slide1:

Last time, we talked about image quality assessment.Today, we begin a new part of the course — focusing on imaging modalities. And the very first modality we will discuss is X-ray imaging.

Before we dive in, let me share a quick study tip. When you are learning a subject as complex as medical imaging, it is really important to build good habits: preview the material before class, review it after class, and pay special attention to the key ideas. I will repeat important concepts as needed to reinforce your understanding.

Earlier in the course, I showed you a poster that summarized the foundation of medical imaging. I also created a similar poster for my graduate-level class, and students really liked it. These visual maps help connect the big picture with the details.

Tools like this can make complicated subjects easier to follow.

So with that in mind, let's move forward and begin our journey into X-ray radiography.

slide2:

we are right on schedule, and today marks the beginning of our detailed study of imaging modalities.

slide3:

At the end of the course, you can create a summary poster for X-ray. It would be a nice exercise to build your own poster that highlights the main points of each modality. Of course, the poster I made for my graduate class goes much deeper, with governing equations. But here, for the undergraduate level, our goal is to provide a clear and introductory exposure.

So today, let me give you an outline. I think it's always good to start with a roadmap, so you know where we are heading. For X-ray imaging, I'll use this outline format to guide our discussion.

We will cover three fundamental aspects: the source, the attenuation, and the detector. This sequence is very logical. First, you need an X-ray source. Without a source, nothing begins — it's like trying to cook without rice. Next, those X-rays must interact with the human body. If there is no interaction, the photons just pass through like air, and we gain no information. If the body absorbs all the X-rays completely, again, no information escapes. The useful case is partial absorption: different tissues absorb X-rays to different degrees. That difference creates contrast, and contrast is what carries the diagnostic information.

Finally, we need a detector to capture that information. The X-ray photons that escape the body must be recorded so that we can reconstruct an image. So in summary, the imaging chain has three main parts: the source, the interaction with the body, and the detector. Together, they form the foundation of X-ray radiography.

slide4:

The story of X-ray discovery is actually a wonderful example of serendipity in science. In 1895, Wilhelm Conrad Röntgen was experimenting with a cathode ray tube — essentially an early electron gun — when he noticed something unexpected. A nearby fluorescent screen began to glow, even though the tube was enclosed in a box. By accident, he had discovered a new kind of invisible ray that could pass through many opaque materials. He called it the "X-ray," where the letter X simply meant "unknown."

One of the first images he captured was of his wife's hand, showing her bones and even her wedding ring. That single picture immediately revealed the enormous potential of this discovery for medicine.

Röntgen's work was recognized worldwide, and in 1901, he was awarded the very first Nobel Prize in Physics.

So from a chance observation in the laboratory came one of the most important tools in medical diagnosis — the X-ray.

slide5:

Röntgen made his discovery while experimenting with a cathode tube. At that time, he didn't know what was happening — nobody did. He placed the tube inside a box, and unexpectedly, a nearby fluorescent screen lit up. When he looked closer, he noticed that objects could block or shape this invisible radiation. For example, a metallic cross placed in front of the tube cast a clear shadow on the screen. This was something completely new — it couldn't be explained without assuming that some kind of invisible physical rays were being produced.

Because the nature of these rays was unknown, he called them "X-rays." Later, they became known as Röntgen rays, in honor of his discovery.

This breakthrough was so important that Röntgen was awarded the very first Nobel Prize in Physics. Today, of course, we understand X-rays in much greater detail, but the original observation was truly remarkable for its time.

slide6:

X-rays are a special type of electromagnetic radiation. They occupy a specific portion of the electromagnetic spectrum.

The part of the spectrum we are most familiar with is visible light, because we can see it with our eyes. But just because we cannot see other forms of electromagnetic radiation does not mean they don't exist. Think about your smartphone. You don't see the electromagnetic waves, but you constantly receive emails, messages, and photos from anywhere in the world. These are all carried by electromagnetic radiation.

Electromagnetic waves are beautifully described by Maxwell's equations. They are sinusoidal waves traveling through space, with an electric field and a magnetic field that are perpendicular to each other. Together, they form the foundation for how light and all electromagnetic radiation behave.

Now, across this wide spectrum, different portions can be used for different purposes. For example, X-rays are used for radiography and CT imaging. Gamma rays are used in nuclear imaging, like PET and SPECT. Infrared and visible light are used for optical imaging and pathology. Radio waves and microwaves are used in magnetic resonance imaging. So, depending on the wavelength and frequency, each portion of the spectrum serves a different role in medicine.

To describe the energy of electromagnetic radiation, we use a fundamental relationship:E equals h times f, which is also equal to h c over lambda.

Here, h is Planck's constant, f is frequency, c is the speed of light, and lambda is the wavelength. From this equation, you can see that energy is directly proportional to frequency and inversely proportional to wavelength.

In medical imaging, we work with very small wavelengths, usually measured in nanometers and angstroms. One nanometer is 10 to the power of minus 9 meters. One angstrom, written as a small circle above the letter A, is even smaller — 10 to the power of minus 10 meters.

So this slide shows you clearly where X-rays fit in the electromagnetic spectrum, and why their energy levels make them so useful for medical imaging.

slide7:

If we narrow down within the electromagnetic spectrum, we can clearly see where X-rays sit. Their wavelengths are extremely small, ranging roughly from one-tenth of a nanometer up to ten nanometers. Within this narrow window, we classify radiation as X-rays.

But not all X-rays are the same. When the frequency is higher — and remember, higher frequency means higher energy — we call them hard X-rays. When the frequency is lower, corresponding to longer wavelengths, the energy is weaker, and we call them soft X-rays.

Different kinds of X-rays are useful for different applications. For example, in mammography, we use relatively soft X-rays because we want sensitivity to soft tissue structures. But if we want to penetrate thicker parts of the body or denser tissues like bone, we need hard X-rays, with energies that can go above 100 kilo-electron-volts.

In non-medical settings, even higher-energy X-rays are used. For example, in airport security screening, luggage may be very thick or contain dense objects. In those cases, stronger X-rays are required; otherwise, all the radiation would be absorbed, and no useful information would come out. So, depending on the wavelength and energy, X-rays can be tuned for very different purposes — from delicate imaging in medicine to powerful scanning in security.

slide8:

We saw this briefly back in the first lecture, but let's look again. This is a standard X-ray radiograph — in this case, a chest X-ray. It's a single two-dimensional projection of the body, captured on film or in digital form.

What you see here is everything superimposed: the heart, the lungs, the ribs — all layered on top of each other in the same image. Because of that overlap, radiography is not always the clearest view. Still, it provides valuable insight into the human body.

For example, if a patient has a broken bone or if there is a foreign object like a bullet, it shows up immediately. That's why radiography is so important in emergency rooms — it's fast, widely available, and gives doctors a quick look inside.

But keep in mind, this is only a projected view. Moving from projection imaging, like radiography, to cross-sectional imaging, like CT, was a huge leap forward in the field of medical imaging.

slide9:

Here we have X-ray computed tomography, or CT. Unlike a simple radiograph, CT allows us to see a clear cross-sectional view of the body. With CT, the resolution can be extremely fine, down to about 300 microns, which is three-tenths of a millimeter. That means we can visualize very small structures and details inside the body.

How do we actually make a cross-sectional image from X-rays? That is the central question of CT, and it will be the focus of our next lecture. For now, I just want to give you a sense of how powerful this technology is.

When you see images like this, you realize how useful X-ray imaging can be. Hopefully, this motivates you to dig deeper — to understand how the X-ray source, the detectors, and the reconstruction process all work together to extract as much information as possible from the human body.

slide10:

Looking into the future, X-ray imaging may become far more portable and accessible. Imagine carrying a powerful but compact X-ray device — perhaps even something integrated into a smartphone. It may sound futuristic, but nothing is stopping us from thinking outside the box.

One open question is how to make X-ray imaging as cheap and convenient as possible. Believe it or not, there are papers suggesting that even something as simple as peeling Scotch tape can generate X-rays under certain conditions. If you're curious, you can look up this fascinating idea.

So while today we use large, expensive machines, the future may bring surprising and creative ways to generate and use X-rays. New ideas like these could turn out to be very important for both medicine and technology.

slide11:

You may remember this TED-Ed video, "How X-rays See Through Your Skin," which I showed you earlier. When I first shared it, it already had a large number of views, and since then, the count has only grown. That tells you how many people around the world are curious about the history and the science of X-rays.

It's a short, engaging explanation of how X-rays work, and I encourage you to watch it again after today's lecture. With the background you now have, you may notice details you didn't fully appreciate before. Revisiting resources like this, once you have more context, is a great way to deepen your understanding.

slide12:

Medical X-rays are essential. Techniques such as plain radiographs and CT scans are the backbone of modern radiology. In fact, some of my colleagues and mentors often say: if you had to pick just one imaging modality as the cornerstone of modern medicine, it would be X-rays. If you take away X-ray imaging, hospitals simply cannot function effectively. If you take away any other modality, they can still manage, as long as X-rays remain available.

There are many reasons for this. X-ray imaging is relatively cheap and widely available. It provides high resolution and fast speed, which are crucial in clinical practice. Another advantage, and one I want to emphasize, is its high geometric accuracy.

What does that mean? Think about MRI: it produces beautiful images, but sometimes those images are geometrically distorted. A tumor might appear shifted, much like how a pencil looks bent when you place it in a glass of water. If a surgeon were to rely on that distorted MRI image without correction, they might target the wrong spot. CT and X-ray imaging, by contrast, give you precise geometric accuracy. This accuracy allows us to deliver radiation beams for therapy or guide surgical procedures with confidence.

X-ray imaging is also both sensitive and specific. Sensitivity means that if the disease is present, the imaging will detect it. Specificity means that if the disease is not present, the imaging will not falsely indicate it. So in many cases, X-rays provide a very reliable diagnostic tool. And even when they cannot fully answer the question, they guide us toward which additional imaging modality to use next.

Of course, there are downsides. The most important is that X-rays are a form of ionizing radiation, which means there is some risk. But if the radiation dose is controlled, the risk is manageable. To put it in perspective, a typical CT scan delivers a dose roughly comparable to what you would receive on an international flight, for example, from the United States to China. If you are comfortable taking that trip, then you should not worry too much about a single CT scan. Still, we always work hard to minimize the dose, and over the past decade, remarkable progress has been made in reducing it — while maintaining high diagnostic quality.

Another limitation is that X-rays have poor contrast for soft tissue. Structures like the brain, liver, or tumors are not as clearly distinguished as they are with MRI. In many cases, we need to inject contrast agents to improve visibility.

Finally, from a global perspective, the medical X-ray market is huge and continues to grow — billions of dollars worldwide, with steady expansion each year. This reflects not only how important X-rays are today, but also how much innovation continues in this field.

slide13:

Medical X-ray imaging generally comes in two main forms.

On the left, you see the simpler and more cost-effective version: a single projection image, like a chest X-ray. It captures one view at a time, and everything inside the body is superimposed together in that picture.

On the right, the principle is extended. At any single instant, you still record just one projection. But instead of stopping there, the X-ray source and detector rotate around the patient. At the same time, the patient table can also move forward or backward. By acquiring projections from many different angles, we collect enough information to reconstruct detailed cross-sectional slices. When these slices are stacked together, they form a volumetric image of the entire body.

So, while projection radiography gives us a fast overview, computed tomography transforms those multiple views into a three-dimensional representation — a much richer source of information.

slide14:

Here is another common application of X-ray imaging. In addition to medicine, you encounter X-rays almost every time you go through an airport. Large machines like this are used routinely to screen luggage.

X-ray imaging in this context is designed to detect dangerous items such as explosives or weapons. It has become a standard part of modern security systems, and most of us are so accustomed to it that we hardly think about the technology at work.

So, this concludes our general introduction. At this point, you should have a clear idea of what X-rays are, where they fall in the electromagnetic spectrum, and some of their major practical uses — in both healthcare and security.

Now we move on to the more detailed and important part of today's lecture: the X-ray imaging chain — the source, the interactions, and the detector.

slide15:

Now, once you have both the X-ray source and the detector, you can build an imaging system. Simply put, if you hold them together in the right configuration, you can start performing medical X-ray imaging.

At the heart of the system is the X-ray tube. This is the device that actually generates the X-rays. We will look at its working principle in more detail, but for now, just remember: the X-ray tube is the source that makes the entire imaging chain possible.

So, the first step in understanding X-ray imaging is to understand the X-ray source.

slide16:

Now let's take a closer look at how an X-ray tube works, because this is the heart of the X-ray source. The figure you see here illustrates the essential components.

At one end, we have a filament, which is usually made of tungsten wire. The filament is coiled so that it has more surface area, and when we pass a strong current through it, the wire becomes extremely hot. This heating process frees up electrons — they become "boiled off" from the surface of the filament.

Next, we apply a very high voltage potential between the filament, which serves as the cathode, and the target, which is called the anode. The anode is also made of tungsten. The electric field between the cathode and the anode accelerates the electrons at very high speeds toward the tungsten target.

When these high-energy electrons strike the anode, their interactions with the tungsten atoms produce X-rays. This process also generates an enormous amount of heat. In fact, more than 99 percent of the energy is converted into heat, and less than one percent actually becomes X-rays.

Managing this heat is a major engineering challenge. The inside of the tube is kept under a high vacuum so that electrons can move freely without colliding with air molecules. But even in a vacuum, the tungsten target would quickly overheat if it were not carefully designed. The solution is to use a rotating anode. By spinning the target rapidly, the electron beam does not strike the same spot continuously. Instead, the heat is spread over a larger area, preventing the tungsten from melting.

Modern X-ray tubes are sealed inside durable glass or ceramic housings and are manufactured with extremely high precision. My colleagues at GE Global Research, for example, have worked on these tubes for many years, refining the design step by step. While the basic principle has remained the same since Röntgen's time, continuous improvements have made the tubes more efficient, more powerful, and capable of delivering the high flux of X-rays needed in today's imaging systems.

So in summary: heat the filament, release electrons, accelerate them with high voltage toward a tungsten anode, and X-rays are generated at the point of impact.

slide17:

Now, let's ask the question: why is it that an X-ray tube can actually generate X-rays? The answer lies in two main physical mechanisms. This is what I call our first "red diamond," or key concept.

The first mechanism is called Bremsstrahlung radiation, a German word meaning "braking radiation." Here's what happens: a primary electron beam, accelerated toward the target, passes close to the nucleus of a tungsten atom. Because the nucleus is positively charged and the electron is negatively charged, there is an attractive force. This force bends the trajectory of the electron. Whenever a charged particle is accelerated or decelerated, it gives off radiation. In this case, that radiation is X-rays. Since the amount of bending can vary, the emitted X-rays span a continuous range of energies. That's why Bremsstrahlung is also called continuous radiation.

The second mechanism is called characteristic radiation, sometimes explained in terms of the photoelectric effect. In this case, a high-energy primary electron collides with an inner-shell electron of the tungsten atom and kicks it out of its orbit. That leaves a vacancy in the inner shell. An outer-shell electron then drops down to fill the gap. The energy difference between the two shells is released in the form of an X-ray photon. Because the energy levels of atomic shells are discrete, the emitted X-ray has a very specific energy — a "characteristic" value that depends on the target material.

So, to summarize: Bremsstrahlung gives us a broad, continuous X-ray spectrum, while characteristic radiation adds sharp peaks at specific energies. Together, these two processes explain how an X-ray tube produces the mixture of radiation that we use for imaging.

slide18:

Now that we've discussed how X-rays are generated, let's look at the spectrum they produce.

First, notice the broad, continuous curve — this comes from Bremsstrahlung radiation. You can see that at lower photon energies, the intensity is relatively high, and as the energy increases, the intensity decreases. However, the very lowest-energy X-rays don't actually appear in the output, because they are absorbed inside the tungsten target itself. Since the radiation is produced beneath the surface, the tungsten material acts as a natural filter, removing the very soft X-rays before they can escape.

Next, at higher energies — around 70 kilo-electron-volts and above, for tungsten — you start to see sharp peaks. These are the characteristic X-rays, labeled here as K-alpha and K-beta. They occur only when the incoming electron energy is high enough to knock out inner-shell electrons, allowing outer electrons to fall into those vacancies and release very specific amounts of energy.

The maximum X-ray energy you can obtain is determined by the tube voltage. For example, if the tube is operated at 120 kilovolts, then the maximum possible photon energy is 120 kilo-electron-volts. However, those very high-energy photons are less common compared to the middle range.

Finally, the intensity of the spectrum — meaning the number of X-ray photons produced — depends strongly on the tube current. A higher current means more electrons striking the target, which produces more X-ray photons across the spectrum.

So the shape of the X-ray spectrum is governed by three key factors: the filtering effect of the target, the discrete characteristic peaks, and the limits set by the tube voltage and current.

And keep in mind — this will also connect directly to one of your homework problems, so be sure to review it carefully.

slide19:

This slide, taken from your textbook, shows some important parameters of the X-ray tube.

Let's start with the cathode. In practice, the cathode is often shaped in a particular way and charged so that it helps to focus the electron beam. Sometimes, additional focusing coils are also placed in the middle region. The purpose is always the same — to narrow the electron beam and concentrate it onto a small area of the anode. That small region becomes the X-ray focal spot.

Why is the focal spot size important? If the spot is too large, X-rays are emitted from a broad area, and the resulting image becomes blurry, similar to motion blur in photography. A smaller focal spot, on the other hand, produces much sharper images.

There's also a clever design trick here called the line-focus principle. The anode surface is angled at some angle, theta. The electrons strike the sloped surface, and X-rays are emitted. Because of the geometry, the apparent or effective focal spot size is equal to the actual size times the sine of theta. This way, you can achieve a smaller effective focal spot while still spreading the heat over a larger physical area. It's a neat compromise between sharp imaging and thermal management.

Finally, the field of view, or coverage, depends on the distance from the focal spot to the patient. As the patient moves farther away, the coverage widens. The maximum coverage is limited by the anode angle.

So, to summarize: focusing gives us a sharp spot, the angled anode helps reduce blur while managing heat, and the geometry of the tube determines how wide a field of view we can capture.

slide20:

We've now reviewed the conventional X-ray tube, its parameters, and how it works. Let me briefly mention a new type of X-ray tube, which represents a more advanced and expensive technology.

Instead of using a solid tungsten anode, this design uses a liquid metal anode. A thin jet of liquid metal flows continuously downward. An electron beam, generated in a way similar to what we discussed before, strikes the liquid metal stream. The interaction between the electrons and the liquid target produces X-rays.

The advantage of this design is that the liquid metal can absorb and dissipate heat much more effectively than a solid target. As a result, the system can generate a much brighter X-ray beam — higher intensity and higher flux — without destroying the anode.

These tubes are very costly, often priced in the range of hundreds of thousands of dollars. We won't test you on the details of this system, but it's good to be aware of such innovations. They illustrate how researchers continue to push the limits of X-ray technology, refining designs to achieve better performance.

I should also add a deeper point for you to think about. When we describe Bremsstrahlung radiation, it may sound puzzling: electrons are attracted to the nucleus, so at one stage they seem to gain energy, while at another stage they lose energy. How can net radiation be produced if these processes are symmetric? The key is that the trajectory of the electron is bent. That bending means the electron is accelerated in a new direction, and acceleration of charge always produces radiation. This is the fundamental physical basis — conservation laws and electromagnetic theory together explain why X-rays are emitted.

So, while the equations provide a clear description, it's important to connect them with the underlying physics. That's what allows us to understand not just how X-rays are generated, but also how new designs, like this liquid metal tube, can open the door to brighter and more efficient X-ray sources.

slide21:

By now, you understand how X-rays are generated. That knowledge positions you well, not only for deeper research in this field but also for practical discussions, even in job interviews.

Now we move to the next important part of the imaging chain — X-ray attenuation. Once X-rays are produced by the tube, or in some cases by a synchrotron radiation source, we use them to probe an object or a patient. As the X-rays travel through matter, some of them are absorbed, some are scattered, and some pass through with little interaction.

It's in this process — the way X-rays are attenuated by different tissues or materials — that valuable diagnostic information is encoded. Understanding attenuation is crucial because it directly determines the contrast and quality of the images we can produce.

slide22:

So, how exactly do X-rays get attenuated?

Think about it this way: when an X-ray beam passes through the body, several outcomes are possible. If there is no interaction at all, the X-rays travel straight through as if the body were completely transparent. That doesn't help us, because no useful information is recorded.

At the other extreme, if there is total absorption, none of the X-rays make it out. That also gives us no signal.

The most useful case is partial absorption, and that occurs through two main processes: scattering and the photoelectric effect. We'll talk about these in more detail on the next slides.

Now, at the atomic level, here's what is happening: X-rays are photons, tiny packets of energy. When they strike the body, they may interact with molecules, atoms, or even inner electrons. Some photons pass through unchanged, forming the primary beam. Others are absorbed completely, transferring all their energy to the material. And still others are scattered — deflected in new directions, sometimes losing energy in the process.

It's the balance of these interactions that produces image contrast. Different tissues in the body — bone, muscle, fat, lung — attenuate X-rays to different degrees, which is why we can see structures in an X-ray image.

slide23:

Here we see the overall picture of how X-ray photons interact with matter. At the atomic level, these interactions can be grouped into four main categories.

The first case is no interaction. The photon simply passes through the material unchanged. You could say the photon gets "lucky" — it doesn't collide with anything. That photon contributes to the primary beam that reaches the detector.

The second case is total absorption, also known as the photoelectric effect. Here, the photon is completely absorbed by an atom. It transfers all its energy to an inner electron, which is ejected as a photoelectron. The atom is left with a vacancy, and when electrons from higher shells fall into that vacancy, characteristic X-rays can be emitted.

The other two cases involve scattering. In coherent, or Rayleigh scattering, the photon is deflected at a very small angle, but its energy remains the same. Because the frequency and phase of the photon do not change, we call this coherent scattering.

In contrast, incoherent scattering, also called Compton scattering, happens when the photon collides with an outer-shell electron and transfers part of its energy to it. The electron is ejected, and the photon is scattered at a larger angle with reduced energy. The energy lost by the photon is taken up by the electron and eventually dissipates as heat.

For medical imaging, the two most important interactions are the photoelectric effect and Compton scattering. The photoelectric effect gives rise to absorption contrast, while Compton scattering dominates at higher energies and contributes to image noise and radiation dose.

So, to summarize: no interaction, total absorption, coherent scattering, and Compton scattering — these four categories describe the essential ways X-rays interact with tissue.

slide24:

Here is a summary of the different types of X-ray interactions with matter. For medical X-rays, the two dominant processes are the photoelectric effect and Compton scattering.

Scattering can be divided into Compton and coherent (or Rayleigh) scattering. But in practice, coherent scattering contributes only a very small portion and is usually not considered important for medical imaging.

Now, in clinical X-ray systems, we typically operate with tube voltages up to about 140 kilovolts peak, or kVp. When the energy is below about 80 kVp, the photoelectric effect is the dominant interaction. Above 80 kVp, Compton scattering becomes more significant. This is why, in diagnostic imaging, we see a transition in the relative contributions of these two effects.

Traditionally, scattering was treated mainly as a nuisance because it reduces image contrast. But modern research shows that small-angle Compton scattering actually carries valuable information. In fact, subtle differences in scattering patterns may help distinguish benign from malignant tumors. However, detecting and analyzing this information requires specialized techniques, which go beyond the scope of our class.

In addition to these, there are two higher-energy interactions: pair production and photodisintegration. These processes become relevant only when the photon energy is very high, well above the diagnostic range. For example, in radiation therapy, where high-energy X-rays or proton beams are directed at tumors, pair production and photodisintegration play important roles. In that context, the goal is not diagnosis, but treatment — essentially using radiation to destroy tumor tissue.

So to summarize: in diagnostic radiology, photoelectric absorption and Compton scattering are the key effects. For therapeutic applications at higher energies, pair production and photodisintegration also come into play.

slide25:

Compton scattering has major implications for radiation protection.

When we use X-rays on a patient, most of the useful photons travel through the body and reach the detector to form the image. But many photons do not follow this direct path — instead, they are scattered in different directions. These scattered photons, especially those from Compton scattering, can expose people nearby.

That is why radiation protection is so critical. For patients, the exposure is usually occasional — maybe once a year, or only when necessary. But for radiologists and technicians working in X-ray rooms every day, the exposure could accumulate over time if they are not protected.

To minimize this risk, we use lead shielding, such as lead glass windows and lead aprons. These barriers absorb scattered radiation and protect staff from unnecessary dose. Without them, repeated daily exposure could become dangerous.

So, one of the key lessons is that while X-rays are extremely valuable for diagnosis, we must always use proper shielding and safety measures to protect healthcare workers.

slide26:

As X-rays pass through biological tissue, they are progressively attenuated. This is a cumulative process: the number of photons decreases rapidly as the beam penetrates deeper.

You can think of the tissue as being made up of thin layers. Each layer absorbs a certain fraction of the incoming photons. A useful way to describe this is the half-value layer concept. Suppose you start with 1,000 X-ray photons. After passing through the first half-value layer, about half remain — 500 photons. After a second half-value layer of the same thickness, the number drops to about 250. After a few layers, only a small fraction of the original photons survive.

So attenuation is really a matter of probability: each layer has the same chance of absorbing a photon, and the result is an exponential reduction in the beam.

It's also worth noting that X-ray imaging is not very efficient overall. In the tube, only about one percent of the electron energy is converted into usable X-ray photons — the rest is lost as heat. Then, after filtration and attenuation inside the patient, only a small percentage of the original photons actually reach the detector.

That's why we need two things: smart algorithms to make the most of the limited data, and highly sensitive detectors to record even the relatively small number of photons that emerge.

So now you have the physical intuition. Next, we will introduce the mathematical description of attenuation — Beer's Law — which provides a precise model for this exponential process.

slide27:

Now let's describe X-ray attenuation in a more quantitative way.

Imagine a beam of X-ray photons entering a small region of tissue. You can think of this region as a pixel, a tiny picture element made of uniform material. The number of photons that enter is what we call the input intensity. The number that comes out on the other side is the output intensity.

The relationship between the input and output depends on two key factors. First, the property of the material, which we describe using the linear attenuation coefficient, written as the Greek letter mu. Second, the thickness of the material, which we call delta x.

Putting this together, the output intensity equals the input intensity multiplied by an exponential decay term. In plain language, this means that as the X-ray beam passes through the tissue, the number of photons drops off exponentially with thickness.

This exponential law comes directly from solving a first-order differential equation. The idea is simple: if you look at a very thin layer, only a small fraction of the photons are absorbed, proportional to the thickness. Stack up many thin layers, and the result naturally becomes exponential.

In fact, this same mathematical behavior appears in many other fields — chemical reactions, radioactive decay, and even population growth or decline.

Now, in the human body, the situation is more complex because the X-ray beam usually passes through several tissues with different properties. In that case, the output is determined by the combined attenuation from all of them along the path.

This exponential attenuation law is fundamental to X-ray imaging. It is the mathematical backbone of both radiography and CT, because it precisely tells us how X-ray photons are reduced as they traverse the body.

slide28:

Now let's extend the exponential attenuation model to the case where the X-ray beam passes through multiple elements along its path.

Suppose the beam is divided into many small steps. For each step, there is a local attenuation coefficient and a thickness. When we add up all of these contributions, we arrive at a cumulative relationship.

This summation has a special name: we call it a ray-sum. The term simply means we are summing up contributions along a single ray of X-rays as it travels through the object.

Now, notice what is known and what is unknown. The thickness of each element is something we set in the model. What we don't know are the attenuation coefficients themselves — those are the values we are trying to determine. So, what we really have is a linear combination of unknowns. That's why this is described as a linear system.

On the other side of the equation, you see a logarithm. Taking the natural logarithm is just a mathematical way of undoing the exponential law we started with. This step allows us to link the unknown attenuation values inside the body with quantities we can actually measure.

Think about it: the input intensity is known because we generate and calibrate the X-ray beam. The output intensity is measured directly by the detector. The ratio of the two gives us a measurable value, and taking the logarithm turns it into the sum of attenuation along the path.

If we make the elements smaller and smaller, the discrete sum becomes a continuous integral. In other words, in the limit, what we measure is a line integral of the attenuation coefficient along the path of the X-ray.

This is the key principle: with X-rays, we do not directly measure the attenuation at each point inside the body. Instead, we measure ray-sums — overlapping signals that represent line integrals. By collecting many such measurements from different angles, we can mathematically reconstruct the image.

So this linear equation forms the bridge between the unknown interior of the body and the measurable data at the detector.

slide29:

So far, we've discussed how X-rays are attenuated as they pass through tissue.

Now let's see how this attenuation creates contrast in an image. Imagine a piece of tissue that is mostly uniform, but with a small bump or nodule inside it. The X-ray beam passing through the normal tissue travels a distance x.

Its signal is attenuated according to the exponential law:

A = N naught e to the power of minus μ x.

Now consider the ray that passes through the nodule. This ray travels an additional thickness, z. Its output is therefore:

B = N naught e to the power of minus μ (x plus z).

The contrast is defined as the difference between the normal signal and the signal through the nodule, divided by the normal signal. Doing the math, we find:

Contrast = 1 minus e to the power of minus μ z.

Notice that the contrast depends only on μ, the attenuation coefficient, and z, the extra thickness of the nodule.

For example, if the nodule is 1 centimeter thick, and μ equals 1 per centimeter, then the contrast comes out to about 63 percent. But in real soft tissue, μ is much smaller, closer to 0.2 per centimeter. In that case, the contrast becomes significantly lower, making small nodules much harder to detect.

This simple example shows the mechanism of attenuation contrast: the difference in signal arises because one path is slightly longer or passes through tissue with different attenuation properties. But in soft tissues, where μ values are close together, contrast can be very limited — and that is why techniques like contrast agents or advanced imaging methods are often necessary.

slide30:

Now we can summarize the X-ray measurement model, which is one of the most important concepts. That's why I've marked this slide with the largest red diamond.

Let's begin with the monochromatic model. Here, we assume the X-ray beam consists of photons all having the same energy. That's an idealization, but it helps us understand the basics. In this case, the incoming intensity, I naught, is attenuated according to Beer's law. The output, which we call I-g, is simply I naught times e to the power of minus the line integral of μ along the path. This means the measured signal is directly linked to the cumulative attenuation along the ray.

If you know the input I naught and measure the output I-g, then by taking the logarithm, you recover the line integral of the attenuation coefficient. That's the foundation of reconstruction in X-ray imaging and CT.

But in reality, things are more complicated. An actual X-ray tube does not produce a single energy — it generates a spectrum of energies. The detector, in most systems today, does not separate them. Instead, it integrates the total energy deposited, whether from soft X-rays or hard X-rays.

That leads us to the polychromatic model. In this case, we must integrate over the entire energy spectrum. For each energy, attenuation follows the exponential law, but the detector sums over all energies together. The result is still a single number — a grayscale intensity — but it is a weighted sum across the energy distribution.

So, the monochromatic model is simple and idealized, while the polychromatic model describes the true measurement process in clinical CT scanners. And this distinction is very important — it explains why certain artifacts, like beam hardening, occur in practice.

slide31:

A very important concept in X-ray imaging is the K-edge of the linear attenuation coefficient.

Now remember, μ — the attenuation coefficient — is not a single fixed number. It depends strongly on the energy of the X-rays. In general, as energy increases, X-rays penetrate more easily, so the attenuation coefficient gradually decreases.

But at certain energies, something special happens. When the X-ray energy becomes high enough to knock out inner-shell electrons — specifically, K-shell electrons — the attenuation suddenly increases. This produces a sharp jump in the attenuation curve, known as the K-edge.

This feature is element-specific. Each element has its own characteristic K-edge energy. For example, in gold, the K-edge occurs at about 80.7 keV. That means if the X-ray photon energy just crosses this threshold, gold absorbs much more strongly due to the photoelectric effect.

This is very powerful because it allows us to do chemically specific imaging. If we use detectors that can resolve X-ray energy, we can identify where particular elements are present based on their K-edge signatures. This opens the door to material decomposition and even molecular-level imaging.

In current hospital CT systems, detectors typically integrate over all energies. That's like measuring the total area under the attenuation curve — you get a grayscale value, but not the details of the energy dependence. However, with modern energy-sensitive detectors, we can directly measure the shape of the curve. That means we can distinguish materials that would otherwise look similar in a conventional CT scan.

This concept is especially important for contrast agents. Elements like iodine, gadolinium, and gold nanoparticles all have distinct K-edges. With energy-resolved imaging, we can separate their signals, perform quantitative mapping, and even design targeted nanoparticle-based contrast agents for advanced molecular imaging.

So the K-edge is not just a physics detail — it's the foundation for the future of spectral CT and molecular X-ray imaging.

slide32:

Now let's look at how contrast agents take advantage of the K-edge effect we just discussed.

On this plot, we see the mass attenuation curves for different substances: lipid, water, bone, iodine, gadolinium, and gold. Notice that lipid, water, and bone have smooth curves that gradually decrease as energy increases. Without extra information, their signals can overlap, and it's very difficult to tell them apart.

But look closely at iodine, gadolinium, and gold. Each has a sharp jump in attenuation at its characteristic K-edge energy. Iodine has its edge at about 33 keV, gadolinium around 50 keV, and gold at about 80 keV. These discontinuities are like chemical "fingerprints."

If your detector simply integrates all energies, you cannot distinguish whether the bright signal you see comes from iodine, or from gold, or even a mixture of bone and water. But if you can measure the attenuation as a function of energy — in other words, if you can see the shape of this curve — then you can identify which element is responsible.

This is the basis of K-edge imaging. By designing energy-sensitive detectors, we can separate the signals of different contrast agents. Clinically, iodine is already the most common CT contrast agent. Gadolinium, which is widely used in MRI, can also be detected in spectral CT. And gold nanoparticles are being investigated for advanced applications, including molecular imaging and drug delivery, because of their strong K-edge at higher energies.

So the key message here is: traditional CT shows only grayscale intensity, but spectral CT can unlock these curves and make imaging chemically specific. That allows us not only to see anatomy, but also to distinguish materials and track contrast agents at the molecular level.

slide33:

Now we come to the final part of today's lecture — the X-ray detector.

In your textbook, you will see quite a bit of discussion about traditional film and early detection methods. Those are historically important, and I encourage you to read them so you have general knowledge. But in practice, those methods are now outdated. What really matters for us is how X-rays are detected in modern CT scanners and what technologies are shaping the future of X-ray imaging.

Modern detectors are digital. They convert incoming X-ray photons into electrical signals that can be processed by the computer. And today, there are two main classes of detector technologies that you should understand:

Energy-integrating detectors — This is the technology used in nearly all current CT scanners. They simply collect all the incoming photons and integrate their total energy. Whether a photon has high energy or low energy, the detector just sums them up. That gives us grayscale images — the kind you are familiar with in clinical CT. A variation of this is dual-energy imaging, where two different energy spectra are used to give some limited energy discrimination. Dual-energy CT is already widely applied, for example, to distinguish iodine contrast from calcium.

Photon-counting detectors — This is the emerging, next-generation technology. Instead of summing all energies together, photon-counting detectors measure each incoming photon individually and record its energy. That means they can separate photons into different energy bins, giving us spectral information. With this approach, we can do material decomposition, K-edge imaging, and molecular-level X-ray imaging, which we just discussed with gold and gadolinium.

So, to summarize: energy-integrating detectors are the current standard, while photon-counting detectors represent the future. They promise higher resolution, lower dose, and much richer information content.

This transition from energy integration to photon counting is one of the most exciting directions in X-ray imaging research and clinical translation.

slide34:

Energy-integrating detectors are the most common, practical, and mature technology used in modern CT scanners.

Their principle is very simple: they collect the total energy deposited by all X-ray photons that arrive at the detector element. It doesn't matter whether the incoming photons are high-energy or low-energy — in the end, they are all added together.

Now, this technology works in two modes:

Indirect modeIn indirect mode, the X-ray photons first pass through a layer of material called a phosphor, such as cesium iodide. The phosphor absorbs the X-ray energy and converts it into visible light. Each X-ray photon produces many visible light photons. These visible photons then strike a photodiode layer, which converts the light into an electrical current. The current is then read out by the circuit. In this process, higher-energy X-ray photons produce more visible light, but once converted to charge, all signals are summed together. That's why this method is called energy integrating or sometimes current integrating.

Direct modeIn direct mode, there is no intermediate light conversion. Instead, the X-rays interact directly with a semiconductor material, such as amorphous selenium. The X-ray photons generate electron-hole pairs in the material. Under a high-voltage bias, these charges are separated and collected as an electrical current. Again, the detector integrates all charges into one signal, which represents the total X-ray energy deposited.

So in both indirect and direct modes, the detector does not distinguish between high- and low-energy photons. It simply reports a single number, proportional to the total energy of all photons that passed through the patient. This gives us the grayscale images we are familiar with in clinical CT.

This approach is robust, efficient, and reliable — and it has been the backbone of medical CT imaging for decades.

slide35:

Now let's talk about an important concept in X-ray imaging: the anti-scatter grid.

As we discussed earlier, when X-rays interact with the body, two main processes occur — photoelectric absorption and scattering. In medical X-rays, scattering is very common. That means many photons change direction inside the body and then try to reach the detector from different angles.

But here's the problem: what we really want to measure is the primary beam — the photons that travel straight along the original X-ray path. These are the photons that contain the true line integral information we need for accurate imaging. If scattered photons are also recorded, they blur the image and reduce contrast.

To solve this, we use an anti-scatter grid. You can think of it as a sheet with many tiny channels, almost like thousands of pinholes. Only photons traveling in the correct direction, aligned with the primary X-ray beam, can pass through these channels and reach the detector. Scattered photons, coming in at an angle, are blocked by the grid walls, which are made of dense materials such as lead.

This dramatically improves image quality by reducing the effect of scatter. Of course, there are trade-offs. If the grid is very tall or tightly spaced, it rejects more scatter, but it can also block some of the primary beam, which reduces efficiency. Engineers must carefully design the grid's height, spacing, and material to strike a balance between scatter rejection and preserving the useful signal.

So, in summary, the anti-scatter grid is a simple but powerful device that ensures we measure mostly the primary beam, giving us much clearer and more reliable images.

slide36:

Here you can clearly see the difference between using an anti-scatter grid and not using one.

On the left, with the grid, the anatomical structures are much sharper. The boundaries of the ribs, the lung fields, and even the soft tissues are more clearly defined.

On the right, without the grid, the image looks more foggy. That haziness comes from scattered photons that are superimposed on top of the fine structures. Essentially, the scatter washes out important details and reduces contrast.

Think of it like looking out of an airplane. On a clear day, you can see the ground with sharp detail. But on a cloudy day, the clouds blur your view. In the same way, scattered X-ray photons act like a layer of clouds, blurring the image and hiding subtle differences.

This becomes especially important when detecting low-contrast tumors. Without a grid, these tumors might be completely hidden in the fog of scatter. With a grid, however, you remove much of that scatter, revealing the underlying details more accurately.

So, using an anti-scatter grid can make the difference between detecting a disease early and missing it entirely.

slide37:

Now let's dive into dual-energy imaging, which is one of the most important developments in modern CT.

All the major CT manufacturers — Siemens, GE, Philips, and Toshiba — have worked on energy-sensitive imaging. While true multi-energy imaging is still a challenge, they have all developed practical approaches to perform dual-energy CT, which means collecting data using two different X-ray spectra.

Let me explain the main strategies:

Siemens uses dual-source CT. Here, the scanner has two X-ray tubes and two detectors mounted at an angle, each operating at a different voltage. For example, one tube might run at 80 kilovolts, while the other runs at 140 kilovolts. This gives you two spectra simultaneously, which means excellent temporal resolution and minimal motion artifact.

GE uses kVp switching. In this method, there is only one X-ray tube, but the tube rapidly alternates its voltage between high and low values — say, 140 and 80 kilovolts — from one projection to the next. As a result, you still get two distinct spectra, but collected sequentially.

Philips uses dual-layer detectors. Instead of modifying the tube, they modify the detector. The top layer absorbs most of the low-energy X-rays, while the second layer captures the higher-energy photons that pass through. This way, low- and high-energy data are naturally separated at the detector level.

Toshiba uses dual-scan mode. In this case, the scanner acquires two full scans of the same patient: one at high voltage and one at low voltage. While this is simpler in concept, it has drawbacks. It takes longer, and the patient may move between scans. Also, if contrast material is injected, its distribution may change between the two acquisitions, leading to misregistration.

So in practice, the first three technologies — dual-source, kVp switching, and dual-layer detectors — are more effective. They provide two sets of spectral information within the same exam, which can then be used to separate materials, improve contrast, and reduce artifacts.

This is why dual-energy imaging is now widely used in clinical practice. It gives radiologists much more information than standard CT and opens the door to advanced applications like material decomposition and virtual contrast imaging.

slide38:

Before we go further into dual-energy imaging, let's pause for a quick review of some basic chemistry concepts—specifically, the atomic number and the mass number, because these concepts form the physical foundation of dual-energy CT.

The atomic number, usually denoted by Z, is simply the number of protons in the nucleus of an atom. It defines the identity of the element. For example, hydrogen has Z = 1, helium has Z = 2, carbon has Z = 6, and so on. If you change Z, you change the element itself.

The mass number, often written as A, is approximately the total number of protons plus neutrons in the nucleus. Since neutrons contribute to the mass but not the charge, A is usually larger than Z.

Take helium as an example. Helium has Z = 2, meaning two protons. Most stable helium atoms also have two neutrons, giving a mass number of A = 4. That's why you see the helium symbol here written with a 2 at the bottom and a 4 at the top.

So why is this important for us? Because the way X-rays interact with matter—whether through photoelectric absorption or Compton scattering—depends strongly on both the atomic number and the mass of the element. High atomic number elements like iodine, gadolinium, or gold interact with X-rays very differently from soft tissue, which is mostly carbon, oxygen, and hydrogen.

And this difference is exactly what dual-energy imaging exploits: by probing tissues at two different energy levels, we can begin to separate and even identify different materials inside the body.

slide39:

So here is a classic expression for the X-ray attenuation coefficient.

The main idea is that attenuation comes from two parts: one part is the photoelectric effect, and the other part is Compton scattering.

The photoelectric effect is very strongly dependent on energy. In fact, as the X-ray energy increases, the contribution from the photoelectric effect drops very quickly — that's why low-energy X-rays are absorbed much more.

The Compton scattering term, on the other hand, has a more gradual dependence on energy. It comes from what is called the Klein–Nishina function, which describes how photons scatter off electrons. So if you put these two terms together, you can describe the full attenuation as a combination of photoelectric absorption and Compton scattering. Now here's the powerful part: if we measure attenuation at two different X-ray energy spectra, we can separate these two contributions. Once that's done, we can predict how the material would behave at any other energy.

You can think about this in two ways:

Physically, we are separating attenuation into its two main mechanisms: photoelectric and Compton.

Material-wise, we are saying that the body can be thought of as a mixture of two basic materials — for example, water and bone. Any other tissue can be represented as a combination of those two.

That's the foundation of dual-energy CT: by using two X-ray spectra, we can get material-specific information, not just grayscale attenuation.

slide40:

So how do we actually perform a dual-energy measurement?

The key is that we collect data using two different X-ray spectra. Let's call them spectrum one and spectrum two.

Now, recall that the attenuation coefficient has two main components — the photoelectric part and the Compton scattering part. When we take the line integral along an X-ray path, we can separate these contributions into two quantities. We usually call them A1 and A2.

So what happens is this: if we measure the transmitted X-ray signal with spectrum one, we get one equation that depends on A1 and A2. If we measure again with spectrum two, we get a second equation.

Now we have two equations and two unknowns. That means we can solve for A1 and A2. And once we know A1 and A2, we can reconstruct the attenuation coefficient at any X-ray energy.

That's the whole idea of dual-energy CT. By using two spectra, we get enough independent information to separate the photoelectric and Compton contributions. And from there, we can start to do material decomposition — distinguishing water from bone, or even identifying contrast agents like iodine or gold.

So dual-energy imaging is not just a trick with hardware. It has a very solid physical and mathematical foundation.

slide41:

Now let's move to the last topic in this section: photon-counting CT.

Up to this point, we've mostly been talking about conventional detectors. Those are what we call current-integrating detectors. What they really do is add up all the incoming X-ray photons across the entire energy spectrum, and then report one number. In other words, they just give us the area under the curve. That means we only get a grayscale image — essentially black and white — where each pixel represents total intensity.

But imagine if, instead of throwing all the energy information into one basket, we could measure the energy of each photon. That's the principle of photon-counting detectors. Instead of giving you one number, they break the spectrum into many energy bins.

So now, instead of saying "this voxel attenuates X-rays by this much in total," we can say, "here is how it attenuates soft X-rays, medium-energy X-rays, and high-energy X-rays." That gives us a rich spectrum of information.

And here’s the exciting part: with that energy-resolved data, we can make not just one grayscale image, but multiple images across energy ranges.

We can even assign colors to those ranges. The result is what we sometimes call spectral CT or molecular CT imaging, where each material can be distinguished by its unique spectral fingerprint.

So while conventional CT gives us one picture in shades of gray, photon-counting CT has the potential to give us multiple pictures at once, even in color. That’s the future direction of CT technology.

slide42:

Now let’s look a little closer at how a photon-counting detector actually works.

In many ways, the design is similar to direct detection. An incoming X-ray photon interacts with a semiconductor material — typically cadmium telluride, or cadmium zinc telluride. This interaction generates an electron-hole pair. Under an applied bias voltage, those charges are quickly separated, creating a tiny current impulse.

Now here is the critical part: in a photon-counting detector, the electronics are extremely fast and sensitive. Each X-ray photon produces its own electrical pulse. That pulse is then compared against a series of preset thresholds.

If the pulse amplitude falls between two thresholds — say, between the fifth and the sixth — we know immediately that the energy of that X-ray photon lies within that range. In other words, the detector doesn’t just say “I saw a photon.” It also says, “I saw a photon in this specific energy band.”

What you see in this diagram is the circuit pathway: the signal is amplified, shaped, compared against thresholds, and then counted in real time. This entire process happens at the pixel level. And today’s state-of-the-art photon-counting detectors can have pixels as small as 55 microns. Each pixel has its own miniature readout circuit, often called an ASIC — an application-specific integrated circuit. Inside each ASIC are hundreds or thousands of transistors, all designed for high-speed, low-noise performance.

This is truly high-tech engineering. And the result is not just a count of how many photons arrive, but a spectrum of how those photons are distributed across energy ranges. That is what enables spectral CT — the next generation beyond conventional CT.

slide43:

Now we come to one of the most exciting applications of photon-counting CT, which is K-edge imaging.

If we are dealing only with the human body — tissues like bone, soft tissue, and water — then dual-energy CT is usually sufficient. With two basis materials, two measurements are enough to solve the equations, and you can predict the X-ray response at any energy level. That is why dual-energy imaging has been so successful in the clinic.

But the moment we introduce contrast agents — iodine, gadolinium, or gold nanoparticles — the situation changes. These materials have sharp K-edges at very specific energies. To detect and separate them reliably, we need true energy-resolved imaging.

That is exactly what a photon-counting detector provides. Instead of just two broad spectra, it measures photons across multiple energy bins. This allows us to do material decomposition and create chemically specific maps. In other words, we can actually tell apart iodine from calcium, or gold nanoparticles from bone, based on their spectral signatures.

And this opens the door to something much more powerful: molecular and cellular imaging with CT. If you attach nanoparticles to specific molecular targets, photon-counting CT can visualize not just anatomy, but biology — at the molecular level.

Of course, there are additional benefits. Photon-counting improves dose efficiency, reduces beam hardening, and increases spatial resolution. But in my mind, the major breakthrough is K-edge imaging: the ability to map contrast agents and even nanoparticles with chemical specificity.

This is why the field is so excited about photon-counting CT — it takes X-ray imaging beyond anatomy and into biology.

slide44:

And let me close with this very exciting development.

Just last year, a major paper was published in the journal Radiology. This study, conducted by a team of physicians in collaboration with Siemens, reported the first human experience with contrast-enhanced photon-counting CT.

What does this mean? For the very first time, we are seeing photon-counting CT applied in real clinical imaging of patients — not just phantoms or animal models. The results demonstrated the feasibility and advantages of this new detector technology, including improved contrast, better noise performance, and the ability to separate materials with high precision.

So this is a milestone. What we have been discussing — from X-ray physics, to attenuation, to detector evolution, to dual-energy imaging, and finally photon-counting — is not just a research vision anymore. It is becoming a clinical reality.

This is where the field is going. And as engineers and scientists, your understanding of these fundamental principles will position you to contribute to the next generation of medical imaging.

slide45:

Here we see the landmark result.

These are the very first human CT images acquired with a photon-counting detector. In the top row, you can compare conventional energy-integrating CT with photon-counting CT. At this stage, the overall image quality and diagnostic value are already comparable to the best systems currently in use.

But what’s truly remarkable lies in the bottom row. With photon-counting CT, we can separate energy information and create maps — for example, iodine concentration maps — directly from the same scan. This means we are no longer limited to grayscale attenuation images. Instead, we can extract quantitative spectral information, detect contrast agents with high specificity, and even generate functional or molecular-level insights.

So this is not just an incremental improvement. This represents a paradigm shift. We are moving from conventional black-and-white CT into a new era of spectral and molecular CT imaging, powered by photon-counting technology.

The field is advancing rapidly, and these first results are just the beginning. With further improvements, photon-counting CT promises to fundamentally transform the way we diagnose and study human disease.

slide46:

Here is an inspiring analogy from biology. This is a microscopic view of the vertebrate retina — truly an amazing design. You can see the two main types of photoreceptors: the rods and the cones.

Rods are highly sensitive to light intensity. They help us see in dim light, but they do not provide color information. In a way, they are like the energy-integrating detectors we discussed earlier — they sum up all incoming light, or in our case, x-ray energy.

Cones, on the other hand, are responsible for color vision. They are less sensitive to intensity but can distinguish between red, green, and blue. In other words, they are like photon-counting detectors — slower, but capable of providing spectral, or energy-specific, information.

So why does the human retina use both? Because together, they give us a complete picture. Rods allow us to detect faint light with high sensitivity, while cones let us see rich color details.

This hybrid approach is exactly the idea being explored in modern CT technology — combining energy-integrating detectors for strong, high-flux signals with photon-counting detectors for detailed spectral information. Nature shows us that using both types of detectors in parallel can be the most efficient design.

slide47:

So here is the idea we are pursuing — a hybrid detector array.

Instead of using only one type of detector, why not combine them? In this design, most pixels are traditional energy-integrating detectors, which provide strong grayscale signals with high efficiency. But we strategically insert photon-counting detectors at certain positions. These pixels give us the spectral information — the energy-resolved measurements that allow for material decomposition, K-edge imaging, and molecular contrast. By mixing the two, we get the best of both worlds: the robustness and speed of energy-integrating detectors, and the precision and specificity of photon-counting detectors.

This concept is similar to the retina analogy I mentioned earlier — rods and cones working together. The rods give high sensitivity, the cones give color. Likewise, here the integrating detectors give us the structural information, while the photon-counting detectors add rich spectral detail.

We have already published several papers on this hybrid design, and it represents an exciting direction for the next generation of CT technology.

slide48:

Here I want to highlight the work of one of my former PhD students, James Bennett.

James focused on the problem of cost and practicality. As you know, photon-counting detectors are still very expensive. If we want to move this technology from the lab to widespread clinical use, we need solutions that balance performance with cost-effectiveness.

His dissertation, later published in IEEE Transactions on Biomedical Engineering, presented a clever system design for hybrid spectral micro-CT. He demonstrated how combining energy-integrating detectors with photon-counting detectors can achieve many of the benefits of full spectral imaging — but at a fraction of the cost.

The results were promising: improved contrast resolution, better spatial resolution, and lower radiation dose. This work was an important step toward making spectral CT more practical and accessible for real biomedical applications.

So this is a good example of how doctoral research can contribute to both advancing the science and solving real-world challenges in medical imaging.

slide49:

Here is another example of our research direction. This work, also involving James and several collaborators, focused on image reconstruction for hybrid true-color micro-CT.

The goal was to demonstrate that when you combine hybrid detectors with proper reconstruction algorithms, you can actually generate color CT images. Instead of the usual grayscale representation, here you see different contrast agents represented in distinct colors.

The figure on the left shows a phantom study — different inserts with iodine, gold, and other materials, clearly separated into unique color channels. On the right, you can see reconstructions from biological samples, where regions of interest are highlighted, again providing more specific information than traditional CT.

So this is really a proof of concept: with hybrid spectral CT, we can move beyond just structure and begin to recover molecular or functional information. This was an important step toward what we now call “true-color CT.”

slide50:

Alright, before we wrap up today’s lecture, let me leave you with the homework assignment.

From now on, we’ll be following the green textbook as our main reference. Each week I’ll select a few problems that are especially useful for reinforcing the key concepts we’ve covered.

For this week, please work on the following three problems:

Page 50, Problem 1.1

Page 51, Problem 1.3

Page 51, Problem 1.6

These are not particularly difficult, but they will give you a chance to apply what we've discussed in class and make sure the ideas really sink in.

So much for today's lecture — thank you, and I'll see you in the next class.