked at previously. By solving these equations and plotting with MATLAB, you can reproduce the curves and understand the dynamics involved. Normally, we conduct what's called a "milking operation" at peak activity. This produces the types of curves shown here. So when starting from the Mo-99 source, you continually elute Tc-99m for SPECT imaging.

Can anyone tell me why the decay rate of the daughter, which is Tc-99m, is governed by the decay constant or rate of the parent? Do you see this in the equations, or is there something confusing about it? Why do we have this rate relationship? Here's a way to think about it: When learning medical imaging, always try to simplify complicated concepts into simple analogies for better understanding.

For example, imagine parents regularly giving money to a daughter. The parents have a certain amount of fortune and give it at a certain rate. The daughter is very active and immediately spends anything she gets—whether it's on an iPhone or movies, or travel. So, the parents' money declines slowly over time, but the daughter always spends fast, as soon as she receives it. In this analogy, the spending rate is governed by the rate at which the parents give money. This heuristic gives us a simple mental picture: the rate at which the daughter "decays" or spends is determined by the parent's supply rate.

Returning to the technical aspect, if you look at equation 2.17 from the textbook, it may not be immediately clear. But if you break it down and consider the relationship between parent and daughter, it becomes much easier to understand why Tc-99m's activity is governed by Mo-99's decay rate.

slide4:

Now, let’s talk about dead time in nuclear imaging systems. If the efficiency of our system is high, and we inject a large dose of radiopharmaceutical, the number of gamma rays hitting the scintillation crystal can actually be more than the system can handle. This happens because the gamma camera needs some time to recover between events. If two scintillation events happen too close together, they can't both get recorded.

The total dead time, which we call "tau," is given by this equation: tau equals capital N minus small n, divided by small n times capital N. So, let me say it again: tau equals capital N minus small n, divided by small n times capital N.

Here, "capital N" is the true count rate, which means how many scintillations actually happen every second. ”small n" is the observed count rate, so that’s how many the system actually records. If you take one divided by a small n, you get the average time spent on each counted event. If you take one divided by capital N, that’s what you'd have in an ideal system without any dead time. The difference between these two tells us how much measurement is lost because of dead time.

Of course, what we want is for dead time—again, "tau"—to be zero, so every real event is counted. But in reality, as count rates go up, the number of missed events goes up too.

Always try to use simple ideas to help you understand, and don’t get lost in complicated formulas. For example, if a student brings me a huge program and says the results are wrong, I tell them, break it into small pieces—each piece with clear inputs and outputs. Debug those parts one by one. This makes understanding the whole problem much easier and helps you find mistakes faster.

slide5:

Let’s try to keep a simple, heuristic picture in your mind. Having a clear mental image is key to understanding any concept well. Now, let’s talk about our outline and what we’ll cover today. We’ll go through single-photon imaging, computed tomography, and positron emission imaging. These are the two main nuclear imaging modalities for tomographic reconstruction.

So, with single photon imaging, we deal with gamma rays coming from the patient, and with positron emission, we’re focused on positrons. The positron doesn’t stay around for long—it quickly finds a nearby electron. When a positron meets an electron, we use a big word called annihilation, and together they create a pair of gamma-ray photons that shoot out in opposite directions at the speed of light.

So you have two kinds of imaging. With single photon emission computed tomography, or SPECT, we use single photons—pretty obvious from the name. And with positron emission tomography, or PET, it’s actually two photons created from one positron and one electron. Both types of imaging are deeply governed by statistics.

As background, I really want us to review statistical distributions. This part will be brief—just a rapid overview. If you want to dig into details, you can do that later. After the statistics part, we’ll talk about data models. For either single-photon or positron imaging, you need a mathematical model for how you collect data. Last time, I explained the physical concepts and chemical mechanisms, and we talked about collimation. But with formulas, you can actually write computer programs to do image reconstruction based on these models.

So the data modeling part is about writing down the equations. Given a distribution, what measurements do you expect to see? Once you have a forward model, you want to invert the process. That lets you recover the underlying distribution—the actual tomographic images. Now, these images aren't anatomical details or structures, but they show the distribution of radioactive features that you introduced into the patient.

To do this, we need to compensate for attenuation, and we have two options: deterministic and statistical reconstruction. We'll talk about those later. Then, finally, I'll discuss the scanner architecture. You know, in a CT scanner, you have X-ray tubes, X-ray detectors, high voltage, a transformer, a cooling system, a slip ring, and so on—all packed inside the gantry. A nuclear scanner might look similar on the outside, but inside, it's quite different, and I'll explain how.

I'll also mention correction methods, especially scatter correction and how we deal with random counts—how we detect them and remove them. These corrections are much different from how CT works. Finally, I'll mention a bit about the Fully3D conference, which focuses on advanced topics in medical imaging reconstruction, with an emphasis on X-ray CT and related techniques. We're organizing that meeting, and it's a good chance to learn more. If time permits, I'll show you today's homework questions and give comments. That's pretty much the plan—just to give you the big picture, so you can see where each part fits in.

slide6:

Statistical distributions are very important, especially for understanding noise and random variables over different intervals. The simplest case is the uniform distribution. With a uniform distribution, every point in the interval, like zero to one or zero to five, has the same chance of being picked. Every number in that range is equally likely. That's one simple kind of distribution.

Next, we have the Gaussian distribution, and after that, the Poisson distribution. The Gaussian distribution, also called the normal distribution, is your classic bell-shaped curve. You usually describe it with a mean—called mu—and a standard deviation, sigma. If mu is zero and sigma is one, that's the standard normal distribution. Mathematically, you see a scaling factor and an exponential factor involving mu and sigma. This type of distribution is used everywhere in engineering, noise measurement, estimation, and more.

If you have lots of independent random variables, each contributing a bit—maybe some are uniform, maybe some are not—when you put them all together, the outcome will tend to look like a Gaussian curve. This is a fundamental result from statistics, and it's why you see the bell-shaped curve over and over in reality—lots of independent factors add up to give something that looks Gaussian.

The Poisson distribution is related to the Gaussian, but looks very different, especially with small means. The Poisson describes random variables that give you counts—like one, two, three, and so on—with different probabilities for each count. It's controlled by the mean, which we call lambda. If that mean is small, you get a very asymmetric distribution. But as lambda grows larger, the curve starts to look more and more like a Gaussian—eventually, if lambda is really big, the Poisson distribution and the Gaussian distribution look nearly the same.

For the Gaussian, you have mean and variance, with variance noted as sigma squared. For Poisson, both mean and variance are just lambda. You can actually derive the Gaussian by combining lots of individual random distributions—through convolution—and prove it's a bell-shaped curve.

So, in summary, we're going to see both Gaussian and Poisson noise often in imaging, and understanding how they behave is key to handling measurement and image reconstruction accurately.

slide7:

And Poisson noise, Poisson distribution, you would need some high school statistics, you learned permutations and combinations. And this permutation, you just see the different arrangement, roughly speaking, a combination, you need to just remove those, those essentially the same arrangement like ABC and ACB.

In terms of combination, you think it's the same thing because this is same three individuals are involved in this activity. So you need just the distinction between permutation and combination. So you can review this slide, you will recall what you learned.

slide8:

And by nominal formula, basically, you have a series of events. And for an easy event, it's a probability stake. And say you're a toaster coin, and it's a certain probability P, and you see one result. And the other possibility will be Q. Q equals 1 minus P.

So you keep doing multiple experiments. And you know how many times you see high, how many times you see tails, this is a good model. And like the example shown here, if you purchase 10 computers. And individually, you know the defective rate is 2%. If you buy 10 computers in a row, what's the chance you get to buy the computer, then you can use by nominal formula. Basically, something like a toast coin. And the probability for you to see one side may not be exactly half. And in general, you can call it P and Q, Q equal to 1 minus P. And the personal distribution is derived from this binomial distribution. One, the number n is very big.

slide9:

Let's go through the math a bit. I put a green button on this slide to help you remember the key idea: to understand Poisson and Gaussian noise, you need a good grip on the binomial formula, but also on permutations and combinations. When you have a large number of events, you can use some analytic steps to show that the distribution will approach the Poisson form.

For example, in gamma-ray emission—whether it's single gamma-ray emission or paired photons—the process follows the Poisson distribution model, which uses the parameter lambda. So what does this mean in practice? Let's say you're counting gamma-ray photons during a scan. The whole time period gets cut into many small intervals. In each tiny interval, the probability that you get a photon is small, call that P, and the probability that you don't is Q, which is just one minus P. Whether you get a photon or not in any interval is random; there's no way to predict it exactly.

This randomness is at the heart of nuclear physics. The emission is not deterministic—it just follows probabilities. Einstein famously said, "God does not play dice," but some would argue that God indeed lets randomness play a role in the universe. In our context, the emission of gamma-ray photons is governed by pure chance, following the Poisson distribution. When you zoom in to those tiny intervals, you're really

dealing with just the possibility—P or Q—of a gamma-ray being emitted, and it's entirely a game of probability.

slide10:

Let me just sum up: this is a statistical model, called the Poisson model, and you should know a little bit about it, because it also applies to X-ray emission. When you hit tungsten with an electron beam, you get a certain number of X-ray photons coming out, and that X-ray emission also follows the Poisson distribution.

When you detect X-ray photons, sometimes you use photon-counting detectors, which are the latest technology. But most common X-ray detectors are still current-based or energy-integrating types. In traditional X-ray detection, besides Poisson noise, you also get electronic noise from the semiconductor detector design. That electronic noise is modeled by the Gaussian distribution.

So for X-ray CT data acquisition, the full statistical model is actually a combination—Poisson noise plus Gaussian noise. This is just a side note to help you see the big picture. Now, let's look at data models for single photon emission and paired photon emission, and how these statistical concepts fit in.

slide11:

Let's look at this slide—you see a diamond symbol, so pay attention. For nuclear emission, the energy of gamma-ray photons is usually in the range of about 60 to 600 keV, and most importantly, about ninety percent of applications use the isotope technetium-99m. This isotope, Tc-99m, has an energy of 140 keV. Now, keep that number in mind—sometimes I forget, but it's 140 keV, just a little higher than the energy used for X-rays. X-ray energy is usually between 80 keV and up to around 130 keV, so this radiotracer energy is actually quite comparable to the X-ray photon energy range. Now, for positron emission, you'll notice in the homework there's a question about the gamma-ray energy produced by positron emission. The answer is 511 keV, which is much higher.

For single photon emission, here's how we model it. Let's say the source is at position "a." At location "a" in the body, you have a certain amount of radiotracer—a concentration or count of radioactive atoms—denoted as N(a). You collimate the gamma-ray flux at this spot, and you look down that line to see how much signal flows toward the detector. Now, the gamma-ray can be emitted in any direction, so a scaling factor is applied, proportional to the size of that mechanical collimator. At the detector, you measure how many gamma-ray photons you can receive.

If the gamma-ray source were in a vacuum, you'd get just your initial concentration, scaled by the geometric factor—that's determined by how big the collimator opening is. In reality, the radiotracer is inside the human body, so the photons move from "a" toward the gamma camera detector at position "d." Along the way, the photons encounter tissue and undergo attenuation. For gamma-ray photons at 140 keV, their energy is just above the high end used for CT x-rays, but the physics is similar—they're tiny waves.

The formula you see uses an exponential factor, e to the negative integral from "a" to "d" of mu(s) ds. Here, mu is the attenuation coefficient, which depends on position in the body. This exponential term describes how the signal from the source gets weighted by attenuation—it's a line integral, but not a full line like in CT, more of a partial line integral. This is what makes SPECT reconstruction different from X-ray CT reconstruction—so this data model is straightforward but very important.

Now, let’s look at paired photon emission. Here, an event occurs: a positron is emitted, finds a nearby electron, and annihilates. That results in two gamma-ray photons flying off in opposite directions. One gamma-ray photon travels from position "a" to detector d1, the other goes from "a" to detector d2. Both photons are attenuated along their paths. The attenuation factor for each path acts like a probability—the chance that the photon actually reaches the detector.

For a paired event to be registered, both photons have to hit their detectors at the same time—so the total probability is the product of two individual probabilities. Both exponential attenuation factors multiply together, so in the exponent, you just add up the two integrals—from d1 to a, and from a to d2. Together, these cover the entire path from d1 to d2, resulting in a traditional line integral. And the nice thing is, this line integral does not depend on the original source position "a"—it's only about the line between d1 and d2.

So, you have two models here: one where the weighting depends on the source location, and one where the weighting—the response along the line—remains constant, regardless of where the gamma source is. For paired photon events, the formula has constant weight; for single photon events, the weight changes based on where the radiotracer is located. And if the radiotracer moves, the weighting factor changes with it. Any questions?

slide12:

Let’s keep going and look further into this constant weight idea. You’ll notice I put a green button here, and what I want to show is how you can eliminate the location dependence—meaning the weighting factor no longer depends on the source position "a."

When you have two factors, each one representing a probability, you multiply them together. The source position cancels out in the exponent when you do this multiplication. You can actually apply this same concept to SPECT imaging. In SPECT, which uses single photon emission, if you put detectors on both sides and take two measurements—one on each side—you can multiply those measurements together.

When you multiply the measured signals, you end up with the concentration of the radiotracer squared—it appears as a square here because you measured along both paths. Technically, you’re combining two partial attenuation factors: one from "a" to detector d1, and one from "a" to detector d2. When you multiply them, the exponentials add up, and the weighting factor becomes constant with respect to the source location.

So it seems, at first glance, like SPECT imaging can have a constant weighting factor for these combined measurements, just like paired photon emission in PET, which already has constant weighting. But, be careful—real life is not always this simple. There are practical limitations and differences in how signals behave, so always question whether these theoretical simplifications hold up in practice.

slide13:

Let’s look closer—if you have only one radiotracer concentration at a single source, it’s pretty simple. But when you have two sources, or more specifically, two gamma resources, things get complicated. If you try to multiply the measurements, like I showed on the previous slide, you end up expanding two terms from each measurement. That gives you four terms in total.

The first two terms are nice—they give you a constant weighting factor, and the line integral of the concentrations squared. That part would make things look easy, as long as the weighting factor stays constant. But the problem is, you also get cross-product terms when expanding. These cross terms involve measurements from one source to the detector, multiplied by measurements from the second source to the other detector, so you get combinations like from d1 to a and from d2 to b, where a and b are different source positions.

These cross terms are messy and make modeling in SPECT much more complicated. The cross terms mean the data model for SPECT includes weighting factors that still depend on the source positions, even when you know the attenuation background. This partial attenuation messes up the tidy modeling you get with paired photon events in PET; for PET or paired emission, the modeling is straightforward, and you get constant weighting. That's why, for SPECT, you can't just ignore these source-position-dependent factors—they're built into the math and must be dealt with when reconstructing images.

slide14:

When we use a ring detector in PET, the key is to perform coincidence detection. This means we're looking for events where two detectors on opposite sides register gamma photons at the same time—those photons come from an annihilation happening somewhere inside the ring. You can determine a line integral along the green line shown here, or along the blue line for different angles. By rotating the ring or changing the projection angle, you can collect projections from multiple directions.

When you arrange all these projections by angle, you get what's called a sinogram, very similar to what you see in CT. As long as you know the attenuation and apply the correct compensation factors, you can accumulate your data into these line integrals. That's the essential idea behind PET and SPECT data acquisition—grouping your measurements along lines of response and sorting them by angle into a sinogram.

slide15:

Earlier PET scanners used planar imaging mode, which meant you had mechanical collimators defining the imaging planes. These collimators would shape the radioactivity that could reach the detectors, so the ring only collected signals from a thin slice, or plane, of the patient. The detectors were arranged in rings around the patient, and you'd have different parts of the ring recording data for the left and right sides.

Modern PET scanners, though, have removed these mechanical collimators. Instead, you use coincidence detection in three dimensions. This 3D orientation means you can detect gamma photons from many directions, and you're not limited by a physical collimator—anywhere in the field of view, you can pick up valid annihilation events.

slide16:

Visually, you see that 3D PET systems have a clear benefit because they don't use mechanical collimators. With no collimators blocking the gamma rays, you can capture a much greater number of photons, which raises the system's sensitivity and improves image quality. The drawback is that, because the system has a

wide opening, it can also pick up more scattered and random events. These unwanted signals can interfere with the data and reduce the quality of your images.

Later, we'll cover what scattered and random events actually are, and how we try to remove them in post-processing. For 2D PET imaging, where you have strict collimation, only photons from a certain plane are detected, and scattered photons from other directions are blocked. But in 3D PET, you let in more data, and while that's a benefit for sensitivity, it can also make noise and unwanted signals more of a problem.

slide17:

When we talk about PET or paired photon emission models, the weighting factor for attenuation is constant along a given line of response. That means when you do an integral over all source distributions, you can pull this constant weighting factor outside the integral—it doesn't depend on the source position.

You'll see how this works more clearly in the next slide. So, let's shift our focus to image reconstruction, where these ideas become important.

slide18:

Now we're talking about the mathematical model for image reconstruction—the core idea is how to represent our data. For PET reconstruction, we usually assume that the number of detected photons is large, so the probability averages out and the measurements closely reflect the mean values.

That lets us use a deterministic model for reconstruction. For a given line path or line of response, the measured data will primarily reflect the radiotracer concentration along that line, so we can treat the whole reconstruction process as solving for concentrations along known paths.

slide19:

There's a lot of research going on in this area. I should put another green button here for curiosity's sake—because some of these ideas can be solved analytically when you assume the attenuation coefficient, mu, is constant. You can use Fourier analysis, closed-form solutions, or algebraic approaches, like solving systems of linear equations.

Just to point out, for those who are curious, there's a paper shown here about using analytic techniques and the Fourier slice theorem for SPECT. These methods are more complicated than CT, but they show that there are options for solving the reconstruction problem with closed-form solutions or advanced mathematics. Next, let's look at deterministic PET reconstruction.

slide20:

In PET, you see that the data model has a uniform weighting factor multiplied by the concentration at each location. In previous slides, I mentioned how, when you measure along a longer line, you need to take account of the location, but for PET, you measure all coincident detections. Effectively, you sum or integrate

all the individual concentrations, called lambda. The important point is that the attenuation factor remains constant—it doesn't depend on the source position, so it can be factored out.

If the attenuation coefficient, mu, is known, then it becomes an unknown factor in your measurements. But once you know mu—like if you assume the patient's body is mostly water, making mu equal to the value for water, or if you measure mu with a dedicated transmission scan—you can normalize your measurements, and you end up with a clean line integral. With lots of gamma photons, the measurements really do reflect the mean, and you can use deterministic reconstruction, like filtered backprojection.

If you know mu from a special transmission scan or by assuming it's the value for water, then your reconstruction is straightforward. An interesting twist is the time-of-flight approach, which uses the difference in arrival times of photons at the detectors to estimate where the annihilation occurred. By measuring these time differences, you can gain additional information about where the event happened—close to one detector or the other. This time-of-flight technique is advanced, and although we don't have time to fully cover it here, just know it's one way to estimate mu and improve PET reconstructions.

slide21:

Let me explain attenuation correction and why it's so important for PET-CT imaging. When you want to measure the attenuation coefficient, or mu, for PET, you often use an individual CT scan to map out the mu distribution inside the patient's body.

But here's a key point: the photon energy in PET is much different from CT. PET uses gamma photons at 511 keV, while X-ray CT scans use tube voltages that go up to about 140 kVp. The actual X-ray beam covers a spectrum from around 30 keV to 140 keV, but all these energies can be summed up as a single effective energy, usually around 70 keV. So, when you measure attenuation with CT, the coefficient you get is really for that effective energy, about 70 keV.

For PET attenuation correction, you need the mu value at 511 keV, because those more energetic gamma photons interact differently with tissue. To bridge this gap, you use a piecewise linear mapping. That means for any CT-measured attenuation coefficient at 70 keV, you estimate the corresponding coefficient for PET at 511 keV. This mapping follows a pretty reliable, piecewise linear relationship: tissues that have low attenuation at CT energies also have low attenuation at PET energies, and vice versa.

Why do we combine PET and CT? The CT scan provides anatomical structure, while the PET scan gives functional information. Having both types of data in a single framework is powerful—you get two kinds of complementary information. But beyond that, CT also provides the key information needed for attenuation correction in PET. With the mu values from CT, you can calculate the scaling factor, making it possible to convert PET data into proper line integrals for accurate image reconstruction.

slide22:

Once you convert PET data into line integrals, you can remove the attenuation factor, and that will make PET imaging more accurate. Quantitative information can be extracted this way.

Like I said here, the PET result without attenuation correction shows a very dark side. But when you apply the correct attenuation coefficient, you see a much more uniform appearance. The result is more accurate for concentrations—for example, in the liver—making the measurements statistically meaningful. That's

one reason I covered the deterministic reconstruction process first. Now I'm giving you a high-level idea about statistical reconstruction. So, if you have lambda—capital lambda—as the image to be estimated, here's how the statistical approach works. We are given noisy measurements, and we want to estimate the image from those measurements. The measured quantity, called q, is what actually comes from the system, whether you use mechanical collimators or coincidence detection. The measurement is very noisy because the gamma-ray tracer concentration isn't perfect, and the decay rate is reasonable, but the noise is not normal anymore. It's not the normal noise you see in CT reconstruction. The underlying source distribution can be modeled statistically; the image represents the source distribution, broken down into small pixels, where each pixel acts like a small light bulb.

So far, we've talked about image reconstruction for both SPECT and PET. This assumes your data model is good enough—meaning you have a large number of gamma photons captured. In that situation, the measurements reflect the actual statistical mean, so you get a formula—the data model.

The data model is tricky for single photon emission and different for paired photon emission because the weighting factor changes. For single photon emission, the weighting factor fundamentally depends on the location of the active source element. The weighting factor always depends on location.

I explained that this location dependence cannot be easily removed. When you try to multiply paired measurements together, cross terms appear and you can't easily get rid of the spatial dependence.

But for PET photon emission, you're actually able to account for the single source distribution, with a probability factor in the formula. If you have multiple radioactive source points along the same line, you just form the line integral by adding them up. The attenuation factor for any radioactive element along that line is the same.

So you can factor it out. As long as you know mu—the attenuation coefficient—you know the attenuation factor, and you get a neat line integral. Then you use filtered back projection, just like in CT. This is a very nice feature for image reconstruction. If you haven't got the point yet, review the slides—I've been busy these past few days and haven't uploaded this part yet, but I'll refine it and upload it soon. It's important to review the slides and follow the argument—this way, you'll truly understand the data model for SPECT and PET. That's key. Once you know the model, for PET in particular, you can use filtered back projection.

That's why nuclear tomography comes after CT, because you need to have some background knowledge from CT. For SPECT imaging, you have a location-dependent attenuation factor. But if you know mu, and if you measure it—like in an individual transmission scan—you can still convert it into the attenuation rate for your data.

You can turn each measurement into something that forms a line integral—a sum. Then you'll have a system of linear equations, which is just like CT—a system of linear equations to solve. You can use a real high-performance computer for this. Linear algebra gives you the tools to measure and solve these equations, and we know how to do that both numerically and analytically.

But what happens if you don't have enough flux? That's a big problem. I mentioned it the other day: X-ray flux is high, but for nuclear imaging, it's much lower. That's because you're introducing radioisotopes. The number of emitted gamma-ray photons isn't ideal, because you don't want to introduce too much radioactive material into the human body. You need to use the minimum amount for safety. On the other hand, radioisotopes are incredibly sensitive. Any signal you detect comes from the radioisotopes, so you get very high sensitivity.

slide23:

We begin with the idea that each gamma ray measurement is modeled independently. The underlying framework is a plasma model. Consider an image made up of pixels. Each pixel, along with its neighbors, contributes independently—meaning the measurement at one pixel does not directly depend on another. To capture this mathematically, we partition time into small intervals—for example, 10 milliseconds each. One interval, the next interval, and another further away all behave independently. This independence is crucial because it allows us to express the probability of the entire measurement as the product of many small probabilities. Each detector, each pixel, each time interval contributes its own probability, and multiplying them together gives the overall likelihood.

This setup naturally leads to probabilistic modeling. Bayes' rule, which many of us first learned in high school, becomes central here. Imagine a patient comes in, and we perform data acquisition using a camera or a PET scanner. The raw measurements are then organized into sinograms in 2D or 3D. To understand error and uncertainty, we rely on the power of mathematical notation, which is both compact and expressive. In fact, the measured counts can be modeled as random variables following a Poisson distribution. The problem can then be phrased as: given the measured data, what is the most likely underlying image, denoted by capital lambda (Λ)?

This probability depends on several factors. First, assume the true image is known, along with the source distribution and any background activity. If we also know the imaging geometry, then we can predict what each detector should record. For example, a detector may be expected to receive around 3,000 gamma photons during a five-minute scan. The difficulty is that in reality, the image is not known—we only have the measured data. That's why we must also consider prior knowledge about the image distribution. For younger patients without cancer, we may expect little to no radioactive uptake in major organs. For older patients, there may be a higher probability of uptake, for instance, in the colon if there is a tumor. These expectations form a prior distribution on the image.

However, in many cases, especially when data is scarce, we may not want to assume anything specific. In that case, we take the prior distribution to be uniform, meaning every image is equally likely. This constant prior does not affect the statistical estimation, since it cancels out in maximization.

From here, two key concepts arise: maximum a posteriori (MAP) estimation and maximum likelihood (ML) estimation. MAP seeks the image that maximizes the posterior probability, combining both the likelihood of the measured data given the image and the prior probability of the image itself. If we assume the prior is uniform, MAP reduces to maximum likelihood estimation. In ML, we search for the image that makes the measured data most probable.

The process is iterative: we try candidate images, compute the probability that they would have produced the measured data, and adjust until we find one with the highest probability. That image is then reported as our reconstruction. In summary, if the prior distribution is non-informative, MAP and ML give the same result. Otherwise, MAP incorporates prior knowledge to guide the estimation.

slide24:

The IMIL approach is often used because it naturally incorporates statistical knowledge. To explain this more simply, let us recall the idea of CT reconstruction. Imagine you have multiple views of an object. One view

looks like this, another view looks different. Now, you try to reconstruct the underlying image. Suppose the underlying image is a star. When you trace the star vertically, the profile does not look right. When you trace it horizontally, again, it does not match. Then, you try using a different model—say, an ellipse. If you trace the ellipse horizontally, you obtain a profile that matches the measured data more closely. This suggests that the elliptical model fits best. In this case, you would report the ellipse as the reconstructed image.

This is essentially what maximum likelihood estimation does: you keep trying different underlying images, generate simulated measurements, and compare them to the actual data. When the simulated measurements maximize the probability of matching the real ones, that candidate is taken as the true image. This approach is fundamentally different from deterministic CT reconstruction methods. Instead of assuming the data is exact, it operates within a statistical framework.

In this framework, measurements are modeled with a Poisson distribution. Because of the independence of measurements, the overall likelihood can be expressed as a product of individual probabilities. By taking the natural logarithm, this product is converted into a summation, making the formulation simpler and more practical to compute. The key point is that the statistical model accounts for the Poisson noise inherent in the data, which deterministic methods ignore.

To see the advantage, consider a phantom experiment. If you introduce radioactive tracers into the phantom and reconstruct with filtered back projection, you are assuming no noise—that every measurement is an exact line integral. But in reality, measurements always have statistical fluctuations. This mismatch leads to streaking artifacts and random noise in the reconstructed images. By contrast, maximum likelihood reconstruction explains the fluctuations more accurately within the statistical model. As a result, the reconstructed images show fewer artifacts and improved quality, as demonstrated in the phantom example.

slide25:

This example shows patient reconstruction results. When using filtered back projection, the images clearly display streak artifacts and significant noise. These issues appear across all views—coronal, transaxial, and sagittal. Overall, the images produced by filtered back projection are noisy and contain strong artifacts that reduce clarity and diagnostic value.

In contrast, maximum likelihood iterative reconstruction directly incorporates statistical knowledge into the reconstruction process. This approach produces images that are both visually and quantitatively improved. Noise is greatly reduced, artifacts are minimized, and the resulting images are more accurate and reliable for clinical use.

The side-by-side comparison highlights the advantage of maximum likelihood iterative reconstruction over traditional filtered back projection. In real patient data, the iterative approach delivers superior image quality, both in appearance and in diagnostic accuracy.

slide26:

Let’s talk again about PET-CT, which we’ve mentioned several times. First, the CT image provides structural and anatomical information at a high spatial resolution. It shows the fine details and allows us to see

different tissues clearly. On the PET image, you see the distribution of the radiotracer, and in this case, there is a lot of accumulation in one organ. This means the vasculature there is rich, actively taking in nutrients and glucose, which is a sign of rapid tumor growth. The radiotracer—often a form of glucose—is picked up as the tumor cells consume more nutrients than normal tissues. All these radiotracer molecules gather in regions of heightened metabolism, typically where tumors are present.

When you superimpose the PET image over the CT image, you get excellent anatomical context. CT provides what's called the linear attenuation coefficient, and this is used to correct for tissue attenuation in PET. If you use only a uniform attenuation correction, the lungs—which actually have very low tissue density—can be misrepresented. Uniform correction assumes the same background for all regions, so in the lung, this leads to overestimation. In reality, the lungs attenuate the PET signal much less. If you apply uniform correction, you falsely report higher activity in the lungs than is accurate, simply because the tissue is less dense.

CT solves this by providing individualized, tomographic detail for each region. The CT number for the lung is low, and when mapped to attenuation at 511 keV—the PET photon energy—you get a precise, much lower correction for lung tissue. This ensures the activity measurements from PET-CT are more accurate.

Looking at the slide, you see areas where the tumor is intensely bright due to high tracer uptake—these signals are not good news because they reflect significant tumor activity. But after treatment, a follow-up PET-CT scan can show decreased brightness, meaning reduced metabolic activity and, therefore, an effective response to treatment.

Finally, while CT is very good for showing anatomical detail, such as bone and tissue interfaces, it is not sensitive to soft tissue contrast. CT excels at showing dense structures, like bone, but metabolic changes and soft tissue differences are much better visualized by PET. The combination of PET and CT gives you both anatomical layout and functional information, making it a powerful tool in medical imaging.

slide27:

Now let’s look at PET-MRI. MRI provides excellent soft tissue contrast, which is clear here—structures within the organs are visible in great detail. However, MRI is not ideal for imaging bone structures, and some areas, like those in the lungs, can be less distinct compared to CT. We'll be diving deeper into MRI starting from the next lecture, and I’ve already prepared a recorded session because there may be time constraints. I encourage you all to preview the MRI material; it’s a different modality altogether. MRI takes longer to acquire data, and its spatial resolution depends on the specific imaging agent and scan parameters, especially for maximum coverage of larger anatomical areas.

By combining PET with MRI—as shown here—you get unique diagnostic information. PET provides molecular and metabolic data, while MRI supplies detailed anatomical views, especially for the brain and soft tissue tumors. This fusion enables more precise diagnosis and assessment of diseases like cancer and neurological disorders. For brain imaging and soft tissue tumors, PET-MRI truly offers cutting-edge capability, making it a highly valuable technology in modern medical imaging.

slide28:

Let’s move into the fourth part, where we discuss system architecture and scatter correction in SPECT and PET imaging.

This topic is pretty straightforward. We focus on how these imaging systems are built and the basic principles behind scatter correction, which is a major technical consideration for accurate image reconstruction.

slide29:

On this slide, you see a great picture of a SPECT scanner. The SPECT camera system is held together by a sturdy structure, and it’s designed to rotate around the patient. This rotation allows the scanner to collect data from multiple orientations, which is conceptually similar to how CT scanners work.

By rotating the detector and collimators around the patient, we get signals from different angles, enabling a full three-dimensional reconstruction. The system also allows for standard longitudinal scanning, which is what you typically see in practical clinical use.

slide30:

Now, building on the idea used in PET-CT, we can talk about SPECT-CT systems. In this setup, there can be two or even three SPECT cameras working at the same time to gather more data. If you use only one planar gamma detector, a lot of photons—especially those coming horizontally and downward—will be wasted, and their information won’t be captured.

But with triple detectors, you can use gamma photons much more efficiently. The CT part of the system gives you precise anatomical information and helps provide accurate attenuation correction for the SPECT data. This means better quantification, improved image quality, and superior diagnostic capability through combined anatomical and functional information.

slide31:

In pilot imaging, you work with a detector ring. This setup can be used for both human studies and small-animal imaging. Small animals are especially important in pharmaceutical research, since drug companies rely on them to test potential new drugs. The development of a drug typically passes through several phases: starting with cell cultures, then moving to small-animal studies, followed by preclinical trials, clinical trials, and eventually, market approval. Out of hundreds of drug candidates tested, only a very small fraction ever reach the market. This long, costly process is one reason why drug development is so expensive.

The pilot imaging ring itself consists of gamma detectors and the associated electronic circuitry. These detectors not only capture signals but also record timing information. When multiple gamma photons are detected within the same time window, the system identifies them as coming from the same event and reports a coincidence along a line of response. In this way, the line of response corresponds to the radio-tracer distribution along that line, which is essentially a line integral. The attenuation coefficient in this formulation is treated as independent and can be factored out.

slide32:

This is a big picture overview. If you look at the components in detail, multiple figures highlight the power and structure of each unit. You start with a scintillation crystal at the front of each module. The scintillation crystal's job is to convert incoming gamma photons into visible light. This visible light is then captured by a photomultiplier tube, or PMT. The crystal and PMT work together—when the scintillation crystal creates visible light, the photomultiplier tube multiplies and amplifies that light signal, converting it into an electrical current. After the signal passes through the tube, it goes on to electronic circuits that process the electrical impulse.

Depending on the amplitude of that impulse, you can determine the energy of the interacting gamma photon. One unit like this—consisting of an array of crystals and a set of PMTs—is called a "block." Several blocks combined form a "bucket," which can include hundreds of crystals. Multiple buckets are arranged together to create the full detector ring, which collects all the imaging data at once. There are additional engineering details in how these elements are assembled and how the signals are managed, but at its core, this system is designed to efficiently detect and process gamma events for reliable image reconstruction.

slide33:

Now let's talk about the desirable properties of crystal materials used in PET scanners. I'll go through these quickly. First, the crystal should have a high density, which gives it a large effective cross-section for interacting with gamma rays. Next, you want a large effective atomic number—this increases gamma-ray detection efficiency by boosting photoelectric absorption. The decay time should be short; you need the crystal to release its light energy very quickly. If the decay time is too long, then each photon event can overlap in time, causing loss in timing resolution and poor image quality.

A high light yield is also important—the crystal must convert as much gamma energy as possible into visible light, maximizing signal for efficient detection. Ideally, the wavelength of the emitted light should be near 400 nanometers, since photomultiplier tubes are most sensitive in this range, making the system more efficient. Additionally, the refractive index should be close to 1.5 to ensure optimal optical coupling between the crystal and the photomultiplier tube.

The crystal must also be non-hygroscopic, meaning it should not absorb moisture from the air. If it does, it can degrade and turn into powder, which would make it unusable for detector fabrication. All these properties are critical for reliable, high-performance medical imaging with PET scanners.

slide34:

All of these are considered desirable properties for PET detector crystals. The table lists the main candidate materials used in PET scanners. You don't need to memorize all the details or specific values here—in clinical practice, it’s more important to understand the general differences and why a particular material might be chosen for a given system.

What matters is knowing that these materials differ in their density, decay time, emission intensity, atomic number, refractive index, and whether they are hygroscopic. Each property affects how well the crystal works for detecting gamma rays and producing high-quality PET images.

slide35:

There aren't many engineering details here, but this is mainly for those working in detector development. The key components for PET detectors start with the scintillation crystals, which are organized into blocks. Each block is coupled with multiple photomultiplier tubes in the background—typically four per block. These blocks are then assembled into a complete detector ring that surrounds the patient. As the PET system operates, the crystals detect incoming gamma photons, and the photomultiplier tubes convert the light produced by these interactions into electrical signals.

The system's processing unit precisely monitors the electrical signals to determine exactly when and where events have occurred, and at what energy. The main gamma photon from a PET scan carries 511 keV energy. But if there is multiple scattering, such as Compton scattering inside the body, energy is lost and can complicate the signal. The electronics and coincidence processing unit integrate all detected events, and finally, this data is sent to the computer system.

After data acquisition, image reconstruction begins—often with methods such as filtered back projection—to produce cross-sectional images. All of these steps come together to make PET imaging possible, from crystal blocks and photomultiplier tubes to data integration and final image creation.

slide36:

The technical advancement in PET imaging has been driven largely by improvements in detector technology. Over the years, we've seen a progression from NaI detectors to BGO, and now to LSO, with each milestone year bringing significant gains in image quality.

These different generations of detectors reflect long-term innovation in materials, engineering, and system design. As you can see, each new detector technology has resulted in clearer, higher-resolution images. This slide is just for your information; I won't spend too much time on it, but it's important to recognize how advancements in detector design have directly impacted the quality of PET images throughout the history of the field.

slide37:

For a system to work at its best, noise, scatter, and random events must be estimated and removed. This slide is a nice summary. A true event means you have a single annihilation, then two photons are emitted and detected by opposite detectors, and captured by the circuit. That's what you really want—these true coincidences are the signals that contribute to accurate image reconstruction.

Scatter events have different mechanisms. Here's an example: a photon is emitted from an event, but gets scattered away before it reaches the detector. The system may report a coincidence along this scattered path, but it's not the true line of response. This isn't the real signal we need, so the scattered photon event should be removed—just like in CT, where scattering also occurs.

Random events are different. In this case, two gamma photons are generated from different events in the body, and one photon is detected in the same time window as another photon from a separate event. The system reports a coincidence, but along a line that doesn't correspond to either true event—this is not accurate, and should be avoided if possible.

In 3D PET imaging, you can see that scattered and random events become more frequent, mainly because sensitivity increases due to the wider aperture and more possibilities for 3D data acquisition. So, it's important to maximize sensitivity while minimizing scatter and random events to maintain image quality.

slide38:

How can you estimate the scattered background? This is actually simple. You set the energy window around the main photopeak. That's the main window right at the peak, and you collect counts there. Then, you also set windows on the left and right—these are the low and high-energy side windows. The idea is that you know the primary gamma photon should be at 511 keV, so its energy falls in the main window.

If there are scattered photons, their energy will be lower than 511 keV, so you'll see them in the lower energy window. If there are random events or noise, these could appear in the higher energy window. By collecting data in both the lower and higher side windows, you have a way to estimate the background scattering.

You then assign weights to the counts in the lower and higher windows—this weighted combination gives you an estimate of the scattered background in the main window. With that estimate, you subtract the scattered contribution from the main window, giving you the corrected signal. This process lets you measure and remove the scatter components, so your final image is much cleaner and more accurate.

slide39:

So this shows you the result. This is the real measurement, and you can see the scattered background as it appears in the data. When you remove the scattered background, you get these corrected measurements, or corrected reconstructions.

This demonstrates how the result changes by addressing the scattering problem. Scatter correction is a key point in improving measurement accuracy and image quality.

slide40:

And for random correction, this is a very clever idea. If you have random events being recorded, they introduce noise into your data. For usual data processing, these random events really complicate things. But we can address this purposely: for each part of the detector, we introduce a delay.

If it's a random coincidence, then the probability of the random event with this result delay is the same as without the delay. This is the key statement—this is the important idea. By introducing a delay in the coincidence detection, the probability line for incident and delayed results is the same, because the random coincidences depend on the average gamma-ray activity, which is assumed to be more or less constant over time. The delay itself doesn't change that rate.

You can then use the measured random events from the delayed window to directly estimate and subtract out the random event contribution from your true data. This approach gives you higher-quality, more quantitative results. That's the whole idea behind random correction

slide41:

There are multiple high-order effects in PET imaging, and this is what I'm showing here with these figures. Ideally, you want to detect coincidences—by definition, in a given time window, you should only have two gamma-ray photons captured. If your time window is short enough, you get a well-controlled count rate, typically up to 100,000 per second, which is the principle behind coincidence detection. However, sometimes within a short time interval, you might have three gamma-ray photons reported simultaneously. How does this happen? For example, the first annihilation event is detected near detector two, so its gamma photon reaches that detector and is recorded.

Within that same time window, another event occurs closer to detector thirteen, and its sister photon travels a longer path. Before the photon from the second event reaches detector five, the third photon from the first event might reach detector ten. So, within the same interval, three photons can be detected as coincidences.

Whenever you get triple coincidences—three detected photons in the same resolving window—they introduce confusion, and you simply discard that data. The system is designed to expect only pairs of coincident photons. If you see this higher-order event, you just throw it away, focusing only on true doubles.

But confusing situations remain: it's hard to know which events should be assigned to which lines of response. That's where the correction factor formula comes in. For every true coincidence measurement, you also need to account for the data lost due to triple or higher-order coincidences. The formula shown here calculates the correction factor—combining both the estimated loss from discarded high-order events and the true measurement—to help restore your measured data to the real level. Sometimes, higher-order events may be very rare, but if your time window is wide enough, you might encounter even more photons simultaneously, which complicates things further. In practice, these cases are always discarded, and the formula allows you to estimate the amount of lost data. You then use this correction factor to scale your results back to the true count, improving accuracy in quantitative PET measurements.

slide42:

So, graphically, you can see this process clearly. The total counts—the overall measured signal or image—actually contain two main components. These two components are scatter background and random counts, along with a variety of other effects and oddities. You can estimate the random count by using a delayed window, and triple coincidences can be corrected as well. There are many correction methods, but the most important thing to understand from this lecture is to have a rough idea of how these corrections work.

All these corrections—including attenuation correction, scatter correction, random correction, and corrections for dead time—are essential to refine the measured signal and make sure your results reflect the true activity as accurately as possible.

slide43:

The goal is to make your line integral data accurately reflect the true line integrals through the body. This allows you to reconstruct images that truly represent the concentration of radiotracer in the organ of interest. You can do this for each organ and at each time point during the PET imaging session. By plotting

these time curves, you can assess perfusion in different regions and extract physiologically relevant parameters.

This process gives you more diagnostic information and helps to interpret the results accurately. These steps—attenuation correction, scatter and random correction, and dead time correction—are all essential for quantitative PET imaging. They work together to ensure that your final images and time-activity curves are both accurate and clinically meaningful, which is the central principle behind quantitative PET analysis.

slide44:

I want to summarize PET and SPECT imaging with this slide. The key point is that the data models are different, and the underlying technology relies on gamma rays, but at very different energies. SPECT uses lower-energy gamma rays, typically around 140 keV. Because of this lower energy, you can use a mechanical collimator—high-density metals effectively block unwanted photons, shaping the beam. But in PET, the gamma ray energy is much higher, at 511 keV. Mechanical collimation simply doesn't work because the energy is too great for physical collimators, so PET uses electronic collimation instead.

With mechanical collimation, a large portion of useful gamma-ray photons is rejected, reducing sensitivity. By contrast, electronic collimation in PET does not block any useful gamma photons. This means PET sensitivity is fundamentally higher than SPECT sensitivity.

Another key feature is time information. In PET, two photons resulting from an annihilation are emitted simultaneously. The system records the precise timing, making time-resolved imaging possible and increasing sensitivity. In SPECT, we don't know precisely when the radiotracer emits its gamma photons inside the body, so there's no available time scale—just as in standard X-ray imaging, where we can't control the emission timing for each photon. Even if you try to time the arrival of X-ray photons at a detector, you don't have a proper time reference. With PET, the emission process gives you this timing, allowing sensitive and advanced time-resolved imaging. This is a crucial, elegant concept behind PET technology.

slide45:

The leading researchers in this field have developed a very large PET detector known as a whole body PET. By making the detector long enough to cover the entire body, the sensitivity is dramatically increased, allowing for even greater sensitivity than conventional PET systems.

With this approach, you can obtain a much more sensitive and comprehensive view of the entire body in a single scan. This technology offers the opportunity to capture detailed, global images of tracer distribution, potentially improving diagnostic capabilities and expanding what can be studied. It's a very interesting and innovative idea that's advancing the field of molecular imaging.

slide46:

When you perform whole-body PET and also whole-body CT, you need to consider the practical challenges and differences. Whole-body CT uses a significant amount of radiation, since it scans the entire patient with X-rays. In contrast, with whole-body PET, as shown in the figure, there is no space to easily position the X-ray components completely around the patient in the same way as traditional CT.

However, researchers are actively thinking about creative solutions, and we actually encourage students to work on the possibility of achieving full-body CT simultaneously with PET, even when scanners are large enough to scan the whole body at once. This might require using new, out-of-the-box ideas—possibly even involving MRI technology—instead of conventional techniques. These are exciting areas for innovation and research, and we can discuss some of those cutting-edge possibilities later on in the course.

slide47:

Another point I want to highlight, as indicated by the green button here, involves using multiple pinholes with a scintillation camera. Traditionally, you would use a single pinhole to capture one image at a time. But in this pioneering work, as shown in the paper, they introduced the concept of using several pinholes—seven in this case. Each pinhole forms a distinct image, producing multiple views of the target simultaneously; for example, you get a double-cone image from one pinhole's geometry, and with seven pinholes, you get seven images at once.

This approach was used to image the heart, allowing for seven simultaneous views, which improved accuracy by acquiring several perspectives at a single time point. Although the reconstructions from those seven views weren't fully impressive with the technology available back then, the idea remains promising. There is real potential: if you collect seven views from one direction and another seven from an orthogonal direction, then with modern signal processing, you may achieve image reconstructions far superior to those early results. This opens up interesting possibilities for future multi-view tomography and advanced image reconstruction techniques.

slide48:

Anyway, so this is just for your information. The slide summarizes key concepts for SPECT and PET imaging: statistical distributions, how data is modeled with single and paired emissions, various reconstruction methods including attenuation compensation, deterministic, and statistical reconstruction, as well as the architecture of SPECT and PET scanners and scatter correction.

Finally, it mentions the fully 3D approaches and highlights the importance of the conference and community in advancing this field.

slide49:

And we had a fully 3D meeting in China. This conference is dedicated to advances in fully three-dimensional image reconstruction in radiology and nuclear medicine, bringing together experts, researchers, and students from around the world to share the latest developments and ideas.

It's a valuable opportunity to connect with the community and stay current in this rapidly evolving field.

slide50:

So this is the website for the Fully3D 2017 meeting, which was held in Xi'an, China from June 18 to 23, 2017. If you are interested, you can take a look for more information about the conference, its location, and the

history of Xi'an as China's oldest ancient capital and the starting point of the Silk Road. It's also home to the famous Terracotta Army, known as the eighth wonder of the world.

slide51:

And in the center of the city, there is this Bell Tower. Every morning, the bell would be rung so life would be synchronized. This is the heart of the city. We enjoyed very much the tour during the conference.

slide52:

So the interesting thing is that I designed a logo specifically for the Fully 3D meeting. The logo is modeled after the iconic Bell Tower in the city, making it a distinctive symbol that connects the event to the cultural heart of Xi'an.

slide53:

Why does a logo look like this? The reason is clear: the meeting covers the full spectrum of CT, SPECT, and PET reconstructions. The logo visually represents these modalities together. For CT, you see the cone-beam geometry at the top, which is standard in modern CT scanning.

For SPECT, there's a model with a collimator, a layer of crystals, and photomultiplier tubes—these components are essential for forming SPECT images. At the base is the PET detector ring, which is fundamental for PET. By including these technical elements, the logo symbolizes the integration of all three imaging modalities, reflecting the focus and breadth of the Fully 3D meeting.

slide54:

Also, you have this movie of the 3D logo that brings together all the elements for Fully3D'17 and represents the spirit of the meeting.

slide55:

For today's homework, I have listed five problems. You are welcome to do more if you wish.

The first problem is about isosensitive imaging, which I briefly introduced earlier. It involves taking two views and compensating in a specific way. The key idea is that if you assume a simple binary case with only one source distribution, you can multiply the views to achieve compensation. However, remember that this is not a universal solution—it applies only under limited assumptions.

The second question asks how many total counts are necessary to achieve 1 percent uniformity in a SPECT image, given a 128 by 128 data matrix. This problem is fairly straightforward. You can solve it using the Poisson distribution along with basic sensitivity concepts. The answer should follow directly from these principles.

The third exercise focuses on reducing statistical noise in an image. Consider what happens if you double the total imaging time, and separately, what happens if you double the mass of tracer injected. In each case, determine by what factor the noise in the image would be reduced.

The fourth problem asks you to compute the energy of gamma photons produced from positron emission. The correct result should be 511 keV. To calculate this, apply the formula E equal m c square If needed, look up the rest mass of the electron and positron.

Finally, the fifth question addresses artificially high levels of radioactivity observed in PET lung images. You are asked to suggest one possible mechanism by which this effect could occur. These exercises are designed to help you practice applying statistical principles, imaging physics, and basic quantitative reasoning to nuclear medicine problems.