slide1:

Welcome to Lecture 16 on Nuclear Physics.

In this course, we focus on understanding nuclear physics in the context of biomedical imaging. The key to learn this material is to go through it more than once — first by following the lecture, and then by reviewing the textbook or slides afterward. Repetition helps clarify the concepts and makes them easier to remember.

Now, when you open the textbook chapter on nuclear medicine or nuclear imaging, you might feel that it contains an overwhelming amount of information. There are details about pharmaceuticals and diseases that may seem excessive. For our purpose, you don't need to memorize those. Instead, we will focus on the physical principles, including the chemistry and instrumentation concepts, essential components, and particularly the mathematical modeling and quantitative relationships.

These are the core concepts you should master well. The clinical terms and long lists are good to know, but not our main focus.

Throughout the slides, you'll notice visual markers such as red diamonds or green buttons. Again, these are meant to highlight what is most important and to guide your attention to the central knowledge in this lecture."

slide2:

As you can see from the course schedule, we are right on track. We've already covered the earlier topics such as introduction, system concepts, Fourier series, and signal processing. More recently, we completed modules on CT reconstruction and MATLAB applications.

Today's lecture is focusing on nuclear physics. In the next lecture, we will explain SPECT and PET imaging principles and systems. After that, we will move on to MRI, ultrasound, and optical imaging, before wrapping up the course with the final exam.

So everything is proceeding according to plan.

slide3:

Now we begin the section on nuclear physics, which marks the start of our discussion on nuclear imaging as an imaging modality. As with every modality we cover, I'll first give you a general perspective — a big picture introduction to provide background and context. After that, we'll focus on the key elements specific to this modality.

In this course, I want you to always pay attention to the outline. The outline is your roadmap. It shows what's important and helps you remember the structure of the topic. As in earlier lectures, I use visual markers like the red diamond and the green button to highlight the most essential points. These are your hints for where to focus.

Now, some of you may feel concerned about prerequisites, such as calculus or differential equations. Don't let that discourage you. Today, it's easier than ever to look up mathematical concepts — a quick search can bring up explanations and definitions right away. What matters most is your willingness to learn. This field is

highly interdisciplinary, and your goal should be to explore, connect, and build your own understanding of the knowledge you need.

So, as we step into nuclear physics, remember: our approach is not to memorize every detail, but to see the story. First, we build a general understanding, then we dive into specific components like radioactivity, tracer production, and data acquisition. Once you see the big picture, the details fall into place more naturally. Nuclear imaging is a fascinating field, and I hope you'll enjoy this journey as much as I do.

slide4:

Let's begin with a bit of history to set the stage. Every imaging modality we study has been shaped by many talented researchers whose discoveries were recognized at the highest levels. Nuclear imaging is no exception. If you look back at the milestones in its history, you'll find several famous names, and among the most remarkable is Marie Curie.

Marie Curie was awarded two Nobel Prizes for her groundbreaking work on radioactivity. In fact, element ninety-six, called curium, was named in her honor. Her influence extended to the next generation as well — her daughter Irène became a distinguished nuclear physicist in her own right and received a Nobel Prize in 1935. Her other daughter, Ève, wrote a highly acclaimed biography of her mother.

So, when we talk about nuclear imaging today, we are building on a legacy that began with pioneers like Marie Curie. Their contributions laid the foundation for our understanding of radioactivity, which is central to this entire field.

slide5:

This is Michael Ter-Pogossian, a very important figure in nuclear physics and especially in nuclear imaging. He is recognized as one of the pioneers of positron emission tomography, or PET, which became the first truly functional brain imaging technology.

PET was groundbreaking because, unlike conventional CT, which shows anatomy, PET allows us to evaluate the brain in action — during mental processes. Later in this course, we will study PET in detail.

I'd like to share a personal connection here. After finishing my PhD, I was hired at the Mallinckrodt Institute of Radiology at Washington University in St. Louis. I started as a medical physicist, doing imaging measurements and research. The leader of our medical physics group at that time was Professor Michael Ter-Pogossian. I was still young and didn't fully realize how great he was — I simply saw him as my immediate supervisor.

Looking back, I recognize his remarkable contributions. Many people felt that he truly deserved a Nobel Prize for his invention of PET and his pioneering results. Unfortunately, he passed away in 1996 from heart disease. Beyond his scientific work, he was also the founding editor-in-chief of the IEEE Transactions on Medical Imaging, which remains the most important journal in our field.

I remember him as not only brilliant but also humorous and down-to-earth. For example, he once told us that doing too much simulation is like having a simulated meal — it might look convincing, but afterward, you're still hungry. That kind of humor made him unforgettable.

slide6:

Here we see X-ray CT, which you are already familiar with from earlier lectures. Let me give you a quick review.

In conventional CT, the system typically uses an energy-integrating detector. This means that all of the incoming X-ray photons are combined, and the result is an averaged measurement over the entire X-ray spectrum. That gives us the grayscale images you know well, where tissues appear in shades from black to white.

Now, with more advanced technology, we can use photon-counting detectors. These detectors are energy-sensitive, which means they can record not only the number of photons but also their energies. By sorting the detected photons into many energy bins, we can see the full X-ray spectrum, including characteristic features called K-edges for different materials.

This approach allows us to separate and identify materials such as gold, iodine, calcium, and fat. In other words, instead of just a grayscale image, we can create color-coded images that reflect the molecular composition of the tissue or phantom. This is called spectral, or molecular, CT.

So, with photon-counting CT, we move beyond structural imaging into the possibility of functional and molecular imaging — a very exciting direction for the future.

slide7:

Today, we move into a very different story — the basic idea of nuclear imaging. This principle is absolutely fundamental.

Just like in X-ray imaging, we are still working with radiation. But here's the key difference: in nuclear imaging, the radiation does not come from an external source outside the body. Instead, it comes from inside the patient, because we introduce a radioactive tracer into the body. That tracer emits gamma rays.

Now, gamma rays and X-rays are very similar. Their energy ranges overlap, although gamma rays are usually considered to be slightly more energetic. Both are forms of electromagnetic radiation, part of the same spectrum. Importantly, like X-rays, gamma rays can penetrate the body.

But if we simply detect gamma rays without direction, we only get a collection of signals with no spatial meaning. To form an image, we need to know where each photon came from. This is why we use a collimator — such as a pinhole collimator or a parallel-hole collimator. The collimator ensures that only photons traveling along certain directions are recorded, giving us the directional information we need for tomographic imaging.

After passing through the collimator, the gamma rays strike a scintillation crystal. This crystal converts the high-energy gamma photons into flashes of visible light. Next, photomultiplier tubes take that light and convert it into electrical signals, amplifying the signal dramatically.

Those electrical signals are then digitized and sent to a computer. One measurement gives us a single projection. By collecting many projections from different angles, we can reconstruct a tomographic image.

It's important to understand what the resulting image represents. It is not showing us anatomical structure, like CT does. Instead, it shows the distribution of the radioactive tracer inside the body. That means the

image reflects physiology — the function of tissues — rather than just anatomy. This is the key and very powerful idea behind nuclear imaging.

slide8:

This slide further illustrates the principle of nuclear imaging. Here we see the use of a parallel-hole collimator. The purpose of the collimator is to make sure that the signals we detect are aligned with the holes. That way, any photon recorded must have come from directly beneath that hole, and this gives us the spatial information we need. Without it, we would not know the photon's origin.

Once the gamma photon passes through the collimator, it hits the scintillation crystal. The crystal converts that high-energy photon into visible light. Then the photomultiplier tubes convert the light into an electrical signal, which can be read by the computer for image formation.

Now, notice what happens after a radioactive tracer is injected into the bloodstream. The tracer travels through the body and tends to accumulate in certain locations, depending on the physiology. In this example, for cardiac imaging, the tracer concentrates in the heart muscle. This allows us to see how it penetrates the cardiac chambers and the muscle wall.

The resulting image reflects the distribution of the tracer. Since the tracer participates in biochemical reactions, nuclear imaging provides insight into physiology — the function of tissues — rather than just anatomy.

This makes nuclear imaging highly complementary to X-ray imaging. While X-ray imaging uses an external source and shows structure, nuclear imaging relies on an internal radioactive tracer and shows biological function. Together, they give us a much more complete picture.

slide9:

The physical principles of nuclear imaging can be summarized as follows.

First, we use an unstable isotope as a tracer molecule, also called a radionuclide. This tracer can be administered in different ways — most commonly through intravenous injection, but in some cases, it can also be taken orally, for example, by drinking a small amount of radioactive liquid.

Once the tracer enters the body, it participates in normal metabolic processes. As it does so, it emits gamma rays from inside the body. These gamma rays are then detected, and the resulting measurements reflect the distribution of the tracer.

What makes this so powerful is that the signal corresponds to metabolism and physiological function, rather than just anatomical structure. In other words, nuclear imaging allows us to see how tissues are working, not just how they look.

This sets the stage for us to compare nuclear imaging with X-ray imaging, and to highlight both the similarities and the important differences between them.

slide10:

Here, we compare X-ray CT with nuclear imaging by looking at their similarities and differences.

On the left, you see a CT scanner and CT images. On the right is a PET scanner, which represents nuclear imaging. Although the images shown here may not be perfectly matched, the idea is clear — both modalities use radiation and produce tomographic images.

Let's begin with the similarities.

First, gamma rays and X-rays have overlapping energy ranges. Their energies are high enough that, when passing through biological tissue, we can often treat them as traveling in straight lines. Of course, scattering occurs, but for image reconstruction, we assume straight-line geometry. That simplifies tomographic reconstruction.

Second, the data from both modalities can be expressed as line integrals. In CT, this is exactly the Radon transform. In nuclear imaging, it is a related but slightly different mathematical formulation. Still, both methods involve line integrals and support tomographic image reconstruction.

Third, both CT and nuclear imaging generate two-dimensional slices, three-dimensional volumes, or even four-dimensional datasets when dynamics are included. So, in terms of imaging geometry and reconstruction, they share a strong common foundation.

Now, let's turn to the differences. In CT, the radiation source is external — the X-ray tube outside the patient. In nuclear imaging, the source is internal — a radioactive tracer introduced into the bloodstream. This tracer accumulates in certain regions, such as the heart or tumors, especially malignant tumors with a high blood supply.

Another difference is the type of information provided. CT primarily gives anatomical information — structural details of organs and tissues. Nuclear imaging, by contrast, provides functional information, showing how tissues are metabolizing or reacting biologically. This makes nuclear imaging highly complementary to CT.

Third, there is a difference in flux and resolution. CT uses a strong external X-ray beam, producing high photon flux and therefore high-resolution images. Nuclear imaging uses only a small injected amount of tracer, so the signal is weak. The data are noisier, and the spatial resolution is lower. Nuclear images may not look as sharp as CT images.

Finally, there is sensitivity. CT detects differences in tissue density, but since most tissues are made of light elements, benign and malignant tumors can appear very similar until they grow large enough to change anatomy. That means CT may detect cancer only at later stages. Nuclear imaging, however, is far more sensitive. Because the tracer accumulates in areas of abnormal metabolism, nuclear imaging can reveal very small or early-stage tumors — often before structural changes appear.

This sensitivity is one of the greatest strengths of nuclear imaging, making it an invaluable tool for early detection and treatment planning.

slide11:

CT and nuclear imaging are highly complementary. CT provides detailed anatomical information, while nuclear imaging shows functional and biochemical processes. So the natural idea is — why not combine them?

That's exactly what happened with the development of hybrid scanners. One of the pioneers in this effort was Simon Cherry, who helped establish PET-CT as a practical technology.

The major strength of nuclear imaging is its ability to label many types of metabolites. This makes it highly sensitive for studying biochemical and physiological processes inside the body. The limitation, of course, is that nuclear imaging by itself has relatively low resolution. That's why it works best when combined with CT or MRI.

CT or MRI gives us the anatomical background — the structural map of tissues. The nuclear tracer distribution can then be precisely superimposed on that background, maximizing the information content. This combination allows us to see both anatomy and function in perfect alignment.

PET-CT was the first successful hybrid scanner. Today, in oncology departments, it is rare to see a stand-alone PET scanner. Almost always, PET is integrated with CT. This avoids the problems of image registration that occur when separate scanners are used. Moving the patient between machines can change posture or positioning, making alignment difficult. With a single integrated scanner, those issues are eliminated.

Later, PET-MRI was also developed, bringing together molecular imaging with the excellent soft-tissue contrast of MRI. And researchers have even envisioned the possibility of a "tri-modality" scanner — PET, CT, and MRI all in one system. That would give us an even more powerful tool for comprehensive imaging.

slide12:

We have now finished the general perspective, so let's move on to radioactivity.

Radioactivity is explained in several sections of the textbook. The reading itself is not especially difficult, but the terminology can feel confusing, especially if you don't have a strong background in chemistry. I've had to go through these materials multiple times myself. Over time, it becomes clearer — what we are really dealing with are physical and chemical phenomena at the atomic level.

So why does radioactivity occur? The reason is that certain isotopes are unstable. These isotopes undergo decay, releasing radiation in the process. The decay can follow different mechanisms, which we will study in detail.

Importantly, radioactive decay follows an exponential curve. This behavior can be modeled mathematically using ordinary differential equations. Once you see the formulation, it's not as difficult as it sounds. With some practice, the mathematics will become straightforward, and it provides a precise way to describe radioactive processes.

slide13:

This slide is mainly a review of basic nuclear physics concepts: the atomic number and the mass number.

The atomic number, denoted by Z, is simply the number of protons in the nucleus of an atom. The mass number, usually written as A, is the total number of nucleons, meaning the number of protons plus the number of neutrons. In most cases, the atomic weight is approximately equal to this mass number.

For example, in helium, the atomic number is 2, because there are two protons. The mass number is 4, because helium has two protons and two neutrons. When an atom is neutral, the number of electrons surrounding the nucleus matches the number of protons in the nucleus. This keeps the charge balanced.

Now, once we understand atomic number and mass number, we can define what isotopes are.

slide14:

Now let's talk about isotopes.

Isotopes of a chemical element all share the same atomic number — meaning they have the same number of protons. But they differ in their mass numbers because the number of neutrons is different. So, isotopes are the same element, but with different neutron counts.

Inside the nucleus, protons and neutrons are held together by the strong nuclear force. When the balance between protons and neutrons is within a stable range, the nucleus remains stable. But if the configuration becomes unbalanced — for example, when there are too many neutrons or too many protons relative to each other — the nucleus can become unstable.

An unstable nucleus will eventually change into a more stable state. This happens through radioactive decay, which follows an exponential curve. We'll explore that decay law in more detail and even derive it mathematically later on.

slide15:

Radioactive isotopes can emit gamma rays. Let's look more closely at what that means.

On the electromagnetic spectrum, X-rays and gamma rays actually overlap in terms of energy. X-rays typically fall within a certain energy range, while gamma rays can extend both lower and higher. In medical imaging, gamma rays are often in the range of hundreds of kilo-electronvolts. For example, 511 kilo-electronvolts is a key energy level used in positron emission tomography, or PET. At lower energies, gamma rays are used in single-photon emission computed tomography, or SPECT.

So in terms of energy, X-rays and gamma rays are quite similar. The main difference lies in how they are generated. X-rays are produced by electron interactions — for instance, when high-energy electrons strike a tungsten target, generating X-rays in the process. In contrast, gamma rays are produced by transitions inside the nucleus itself. When an unstable nucleus changes from a higher energy state to a lower energy state, it releases that excess energy in the form of a gamma photon.

This distinction is important: X-rays come from interactions involving electrons, while gamma rays come directly from nuclear transitions. Both are powerful tools in imaging, and as we move forward, we'll see in more detail how gamma rays are generated and used.

slide16:

Radioactive isotopes can decay through several different mechanisms, and it's important to understand the main ones.

In your textbook, these are grouped into four basic types. The first is alpha decay, the second is beta decay, which includes both beta-minus and beta-plus, or positron, emission — the third is gamma decay, and the fourth is electron capture.

An easy way to remember them is simply in order: alpha, beta, gamma, and then electron capture. Each of these processes describes a different way in which an unstable nucleus transforms into a more stable state, releasing radiation in the process.

We'll go into detail about each type of decay in the following slides.

slide17:

Let's begin with alpha decay.

In alpha decay, an unstable isotope transforms into a daughter nucleus and emits an alpha particle. An alpha particle is essentially the nucleus of a helium atom — two protons and two neutrons bound together.

Alpha particles carry a relatively high energy, typically in the range of 3 to 7 mega-electronvolts. Because of this, they can cause severe tissue damage. However, they cannot travel far — usually only a few millimeters in biological tissue.

For this reason, alpha radiation is not useful for imaging. It doesn't penetrate deeply enough to form images. Instead, it has therapeutic applications. If we can direct alpha-emitting isotopes precisely to cancer cells, they can deliver very localized, highly destructive radiation that kills the tumor while sparing surrounding tissue.

So, while alpha decay is the first mechanism of radioactive decay we study, it is more relevant to radiation therapy than to nuclear imaging.

slide18:

The second major type of decay is called beta decay. There are two forms: beta-minus decay and beta-plus decay.

In beta-minus decay, a neutron inside the nucleus converts into a proton, releasing an electron — which we call a beta-minus particle — along with an antineutrino. This process changes the nucleus into a different element. Beta-minus particles can be damaging to tissue, and in some cases, this type of decay is used for therapy.

In beta-plus decay, a proton in the nucleus converts into a neutron, releasing a positron — which is a positively charged electron — along with a neutrino. The positron does not remain in the body for long. It quickly encounters a nearby electron, and when the two meet, they annihilate each other. This annihilation produces a pair of gamma photons.

These gamma photons do not emit randomly. They travel in opposite directions, almost exactly 180 degrees apart. This property is the basis for positron emission tomography, or PET. By detecting these pairs of photons and knowing they travel in straight, opposite directions, we can reconstruct images of where the tracer accumulated in the body.

Some textbooks separate positron emission from beta decay and list it as a distinct process. Others group it under beta-plus decay. Either way, the important point is this: for imaging purposes, beta-plus decay and gamma decay are useful, while alpha decay and beta-minus decay are not. PET, in particular, relies on beta-plus decay and positron annihilation.

slide19:

Positron emission is another key decay mechanism, and it is especially important for nuclear imaging.

Here's what happens: inside an unstable nucleus, a proton transforms into a neutron. In the process, the nucleus emits a positron — which is a positively charged electron — along with a neutrino. The emitted positron travels only a short distance before it encounters an electron. Because they carry opposite charges, the positron and the electron are naturally attracted to one another. When they meet, they annihilate each other. This annihilation produces two gamma photons, each with an energy of 511 kilo-electronvolts. These photons travel in opposite directions, almost exactly 180 degrees apart.

This property is the basis for positron emission tomography, or PET. By detecting pairs of gamma photons and knowing they must lie along the same line, we can reconstruct the location where the positron annihilation occurred. That gives us an image of the tracer distribution inside the body.

So, compared with alpha and beta-minus decay, positron emission is extremely valuable for imaging — it is the physical principle behind PET, one of the most important functional imaging modalities in medicine.

slide20:

Now let's look at gamma decay.

Gamma decay occurs when a nucleus changes its internal configuration and releases the excess energy in the form of a gamma photon. This process does not change the number of protons or neutrons, so the element remains the same — only the energy state of the nucleus changes.

There are two main categories. The first is immediate gamma decay. In this case, the nucleus drops from a higher energy state directly to a lower state, and a gamma photon is released right away.

The second category is more common and more useful in practice. Here, the nucleus does not release the gamma photon immediately. Instead, it first enters a metastable state. After a short delay, the gamma photon is released, sometimes along with other particles such as an antineutrino.

A classic example of this process is technetium-99m, often abbreviated as Tc-99m. The "m" stands for metastable. This isotope is widely used in nuclear medicine because of its convenient half-life and the gamma photons it emits, which are ideal for imaging.

So, gamma decay is highly relevant to nuclear imaging, especially in the form of isotopes like Tc-99m, which remains one of the most important tracers in clinical practice.

slide21:

The textbook explains gamma decay in more detail, and it makes an important point: no radionuclide decays solely by gamma emission. Instead, gamma rays appear as part of a decay scheme where an intermediate state — called a metastable state — is formed.

A widely used example is technetium-99m, or Tc-99m. It is produced from molybdenum-99, which undergoes beta decay. The daughter nucleus, Tc-99m, is metastable. After a half-life of about six hours, it transitions to the stable form Tc-99, releasing a gamma photon in the process.

The gamma photon from Tc-99m has an energy of about 140 kilo-electronvolts. This energy is ideal for nuclear medicine imaging. If the photon energy is too low, below about 100 keV, most of the photons are absorbed in tissue and never reach the detector. If the energy is too high, above about 200 keV, the photons penetrate the collimator and reduce image quality. That is why the sweet spot for imaging is between 100 and 200 keV — and Tc-99m fits perfectly in this range.

This is why Tc-99m has become the most widely used radionuclide in nuclear medicine, applied in more than 90 percent of diagnostic imaging studies. So now you've seen the three main decay mechanisms: alpha, beta, and gamma. In the next step, we'll discuss the final one — electron capture.

slide22:

The final type of decay we'll discuss is electron capture.

In electron capture, a proton in the nucleus is transformed into a neutron. This happens when the nucleus captures one of the atom's own inner electrons, usually from the K-shell. By capturing this electron, the proton combines with it to form a neutron, and a neutrino is released.

The daughter nucleus is often left in an excited state. To release that excess energy, it emits X-rays or Auger electrons. In some cases, additional radiation is produced through what we call bremsstrahlung — a German word meaning "braking radiation." This process is similar to what we studied earlier in X-ray production, where electrons slow down and emit radiation.

So, electron capture does not directly produce a particle like an alpha or a beta. Instead, it changes the balance inside the nucleus and produces photons — either X-rays or gamma rays — along with other possible emissions. This mechanism adds to the classification of decay types, and it highlights again that photons can be generated in different ways. Whether they are called X-rays or gamma rays depends not on the photon itself, but on how it was produced.

slide23:

Here is another way of classifying radioactive decay. In total, you can think of five types.

First, alpha decay, which typically occurs among the heavier elements. Second, beta-minus decay. Third, positron emission, or beta-plus decay. Fourth, electron capture. And finally, spontaneous fission.

For our purposes, the first four are the most important. Spontaneous fission is more of a nuclear physics phenomenon and is not especially relevant to medical imaging.

So, depending on which textbook you consult, you may see these processes grouped slightly differently. But for this course, just keep in mind the essential ones: alpha, beta, gamma, and electron capture. Within beta decay, remember there are two cases — beta-minus and beta-plus.

With that, you now have a solid overview of the main decay mechanisms we will encounter in nuclear imaging.

slide24:

This table lists some of the most common radioactive tracers used in nuclear medicine. Each tracer has a characteristic half-life and a specific gamma-ray energy.

For example, technetium-99m has a half-life of about six hours and emits gamma rays at 140 kilo-electronvolts, which is ideal for imaging. Gallium-67 has a longer half-life of about three days and multiple gamma energies. Thallium-201, xenon-133, indium-111, and iodine isotopes are also widely used, each with its own half-lives and energy levels.

You don't need to memorize the details of this table. Instead, just recognize that different radionuclides are available for different imaging applications, and their properties — half-life and gamma energy — determine how and where they are best used. This is general knowledge to give you a sense of the variety of tracers available in nuclear medicine.

slide25:

(No notes for this slide)

slide26:

Now let's look at the mathematical details of radioactive decay.

The key point is that the rate of change of the number of radioactive nuclei, which we call capital N, is proportional to how much material is still present.

In symbols, this is written as:

"dee N by dee t equals negative lambda times N."

Here, lambda is called the decay constant. The negative sign simply tells us that the number is going down over time.

If you solve this equation, the solution is an exponential curve. It looks like this:

"N of t equals N naught times e to the power of negative lambda times t."

Here, N of t is the number of nuclei at time t, and N naught is the starting number at time zero.

Now, the half-life — written as t one-half — is the time it takes for the material to decay to one-half of its original amount. To find it, we set N equal to one-half of N naught in the formula. If you solve for time, you get:

“t one-half equals natural log of two, divided by lambda.”

So the half-life depends only on the decay constant.

But in medicine, we also have to think about biological clearance. When we inject a radioactive tracer, the body tries to remove it — for example, through urine. That biological process also follows an exponential law, with its own half-life.

When we put both processes together — the physical decay and the biological clearance — we get the effective half-life.

The formula is:

“One over effective half-life equals one over the physical half-life plus one over the biological half-life.”

This is the equation you really need to remember. It tells us how long a tracer effectively stays in the body, and that is the number we use in nuclear medicine calculations.

slide27:

Now we move on to the third part of our discussion — tracer production.

So far, we’ve gone through the general perspective, and we’ve talked about isotopes and how they decay through different mechanisms. Some of these mechanisms, like positron emission or gamma decay, produce gamma rays that we can detect for imaging purposes. We also saw how radioactive decay can be modeled mathematically.

Now let’s ask the next question: how do we actually produce radioactive tracers?

One of the most practical, popular, and cost-effective methods is through something called a radionuclide generator. These generators are extremely important because they are the most widely used way to produce tracers for clinical nuclear medicine.

The process is often called “milking.” You’ll soon see why this analogy makes sense. We can even describe this process mathematically. It looks a bit more complicated than simple decay, but really, it’s just a system of two equations — one for the parent isotope and one for the daughter isotope. When you couple them together, you can model how tracers are produced over time.

That’s what we’ll explore next.

slide28:

Earlier, we talked about the four main mechanisms of radioactive decay: alpha decay, beta decay, gamma decay, and electron capture. Within beta decay, remember there are two types — beta minus and beta plus. And for imaging purposes, beta plus is the more important one.

Now let’s switch perspective. Instead of asking how isotopes decay, we ask how do we produce these radioactive tracers in the first place? The good news is, this is also easy to remember — there are four basic methods.

According to your textbook, the four ways to produce radionuclides are:

Neutron capture.

Nuclear fission.

Charged-particle bombardment.

And the use of radionuclide generators.

So again, four mechanisms for decay, and four methods for production — a nice symmetry that helps you remember.

Among these, I want to highlight the fourth method — radionuclide generators — because it is the one most widely used in clinical practice. The first two methods require a nuclear reactor, and while they are very important in research and industrial production, they are not as accessible for everyday hospital use. Between those two, nuclear fission is now more popular and more cost-effective compared to neutron capture.

But for us, as medical imaging scientists and engineers, the radionuclide generator is the key player, and that's the one we'll explore in detail next.

slide29:

Now let's take a closer look at the first two methods: neutron capture and nuclear fission. Both of these require a nuclear reactor.

A nuclear reactor is essentially a device that initiates and controls a sustained chain reaction. In a typical power plant, this chain reaction produces heat, which is then transferred into a working fluid like water or gas, and eventually used to generate electricity.

But for our purpose in nuclear medicine, reactors play a different role — they are used to produce radionuclides. In neutron capture, a nucleus absorbs an extra neutron and becomes unstable, creating a useful isotope. In nuclear fission, a heavy nucleus splits into smaller nuclei, and among the products are medically important isotopes such as molybdenum-99, which is the precursor of technetium-99m.

Between the two, nuclear fission is more popular today because it is more efficient and cost-effective than neutron capture. Still, both methods remain central to the production of medical tracers — though they are limited by the need for specialized, large-scale reactor facilities.

slide30:

Let's now underline the second method, which is nuclear fission. The first method, neutron capture, is less commonly used today, so we'll focus more on fission.

What happens in nuclear fission? The nucleus of a heavy atom, such as uranium-235, absorbs a high-energy particle like a neutron. This creates an unstable excited state. As a result, the nucleus splits into two lighter nuclei, along with a few free neutrons, gamma photons, and a large release of energy.

On the right-hand side, you can see the process illustrated. Uranium-235 absorbs a neutron and briefly becomes uranium-236. That unstable nucleus then splits into two smaller fragments, krypton-92 and

barium-141, along with additional neutrons. Those free neutrons can go on to trigger further reactions, creating a chain process.

This principle is very powerful. It's the same physics behind nuclear power plants and, unfortunately, also nuclear weapons. In our context, though, fission is important mainly because it can be used to produce medically useful radionuclides, such as molybdenum-99, which later decays into technetium-99m for imaging.

That said, in this course, we won't emphasize the reactor-based methods too heavily. Instead, we'll focus more on the radionuclide generator, since that is what you are most likely to encounter in clinical imaging practice.

slide31:

The third method of producing radionuclides, and one that is especially important for nuclear imaging, is the use of a cyclotron. A cyclotron is a type of particle accelerator. It works by accelerating charged particles—such as protons—outward in a spiral path under the influence of a strong magnetic field and a rapidly varying electric field.

When these high-energy particles strike a target material, nuclear reactions occur, and radionuclides are produced. One of the most well-known examples is the production of fluorodeoxyglucose, or FDG, which is widely used in PET imaging.

Because PET imaging requires very short-lived tracers like fluorine-18, cyclotrons must be located relatively close to hospitals and imaging centers. This is one of the reasons PET imaging is more costly: it not only requires advanced scanners but also a dedicated cyclotron facility to generate the necessary tracers.

So, while nuclear reactors play a role in producing some radionuclides, cyclotrons are indispensable for producing the positron-emitting isotopes used in modern clinical imaging.

slide32:

Here you see a graphical illustration of a modern cyclotron system. On the right is the cyclotron itself, the particle accelerator that produces the radionuclides. In the middle is the biosynthesizer, which is used to process the raw radioactive material into a usable radiotracer—such as FDG for PET imaging. On the left is the computer terminal, where the entire operation is monitored and controlled.

This layout highlights the workflow: the cyclotron generates the radionuclide, the biosynthesizer prepares it into an injectable tracer, and the computer terminal ensures that everything runs safely and precisely. Together, these components allow us to produce short-lived tracers on demand for clinical use in nuclear imaging.

slide33:

The fourth method of producing radionuclides, and the one most relevant to clinical imaging today, is the use of radionuclide generators.

A very well-known example is the technetium-99m generator. This system is sometimes nicknamed the “moly cow” because it is based on the parent isotope molybdenum-99, which naturally decays to technetium-99m. Just as a cow can be “milked” repeatedly, the generator can be eluted, or “milked,” each day to obtain fresh technetium-99m for imaging.

Technetium-99m is especially important because it is the most widely used radionuclide in single photon emission computed tomography, or SPECT. It provides the right balance of half-life, about six hours, and gamma ray energy, about 140 kilo-electron volts, making it very effective for diagnostic imaging while keeping the patient dose reasonable.

This approach is extremely cost-effective compared to building a nuclear reactor or cyclotron. That is why radionuclide generators remain a backbone of nuclear medicine practice worldwide.

slide34:

Let us now look more carefully at how the technetium generator actually works. This system operates through a two-step process.

The parent isotope is molybdenum-99, which is prepared and loaded into an alumina column inside the generator. This column is shielded with lead for safety. Over time, molybdenum-99 naturally decays to technetium-99m, the daughter isotope. Technetium-99m, however, binds only weakly to the alumina column. That means if we pass a saline solution through the column, it will wash the technetium-99m out, while leaving the molybdenum-99 behind. This daily washing process is often called “milking” the generator, just like milking a cow. Each elution provides a fresh supply of technetium-99m for clinical use.

Once obtained, the technetium-99m is injected into the patient’s body. Inside the body, it undergoes further decay to technetium-99, a stable isotope. Importantly, this decay is accompanied by the emission of gamma photons, which can then be detected by imaging devices such as a gamma camera.

From a modeling perspective, this is a two-step decay chain. The parent, molybdenum-99, decays with its own half-life to produce technetium-99m. Then technetium-99m decays with a shorter half-life into stable technetium-99, releasing gamma radiation in the process. Each of these steps follows an exponential law, and when combined, they produce a dynamic interplay between parent, daughter, and final stable product.

This two-step process is what makes the technetium generator so effective—it provides a steady supply of a short-lived imaging isotope without requiring a reactor or cyclotron on-site.

slide35:

Now, let us connect the generator process to mathematical modeling. To properly describe the production and decay of technetium-99m, we need three first-order differential equations. These equations track the parent isotope molybdenum-99, the daughter isotope technetium-99m, and finally the stable granddaughter technetium-99.

The first equation describes the parent: it decays exponentially at a constant rate, as we already studied. The second equation describes the daughter, technetium-99m. This one is more interesting because it has two contributions: one positive term, which comes from the decay of the parent feeding into the daughter,

and one negative term, which accounts for the daughter decaying further. Finally, the third equation describes the granddaughter, which only has a positive contribution coming from the decay of the daughter.

If you look closely, you will see that the daughter equation is not purely exponential—it is influenced by both the parent and its own decay. To solve this type of system, we use standard methods for ordinary differential equations. Specifically, the daughter solution is expressed as the sum of two parts: a homogeneous solution, which describes the natural decay, and a particular solution, which accounts for the contribution from the parent. These two components combine to give the real physical solution.

The important point here is not to get lost in the algebra, but to recognize the principle: exponential decay chains can be modeled systematically, and by solving these equations, we can predict the time course of technetium-99m production and availability in a generator. This is crucial in practice, because it tells us how much tracer will be available at a given time and when to perform elution for maximum yield.

slide36:

Now let's make sense of the daily yield from a radionuclide generator.

On the left graph, the dashed line shows the parent isotope—molybdenum-99—which slowly decays over time. The solid line shows the daughter isotope—technetium-99m—which gradually builds up as the parent decays. At first, the daughter accumulates rapidly, but eventually, a balance is reached. Why? Because as more daughter atoms appear, they also begin to decay at a faster rate. At some point, the rate of production from the parent equals the rate of decay of the daughter. That balance defines the maximum activity of the daughter.

This is the key point emphasized in the text: the effective decay rate of the daughter is governed not by its own half-life alone, but by the half-life of the parent. The system naturally seeks equilibrium, so the ratio between parent and daughter becomes constant.

On the right-hand graph, you see the practical outcome. Each spike represents an elution, where the daughter isotope is washed out of the generator for clinical use. Typically, this is done once every 24 hours. After each elution, the daughter begins to build up again from the parent, forming another peak. Over the course of days, however, the overall yield decreases because the parent itself is decaying exponentially.

So in practice, you get this repeating cycle: each day you harvest a certain amount of technetium-99m, use it for imaging, and then allow the generator to recharge. After about a week, the parent has decayed so much that the yield becomes too low, and the generator must be replaced.

This is the essence of milking a radionuclide generator—an elegant and cost-effective way to provide a steady supply of short-lived tracers for nuclear medicine.

slide37:

Now let's look at how technetium-99m is actually used in medicine. This table lists some of the most common radiopharmaceuticals and their clinical applications. What you can see immediately is how versatile this isotope is—it can target almost every major system of the body.

For example, technetium-labeled macroaggregated albumin is used for pulmonary perfusion studies, helping us assess blood flow in the lungs. Diphosphonates labeled with technetium are widely applied in

skeletal imaging, especially for detecting bone metastases. For the brain, compounds like glucoheptonate or HMPAO can be used to image tumors or measure brain perfusion. Sulfur colloid labeled with technetium allows us to visualize the liver, spleen, and even sentinel lymph nodes. DTPA is useful for renal function and pulmonary ventilation studies. And compounds like sestamibi are used to evaluate myocardial perfusion, making technetium imaging an essential tool in cardiology.

So what do we learn from this? With just a single isotope—technetium-99m—we can design radiopharmaceuticals to study the lungs, bones, brain, liver, kidneys, heart, and more. This is one of the main reasons technetium-99m became the workhorse of nuclear medicine.

With that, we finish our discussion of tracer production. We have seen the different ways isotopes can be generated, particularly the third and fourth methods—cyclotrons for producing positron emitters used in PET imaging, and generators for producing technetium-99m used in SPECT imaging.

Now, we move into the last part of today's lecture: data acquisition. Here, we'll focus on two main approaches. The first is the gamma camera, which detects single gamma-ray photons and is central to SPECT, or single photon emission computed tomography. The second is coincidence detection, which captures pairs of gamma photons emitted simultaneously, and this is the basis of PET, or positron emission tomography.

As we go through this, keep in mind the difference: SPECT relies on detecting individual gamma photons, while PET relies on detecting photon pairs from positron annihilation. And this difference has profound implications for scanner design and image quality.

slide38:

So that concludes the third part, where we discussed radioactivity and tracer production. We learned about the four methods of producing radionuclides, with special focus on the third and fourth methods—cyclotrons for positron emission tomography, and radionuclide generators for single photon emission tomography. Together, these methods provide the backbone of nuclear medicine.

Now let's move on to the final part of today's lecture: data acquisition. This is where the signals generated by radioactive decay are actually detected and transformed into images. There are two main approaches we need to understand. The first is the gamma camera, which is used in single-photon emission computed tomography, or SPECT. The gamma camera captures individual gamma-ray photons one by one. The second approach is coincidence detection, which is the principle behind positron emission tomography, or PET. In this case, we don't just detect single photons; instead, we detect pairs of photons that are emitted simultaneously, traveling in opposite directions.

Why does this happen? Because when a positron encounters an electron, the two annihilate, producing a pair of gamma photons. Detecting these pairs in coincidence allows PET to provide highly accurate information about the location of the annihilation event inside the body.

So, in summary, SPECT relies on detecting single photons with a gamma camera, while PET relies on detecting photon pairs through coincidence detection. This distinction has a profound effect on how scanners are designed and on the quality of the images they produce.

slide39:

Now let's look more carefully at the principle of the gamma camera, which is one of the most important devices in nuclear medicine imaging. I've placed a red diamond here to highlight that this concept is essential.

The very first step in the system is the collimator. Without collimation, gamma photons would be arriving at the detector from all directions, and we wouldn't know where they came from. The collimator, usually made of heavy metal such as lead, only allows photons traveling along certain directions to pass through. In this way, we can establish the line of response. Only photons traveling more or less straight along the openings are detected, while those coming from oblique angles are blocked. This is very similar in spirit to X-ray imaging, where each measurement corresponds to a line integral along a particular direction.

The choice of collimator is closely linked to photon energy. If the gamma rays have too low an energy, they will be absorbed within the body and never reach the detector. If the energy is too high, the collimator cannot stop the oblique photons, so the image contrast is lost. That is why, for clinical nuclear medicine, we typically work in the range of about 100 to 200 keV — a compromise between tissue penetration and collimator effectiveness.

Once a photon passes the collimator, it strikes the scintillation crystal. This crystal, usually sodium iodide doped with thallium, emits a flash of visible light when it absorbs a gamma photon. That flash is extremely faint, so it is coupled to photomultiplier tubes. The photomultipliers amplify the light signal into an electrical pulse.

Then, through a network — sometimes called the Anger logic or Anger position circuit — the system calculates the location where the photon entered the crystal. A pulse height analyzer checks the photon's energy, ensuring we only accept photons in the desired energy window. Finally, the computer records the position and energy, gradually building up an image of tracer distribution inside the body.

So the key sequence is: collimator to select direction, scintillation crystal to convert gamma to light, photomultipliers to amplify the signal, electronics to analyze energy and position, and finally the computer to assemble the image.

slide40:

(No notes for this slide)

slide41:

Now, let's take a closer look at the different collimator designs and how each one affects imaging performance.

Starting with the parallel-hole collimator shown at the top left, this is the most commonly used design. In this case, the object and the image appear at the same scale. It's simple, robust, and widely applied in routine nuclear medicine imaging.

Next, on the top right, we have the diverging collimator. Here, the holes fan outward. The effect is that the image is smaller than the object. This type is useful when you need to fit a large object, like the entire body, onto a smaller detector field of view.

At the bottom left is the converging collimator. As the name suggests, the holes angle inward, producing a magnified image. This is helpful when you want to examine a small organ, like the heart or the thyroid, in greater detail while still using the same detector.

Finally, at the bottom right, we see the pinhole collimator. This works like a camera obscura, creating a highly magnified but inverted image. Because the aperture is very small, spatial resolution can be excellent, but the downside is that sensitivity is poor—most photons are blocked.

The key idea is always a trade-off. If you make the holes or apertures smaller, you gain higher spatial resolution, but you lose sensitivity and the images become noisier. If you make them larger, you collect more photons, but the resolution decreases. Designing collimators is therefore about finding the right balance between resolution, sensitivity, and the clinical question you are trying to answer.

slide42:

Now, let's take a look at something a bit more advanced and exploratory.

We know that gamma rays are a form of high-energy electromagnetic radiation, very similar to X-rays, although they are produced differently. In X-ray science, researchers have developed special devices called polycapillary lenses. These lenses use bundles of tiny capillaries to guide and refocus X-rays, effectively bending them toward a focal point. The result is much higher spatial resolution, since the rays that scatter off-path are redirected and concentrated at the detector.

You can see the schematic here: the X-rays pass through the polycapillary optics, and instead of traveling in random directions, they are collected and focused. This technique has already been applied in confocal micro X-ray fluorescence, which allows elemental analysis within a very small probe volume—even beneath the surface of a material.

Now, the intriguing idea is whether similar optics could be applied to gamma rays in nuclear medicine. If we could refocus gamma rays in this way, we might overcome some of the limitations of mechanical collimators, which block most photons and reduce sensitivity. Polycapillary focusing would potentially increase both sensitivity and resolution, opening up exciting possibilities for future gamma-ray imaging.

That said, this is not part of the standard clinical toolbox—it's more of a research direction. I included it here to show you that beyond traditional collimation, there are innovative approaches being explored to push the boundaries of nuclear imaging technology.

slide43:

When I visited companies or nuclear detector vendors, they often gave me these as small souvenirs—clear blocks that look just like glass. But these aren't ordinary glass. They are specially manufactured scintillation crystals, the heart of gamma cameras.

What makes them remarkable is that when an invisible gamma ray photon strikes the crystal, it produces a tiny flash of visible light. In other words, they convert high-energy radiation, which our eyes cannot see, into visible photons that we can detect.

If you continuously bombard the crystal with gamma rays, you'll see flashes of light—very faint to the naked eye, but easily measurable with sensitive detectors. This conversion process is the critical first step in turning radioactive emissions inside the body into signals that we can ultimately form into an image.

slide44:

As I mentioned earlier, the flux of gamma rays coming out of the body is very low. That means the signals are extremely weak. To make them useful, we need to amplify them, and that's where the photomultiplier tube, or PMT, comes in.

Here's how it works. When a gamma ray enters the scintillation crystal, it produces a small flash of visible light. That light then strikes the photocathode inside the PMT. The photocathode is a special surface that converts each photon of light into an electron.

Now, a single electron is still a very small signal, so the PMT uses a series of plates called dynodes. Each time an electron strikes a dynode, it knocks loose multiple new electrons. By the time this process is repeated across several dynodes, the signal has been multiplied millions of times.

At the final stage, what started as just one or two photons of light has been amplified into a strong stream of electrons—a measurable current. This current is then sent out of the PMT as an analog signal, which can later be digitized for image reconstruction.

So the PMT performs two essential functions: it converts light into electrons, and then it magnifies that tiny signal step by step until it's strong enough to record. Without this amplification process, the gamma camera simply wouldn't be able to detect enough information to form an image.

slide45:

Now that we know how to detect a signal, the next challenge is to determine where that signal came from. Detection is handled by the photomultiplier tubes, but localization—figuring out the exact position of the gamma ray interaction—is accomplished by what's called the Anger network, named after Hal Anger, who invented the gamma camera.

Here's the idea. The photomultiplier tubes cannot be made infinitely small, so they are spaced apart. This means a single scintillation event—a flash of light in the crystal—will usually be seen by multiple tubes at once, with some receiving stronger signals and others weaker.

The Anger network uses a carefully designed arrangement of resistors, as you see here, to spread out those signals. By analyzing the weighted distribution of the outputs, the system can calculate the event's most likely position. In other words, the network acts like a mathematical averaging system: if one region receives stronger signals, the center of activity is estimated there.

So, although an individual photomultiplier tube only gives you partial information, the combination of many tubes together, processed through the Anger logic, gives a precise two-dimensional location of the gamma ray event. This way, every gamma ray photon detected is not only counted but also assigned a position in the image, allowing us to build up a full spatial map of radioactivity inside the body.

slide46:

Now that we understand how detection and localization work, let's see how they come together in a full imaging system. The result is called planar scintigraphy. Essentially, this is a two-dimensional projection image, much like a standard X-ray radiograph, but instead of using transmitted X-rays, we are mapping the distribution of gamma rays emitted from inside the body.

Here's the sequence of events. First, the gamma rays from the patient pass through the collimator, which selects only those photons traveling in defined directions. Next, the gamma photons interact with the scintillation crystal, producing flashes of visible light. Those light signals are then converted into electrical pulses by the photomultiplier tubes.

The outputs are processed through the position logic circuit, which localizes each detected event in the x and y directions. At the same time, the pulse height analyzer examines the energy of each pulse. This step is important because scattered photons inside the body lose energy, and we don't want to include them—they would blur the image. By setting an energy threshold, the system can reject those lower-energy, scattered signals.

Finally, the gating circuit can synchronize signal acquisition with physiological motion, such as the beating heart or breathing cycle. That way, the images can be captured at consistent phases, reducing motion artifacts.

All of these components work together so that each detected gamma photon is accurately positioned and validated, building up a sharp two-dimensional image of radiotracer distribution inside the patient.

slide47:

Now let's take a closer look at the role of the pulse height analyzer by examining the gamma-ray energy spectrum.

On the left, you see the spectrum of technetium-99m when the camera is exposed without a patient in front of it. The dominant feature here is a sharp photopeak, centered around the expected emission energy. There are also smaller features, such as the iodine escape peak and signals from Compton scatter or lead X-rays. These arise from interactions within the detector itself.

But once you place a patient between the source and the detector, the spectrum changes, as shown on the right. Now the clean photopeak is broadened, and the lower-energy region is filled in. Why? Because inside the patient, photons undergo Compton scattering. Each scattering event deflects the photon and reduces its energy. The larger the scattering angle, the lower the photon's energy when it finally emerges. As a result, the detector sees not only the original full-energy photons but also a cloud of degraded, lower-energy ones.

This is a problem because those scattered photons don't carry accurate positional information—they blur the image. The solution is to use the pulse height analyzer. Since the amplitude of the electronic signal from the photomultiplier tube is proportional to the photon's energy, we can set acceptance windows. Only signals that fall within a narrow energy range around the true photopeak—for example, around 140 keV for technetium-99m—are accepted. Signals that are too low, likely coming from scattered photons, are rejected.

This energy discrimination is what gives nuclear imaging its clarity, much like how anti-scatter grids improve X-ray imaging. By gating and filtering the signals in this way, we make sure that the image is formed primarily from primary, unscattered photons, preserving both resolution and contrast.

slide48:

The point spread function is a useful way to judge how well a gamma camera can resolve detail. The narrower the point spread function, the better the spatial resolution. Now, what actually determines resolution in nuclear imaging? It turns out there are several factors at play, each introducing its own kind of blurring.

First, there's the intrinsic resolution of the gamma camera itself. This is limited by the scintillation crystal and the electronics that localize the signal. A thicker crystal spreads the light more broadly, and that uncertainty shows up as blur in the final image.

Second, the geometry of the collimator plays a major role. Whether you use a parallel-hole, converging, diverging, or pinhole design, the length and diameter of the channels directly control how sharp or blurry the image will be. Longer, narrower channels improve resolution but reduce sensitivity, because many photons are blocked. Shorter, wider channels improve sensitivity but at the cost of blurrier images.

Third, Compton scattering inside the patient adds another layer of uncertainty. The deeper the radioactive source lies within the body, the more gamma photons are likely to scatter before they escape. Scattered photons are misleading—they point in the wrong direction, lowering contrast and degrading resolution.

When we put all these effects together, the total system resolution can be expressed mathematically. Just as when you combine independent sources of random error, you add their variances, here too the variances of the three components—intrinsic resolution, collimator resolution, and Compton scatter—add together. The overall resolution is then the square root of that sum. In other words,

system resolution equals the square root of the sum of the squares of the three main contributions.

This is the fundamental idea: every component of the imaging chain blurs the signal a little bit, and together they define the ultimate sharpness—or fuzziness—of the nuclear medicine image.

slide49:

When we think about the point spread function, or PSF, we are asking: how does the imaging system represent a single point source of radiation?

Ideally, if you had a perfect gamma camera, one point of activity in the patient would appear as one sharp point in the image.

In practice, because of the intrinsic resolution of the scintillation crystal, the photomultiplier tubes, and the electronics, the point source becomes blurred into a Gaussian-shaped distribution. That Gaussian curve is what you would expect from the use of the camera and collimator.

However, once you place the patient between the tracer and the detector, Compton scattering adds a complication. Many photons scatter inside the body before reaching the detector, and those scattered

photons broaden the distribution. Instead of a neat Gaussian, the PSF develops a long tail—photons detected in places far from the true source location.

This long tail is undesirable because it reduces spatial resolution and creates image blur. That is why we use techniques like pulse height analysis to reject scattered photons, keeping only the primary gamma rays.

So the main lesson here is: the Gaussian-like PSF represents the best-case resolution of the gamma camera, but Compton scattering stretches the PSF, adding tails and degrading image quality. Understanding and controlling this spread is fundamental to improving nuclear medicine imaging.

slide50:

Now let's look at another important issue in nuclear imaging systems—dead time.

When a gamma photon interacts with the scintillation crystal, it produces a pulse of light. That pulse has a finite width; it takes some recovery time before the system is ready to register the next event. If two or more photons arrive too close together in time, the system cannot distinguish them. Instead of recording two events, it may only record one, or in some cases, none at all. This limitation is called dead time. It reflects the fact that every detector and its electronics need a short period to reset after each event. If the incoming rate of photons is too high, the detector simply cannot keep up.

Think of it like catching candies falling from the ceiling. If they fall slowly, you can catch and count them one by one. But if they start raining down too quickly, you can't separate them anymore—you miss some. That's exactly what happens here with gamma photons.

Mathematically, we can describe this with the true count rate, denoted by capital N, which is the number of gamma interactions that really happen. Then there's the observed count rate, lowercase n, which is what the detector actually records. Because of dead time losses, n is always less than N.

The system dead time, usually written as the Greek letter tau, can be expressed as the difference between the reciprocal of these two rates:tau equals one over n minus one over N. This formula tells us how much recovery time the system effectively needs between two detected events. In practical gamma cameras, dead time losses are usually not a big problem unless the injected activity is very high—but it's an important limitation to be aware of.

slide51:

Now let's connect what we have learned to one of the most important applications of nuclear imaging—FDG PET for cancer imaging.

PET stands for positron emission tomography. Unlike CT or MRI, which give us structural information, PET provides functional imaging. It shows us how tissues are working metabolically inside the body. The key is the tracer we use. For PET, the most widely used tracer is fluorodeoxyglucose, or FDG. This is a glucose analogue that is produced in a cyclotron. Because FDG behaves like glucose, it is taken up by tissues according to their metabolic activity. Cancer cells are often much more metabolically active than normal cells, so they absorb more FDG. When we image the concentration of FDG, we are essentially mapping the regions of high glucose metabolism. This makes PET a very powerful tool for detecting tumors and also for evaluating whether cancer has metastasized, that is, spread to other sites in the body.

The physics behind PET is tied to beta-plus decay. FDG emits a positron, which quickly encounters a nearby electron. When the positron and electron meet, they annihilate, producing a pair of gamma photons moving in nearly opposite directions. These paired photons are detected simultaneously—that's the coincidence detection we discussed earlier.

This entire principle—positron emission, annihilation, and detection of the paired photons—forms the physical and chemical foundation of PET imaging. Today, FDG PET accounts for about 90 percent of all PET scans, highlighting just how central this method is in cancer diagnosis and treatment planning.

slide52:

Now, let's take a closer look at how coincidence detection works in PET imaging.

When a positron emitted from the tracer annihilates with an electron, the process produces two gamma photons, each with an energy of 511 kilo–electron volts, traveling in nearly opposite directions. The key is that both of these photons need to be detected at the same time to confirm a true event.

Here is how the process works step by step. Each gamma photon first enters a detector crystal. Inside the crystal, the photon interacts and produces visible light, which is then amplified by a photomultiplier tube, giving us an electronic pulse. But not every pulse comes from the correct kind of photon—scattered photons, for example, may have lower energy. This is where the pulse height analyzer comes in. It checks the energy of each detected photon and only accepts those within a narrow energy window around 511 keV. This helps us reject scattered signals.

Now, imagine we have two detectors on opposite sides. If one detector registers a photon and the other also registers a photon of the correct energy, and the two signals arrive within a very short time window—just a few nanoseconds apart—then the coincidence circuit recognizes them as a pair. This means they must have come from the same annihilation event along the line connecting the two detectors.

So, only when both detectors see valid 511 keV photons at nearly the same instant do we count it as a true PET event. This is what allows PET imaging to trace the exact line of response inside the body, leading to the powerful three-dimensional reconstructions that make PET such a valuable tool in clinical imaging.

slide53:

Now, here's another way to visualize coincidence detection. In this diagram, we see detectors arranged around the patient, with a positron annihilation event taking place inside the body. From this event, two gamma photons are emitted in nearly opposite directions.

Each detector continuously records signals as a function of time. Most of the time, what you see are background signals or small pulses. But every now and then, a true gamma photon from an annihilation event strikes the detector, and this produces a sharp pulse.

The key idea is that we are not looking at a single detector alone. Instead, we compare signals across pairs of detectors. If detector i and detector j, located opposite one another, both register a pulse at nearly the same instant, and if both pulses have the correct energy—around 511 keV—then we know with high confidence that a true annihilation event occurred along the line connecting those two detectors.

So, the system continuously searches for these paired pulses. The electronics are designed to be fast and complex enough to check every possible detector pair in real time. Whenever a coincidence is confirmed, the system draws what we call a line of response between the two detectors. That line is where the annihilation must have taken place.

This principle—detecting paired pulses and building lines of response—is the foundation of PET imaging. By collecting many such lines from annihilations throughout the body, we can reconstruct a full three-dimensional image of tracer distribution inside the patient.

slide54:

Up to now, we've said that coincidence detection tells us that an annihilation event must have occurred somewhere along the line connecting two detectors. But that's still quite a lot of uncertainty—it could be anywhere on that line.

Now, with time-of-flight detection, we can do better. The idea is straightforward: when the two gamma photons are produced, they leave the annihilation site at the same instant, traveling in opposite directions at the speed of light. If the annihilation happens exactly in the middle, the photons will reach both detectors at the same time. But if the event happens closer to one detector, the photon on that side will arrive slightly earlier.

So by measuring this tiny arrival time difference, we can estimate where along the line the event occurred. Of course, because our timing measurements are not infinitely precise, there is still some uncertainty. That's why we don't pinpoint a single location, but instead assign a probability distribution—centered closer to one detector or the other, depending on the measured delay. The effect is illustrated here. Without time-of-flight information, the whole line is equally likely, so the back-projected signal is spread out. With time-of-flight, the probability is concentrated in a smaller region along that line. When you combine millions of such measurements, the reconstructed PET image is sharper, with better contrast and less noise.

So the concept is simple: coincidence detection gives us the line, and time-of-flight narrows it down to a region on that line where the paired emission happened. This is why modern PET scanners increasingly use time-of-flight technology—it significantly improves image quality without requiring extra radiation dose.

slide55:

So, here are your exercises for this week.

These questions are designed to help you practice the core ideas we learned today: calculating radioactivity from decay counts, relating activity to half-life and the number of nuclei, analyzing how the signal-to-noise ratio changes with dose and scan time, and modeling the technetium generator with parent–daughter decay equations. You are not limited to only these four questions. If you want to challenge yourself, feel free to explore more problems in the textbook. But these are the ones that will be graded.

Since this is an online course, all discussions and clarifications will happen here. If anything is unclear, simply revisit the lecture materials, the slides, or the recorded explanation—we'll keep everything accessible for you to review at your own pace.

That brings us to the end of today's lecture. Thank you for studying along online, and I'll see you in the next session.