slide1:

Today, we'll be covering the final imaging modality in our course—optical imaging. This topic isn't included in the green textbook, so this lecture stands alone. Think of it as a story—one that unfolds step by step to show you how light can be used to visualize and analyze biological structures and functions.

As we go along, try to follow the flow of this story rather than memorizing isolated facts and details. Optical imaging combines fascinating principles from physics, biology, and engineering, and by learning how these ideas connect, you'll develop a deep intuition for the topic.

Now, let's approach this lecture with curiosity. Enjoy the journey, stay engaged, and you'll come away with a clear picture of how optical imaging works and why it plays such an important role in modern biomedical imaging.

slide2:

As you can see from our course schedule, today's topic is Optical Imaging, listed here as Lecture 23. This is our final imaging modality in the sequence, coming right after the ultrasound sessions. This lecture will connect what we've learned about other imaging techniques to the world of light-based imaging.

slide3:

Today's lecture on optical imaging will be divided into four main parts.We'll begin with optical microscopy, which is the foundation of optical imaging. This remarkable technology was the first to let us see individual cells and even cell division — truly a revolutionary step in science. In this part, we'll discuss the basic idea of visible light as a segment of the electromagnetic wave and how it interacts with biological tissues. Understanding these optical and tissue properties gives us the foundation we need for everything that follows.

Next, we'll move to optical coherence tomography, or OCT. I'll explain its basic principle — a fascinating concept based on light interference — and show you some of its clinical applications. While optical microscopy dominates biology labs, OCT is widely used in medical settings, especially in eye clinics. These two methods — microscopy and OCT — are the most established optical imaging techniques with real clinical and research impact.

After that, we'll briefly explore diffuse optical imaging, including techniques such as diffuse optical spectroscopy (DOS), diffuse optical tomography (DOT), fluorescence molecular tomography (FMT), and bioluminescence tomography (BLT). These are powerful research tools for small-animal and molecular imaging, though they are not yet commonly used in hospitals. Finally, we'll touch on X-ray optical coupling, which includes X-ray luminescence computed tomography (XLCT) and X-ray micro-modulated luminescence tomography (XMLT) — exciting research directions that bridge X-ray and optical techniques.

Because this lecture stands alone without a textbook chapter, I've reorganized the material into a single, coherent story that combines what used to be two separate lectures. Follow along closely — treat this as a guided story rather than a set of isolated facts. At the end, you'll find a homework question asking you to summarize key ideas like OCT, DOT, and BLT. Optionally, you can also transcribe parts of the lecture to help form a draft chapter. My goal is to make the lecture clear, logical, and enjoyable — so if you simply follow the story, you'll understand the essential points.

slide4:

Now, let's begin our story — and at the beginning, there is light.

Light is part of the electromagnetic, or EM, wave spectrum, which spans an enormous range of wavelengths. We've already talked about the shorter wavelengths, such as gamma rays used in nuclear imaging and X-rays used in computed tomography. On the other end, we have radio waves and microwaves, which are much longer and used in applications like MRI and even cooking.

Visible light, however, occupies only a very narrow portion of this vast spectrum. Specifically, it ranges roughly from 400 nanometers to about 1,000 nanometers in wavelength. Within this band, our eyes perceive different colors — violet, blue, green, yellow, and red — depending on the wavelength. You've seen this effect when light passes through a prism, spreading into a rainbow of colors.

Some animals can detect wavelengths beyond our range. For example, certain species can see ultraviolet or even X-ray light, as confirmed by electroretinogram studies. Among humans, individual sensitivity also varies, and a small percentage of people are color-blind, perceiving only shades of gray.

For medical imaging, each modality uses a different part of the spectrum depending on its purpose. In optical imaging, we focus on the portion we can directly see — the visible and near-infrared range. These wavelengths penetrate tissue to a useful depth, making them ideal for studying cells, tissues, and biological processes using light. Understanding this basic physical foundation will help you appreciate how optical imaging works and why it's such a powerful technique.

slide5:

Now let's talk about how light interacts with biological tissue.

Just like X-rays interact with the human body in medical imaging, visible light also interacts with biological tissues similarly. As shown in this diagram, you can see the different behaviors that occur when light reaches the surface of tissue. There's the incident light, which hits the surface, and part of it gets reflected right away. Some of the light, however, penetrates into the tissue and becomes refracted light. Inside the tissue, that light can undergo scattering — bouncing around in different directions. Some of the scattered photons may come back out, which we call remitted light, while others continue to travel through and emerge as transmitted light.

These are the basic mechanisms of light–tissue interaction. You can think of them as reflection, scattering, and transmission. When we send a beam of light into a piece of tissue, some of that light just passes straight through, continuing in roughly the same direction — that's the transmitted portion. The rest is either reflected or scattered randomly. The light intensity that we detect decreases as the beam travels deeper, and that attenuation happens because of two main effects — absorption and scattering.

Absorption means that some of the light energy is absorbed by the tissue and converted into heat. Scattering, on the other hand, means that photons are deflected from their original path in different directions. Together, absorption and scattering reduce the intensity of the transmitted beam — very much like what happens in X-ray imaging. So, when we think about optical imaging, we're dealing with these same physical ideas — transmission, absorption, scattering, and reflection — but now in the visible and near-infrared light range, where tissue behaves quite differently.

So in summary, when you shine light into biological material, part of it is reflected, part of it penetrates and scatters, and part of it is absorbed. The balance among these processes depends on the tissue thickness and optical properties. These interactions are the foundation of optical imaging and determine how deeply light can penetrate and what kind of information we can obtain about biological structures.

slide6:

Now, let's look more closely at scattering in biological tissue.

Scattering happens because of all the small structures within cells and tissues. The exact way light scatters depends on the size, shape, and refractive index of those structures. We can actually classify scattering mechanisms based on the relative scale of the structures compared to the wavelength of light. When the structures are much smaller than the wavelength, we call that Rayleigh scattering. When they're comparable in size to the wavelength, we call that Mie scattering.

I'm not going to dive into the mathematical formulas or the detailed physics behind these mechanisms — that would take us too far. But the key point is this: visible light directly interacts with cellular and molecular features. Every component inside a cell — the membrane, the nucleus, mitochondria, lysosomes, and so on — all participate in scattering. Depending on their size, each component affects light differently. So the light that comes out of tissue actually carries information about the microstructure inside.

This is why optical imaging is so biologically informative. It tells us about cells and molecules directly — something that X-rays or gamma rays can't do. Optical imaging gives us the finest biological resolution among all imaging modalities we have today. The spatial scale of interaction is on the order of micrometers — that's the cellular scale.

A typical cell is about 10 to 30 micrometers in diameter — that's 10 to 30 millionths of a meter. Different cell types vary, of course, and within each cell, there's a huge amount of structure. You have the nucleus, the nucleolus, the chromatin, the mitochondria, the endoplasmic reticulum, the Golgi apparatus, many channels and vesicles — all these substructures scatter light differently. Each layer contributes to the complexity of how light behaves in tissue.

You know, cells are often called the "building blocks of life." If you open any major journal like Science or Nature, you'll see that phrase everywhere — and it's true. Cells form the basic computational and functional units of biological systems. Later, when we talk about machine intelligence, you'll see how we borrow this idea — we use "neurons" in artificial neural networks to mimic the way biological neurons process information.

So, to summarize, scattering happens at multiple levels — from whole cells to tiny organelles — and it gives us an incredible amount of information about biological organization. The optical properties vary with each structure, making tissue highly complex but also rich in information. And just as in X-ray imaging, we have concepts like transmission, absorption, fluorescence, and scattering. When all these are considered together, we call the combined attenuation — it's very similar terminology. So you can see, optical imaging and X-ray imaging share parallel physics, but optical imaging operates at much smaller, biologically relevant scales.

slide7:

Now, in addition to scattering and absorption, we also have another very important mechanism called fluorescence.In X-ray imaging, we already know about X-ray fluorescence — when we send X-rays into a material, they can produce secondary X-rays with different energies. Similarly, in optical imaging, we can have fluorescence in the visible range. Certain biological molecules or engineered proteins can absorb light of one color and emit light of another color with a longer wavelength.

For example, proteins can be tagged with fluorescent markers. One of the most famous is the green fluorescent protein, or GFP. When you genetically attach GFP to a specific protein — let's say to liprin or another biomolecule — they form a pair. The biological part represents your target or biomarker, and GFP serves as the fluorescent reporter. When you shine a blue laser light — about 488 nanometers in wavelength — onto this combination, the GFP absorbs the blue light and emits bright green light. So, when you see that green fluorescence, you know where your biomarker is located.

In other words, by detecting the emitted fluorescence, you can infer the spatial distribution of your biological target. This is the basic principle behind fluorescence imaging. You can even design these proteins so that they target different molecules — for instance, one may emit green, another red, another yellow — allowing you to track multiple biological processes at once. This kind of optical labeling gives you the ability to visualize molecular events inside living systems in real time.

Fluorescence is a very powerful concept. It lets us do what we call "bio-design." Using genetic and bioengineering tools, we can attach different fluorescent proteins — green, pink, red — to different biomolecules. When those molecules move or react inside a cell, we can literally see where and when they act, based on the color of light emitted.

This work was so groundbreaking that it earned a Nobel Prize, awarded to several researchers from Japan and California who helped develop and refine these fluorescent proteins. With fluorescence imaging, we can observe living biological systems directly — not just fixed samples. We simply illuminate the sample with the appropriate light to excite the fluorescent probe, and then detect the emitted light, which is usually at a longer wavelength. Both the excitation and emission are in the visible range, making it easy to detect and analyze.

So overall, fluorescence enables us to perform cellular and molecular imaging in a very direct, intuitive way. You can send in light, stimulate the fluorophore, and watch as it emits a different color of light — giving you a real-time window into biological activity. This is the essence of fluorescence, and it forms one of the most important foundations of modern optical imaging.

slide8:

Now let's talk about bioluminescence. So, what exactly is bioluminescence?You've probably seen it yourself — on a summer night when the air is warm and the grass is glowing here and there with tiny flashing lights. Those little bugs you see flying around with glowing tails — fireflies — are perfect examples of bioluminescent creatures. Or, if you've ever been near the ocean at night, sometimes you can see a faint glow in the water — that's also bioluminescence, produced by microscopic marine organisms.

So bioluminescence simply means light that is produced by a living organism. It's very similar in appearance to fluorescence — both involve light emission from certain molecules or proteins — but the mechanism is quite different. In fluorescence, you must first shine an external light, such as a laser or lamp, to excite the molecules, and then they emit light as they relax back to their original energy state. But in bioluminescence,

you don’t have to do anything at all. The organism itself produces light through a chemical reaction that happens naturally inside its body. No laser, no excitation source — just spontaneous light emission.

Think of those insects again: they’re not being “shined on” by any light source, yet you can clearly see them glowing green or pink as they fly. That’s bioluminescence — living creatures generating visible light on their own. It’s sometimes called cold light, because it doesn’t involve heat like a light bulb does.

From the imaging point of view, this phenomenon is extremely useful. Scientists can isolate the light-emitting proteins from these organisms — the so-called bioluminescent proteins — and then combine them with other target molecules, such as specific genes or cancer markers. When those target genes are expressed in cells, the bioluminescent signal becomes visible. In other words, if a cancer cell expresses a certain gene, it will literally light up — allowing us to observe biological processes directly, inside living tissue.

So, thanks to optical imaging, we can now visualize biological activity in vivo, meaning within a living organism, in real time. This is a major step forward for modern biomedical research. Bioluminescence, like fluorescence, represents another important form of light–tissue interaction — one where living systems themselves emit light through molecular mechanisms. In the next part, we’ll discuss how we can use these natural light-emitting processes to actually perform imaging.

slide9:

Alright, now let’s take a moment to look at something simple and familiar — the optical lens.You’ve seen these before: there are two basic types of lenses — converging lenses and diverging lenses. I think most of you learned this back in high school, right? Maybe in Physics II, or in your earlier science classes. No? Well, in any case, this is something that every student should know at least conceptually, and if you don’t, that’s perfectly fine — we’ll go through it together now.

So, here’s the basic idea. When you send in a parallel beam of light into a converging lens, the rays bend, or refract, toward each other and meet at a single point called the focal point. For a diverging lens, it’s the opposite — the rays spread apart after passing through the lens. If you trace those diverging rays backward, they seem to come from a single point on the same side of the lens — that’s the virtual focal point.

Now, how do we actually form an image using these light rays? Let’s think about a simple example. If you have a real object in front of a converging lens, a ray parallel to the optical axis will bend and pass through the focal point on the other side. Another ray that passes through the optical center of the lens continues in a straight line, unchanged in direction. Where those two rays meet, that’s where the image forms. This is what we call ray tracing, and it’s a fundamental rule of geometrical optics.

The trick to making this work lies in the curvature of the lens surfaces. If you design the curvature properly, you can make all the parallel beams focus precisely at the same point. Usually, we approximate the surface as spherical when the lens is thin — that’s why it’s called the thin lens approximation. So as long as the curvature is correct and the lens is thin, parallel rays will converge or diverge exactly as expected.

There’s also an important rule here: when a light ray passes through the optical center — the very middle of the lens — its direction does not change. It just goes straight through. That’s one of the key rules for tracing rays. Another important principle is reversibility — light follows the same path backward and forward. So, if a ray travels from the object through the lens to form an image, the light could just as easily travel the opposite way, from that image point back through the lens to the object.

If you take the time to think about it carefully, this all makes sense geometrically. The symmetry of the two curved surfaces ensures that when a ray passes through the center, the small refraction at one surface is canceled by the opposite refraction at the other, keeping the ray direction unchanged. That's why it goes straight through without bending. All these properties — convergence, divergence, optical center, and reversibility — together form the essential rules for understanding how lenses create images.

slide10:

Now, let's summarize what we just discussed by looking at the rules for ray tracing.There are a few simple but very important rules that describe how light behaves when passing through a lens. Rule number one says that a ray passing through the center of a lens will not be deflected — it keeps its direction. This happens because the lens is symmetric, and the small amount of bending at one surface is canceled by an equal and opposite bending at the other. So, any ray that goes right through the center remains straight.

Rule number two tells us that rays parallel to the optical axis will converge to the focal point on the opposite side after passing through a converging lens, or will appear to come from a focal point on the same side in the case of a diverging lens. And if you reverse the process — sending rays from the focal point through the lens — they'll come out parallel to the optical axis. This is a beautiful symmetry and is very useful when tracing images.

Then we have rule number three, which extends that idea. Parallel rays entering the lens at different angles all meet at a single focal plane after refraction. This rule also works in reverse. These three rules are the foundation of geometrical optics, and with them, you can quickly predict where and how an image forms.

So, when you combine these ideas, you can understand how a real image forms with a converging lens — the rays physically meet at a point on the opposite side. For a diverging lens, the rays spread apart, and the image appears to form on the same side as the object — that's a virtual image. These concepts are very simple once you visualize how the rays travel.

slide11:

Here we can see some examples of actual lenses used in optical systems.These are the real-world components that apply all those rules we just discussed. You can see lenses of different shapes and sizes — convex, concave, cylindrical, plano-convex, biconvex, and so on. Each type has a specific purpose, but all follow the same physical principles of light refraction and focusing.

If you're patient enough to spend time studying these lenses, you'll find that the underlying physics is straightforward and elegant. What's fascinating is how we can shape these pieces of glass to control light precisely — to focus it, spread it, or redirect it any way we want.

In fact, working with lenses can be quite enjoyable. When you actually draw ray diagrams and see them align perfectly, it's like creating a piece of art. In optical labs, researchers and engineers spend a lot of time polishing lenses and adjusting their curvature to make sure light behaves exactly as intended. It's a combination of science and craftsmanship.

So, remember — even though the optical principles are simple, mastering them opens up an entire world of imaging technology. Everything from microscopes to cameras to telescopes depends on these very same ideas of ray tracing and lens design.

slide12:

Now let's take a look at something even more fun — the zoom lens.You can see in this diagram and the photos on the right how the zoom lens works. It's really cool. The zoom lens is designed so that by moving certain lens elements back and forth, you can continuously change the focal length. That means you can make the image appear closer or farther away without ever changing the focus on your object.

Here you can see the different lens groups — usually labeled L1, L2, L3 — working together. When the internal lenses move relative to each other, the path of the light rays changes. The blue, yellow, and green paths you see here represent different light rays being bent and refocused as the lenses shift. By doing this, the system can magnify what you're looking at — that's the "zoom" part.

So, as you zoom in, you can magnify small details more and more. You start by seeing larger features, then zoom in again to reveal smaller and smaller structures — down to cellular or even sub-cellular levels, depending on the optical system. The magnifying power increases with each step, allowing you to explore fine details that are invisible to the naked eye.

That's the basic idea behind zoom optics. It's not just about bringing things closer — it's about dynamically adjusting magnification while keeping the image in focus. And that ability to continuously zoom while maintaining clarity is one of the most powerful features of modern optical systems, especially in microscopy and photography.

slide13:

Now, let me talk about a very cool idea — the confocal optical microscope.This is truly a clever and elegant design in optical imaging. Let me first describe the essential concept. You can see here in the diagram that we start with a laser source — this is our input light. The laser beam is reflected by a semi-transparent mirror, or what we call a beam splitter, and it's directed downward through a set of lenses that focus the light onto a single tiny spot, known as the focal spot. That blue plane you see here represents the focal plane — the layer of the specimen that's currently in focus.

Now, of course, light doesn't only come from that exact focal plane. Some light is reflected or emitted from regions above or below the focal plane — those are the out-of-focus areas. These are shown here as the dotted lines. The genius of the confocal design lies in how it handles these different signals. Only light from the true focal point is allowed to reach the detector. How? There's a small pinhole aperture placed in front of the detector. This pinhole blocks light coming from out-of-focus regions, allowing only the in-focus light — the sharply focused point — to pass through.

As a result, the microscope effectively captures light from one small point at a time, dramatically increasing the resolution and contrast of the image. That's why we call it confocal: there's a focus point in the illumination path and another in the detection path — both confocal, both matched to the same plane. The effect is that all blurred background light is rejected, giving you an image that's crisp, clean, and precise.

But there's more — because this setup only collects light from one point at a time, we can't capture a full image in one shot. Instead, the system uses an X-Y scanning device to move the laser spot across the sample, scanning point by point and line by line. Each time, one pixel of data is collected, and after scanning

the entire area, you can reconstruct a two-dimensional image. If you then scan multiple focal planes at different depths, you can combine them to form a three-dimensional, volumetric image of the specimen.

This technique revolutionized optical microscopy. You can think of it as starting from the basic microscope — the magnifying glass we all know — and then taking a major leap forward. Traditional microscopes magnify, but they can't reject the out-of-focus blur very well. The confocal microscope solves that problem beautifully by adding precision optics and scanning. It's a milestone in optical imaging, giving us a much higher resolution and clearer view of microscopic structures.

So, in short, the confocal optical microscope is a powerful tool that combines focused laser illumination, spatial filtering, and precise scanning to achieve incredibly sharp images. It's one of those designs that feels both simple and brilliant once you understand it — truly a landmark in modern biomedical imaging.

slide14:

If you're interested, you can click the link on this slide to watch an actual cell division video.What you're seeing in that video is an image captured using the confocal microscopy principle that we just talked about. The image clearly shows living cells dividing, and the red-colored regions represent the nuclei or specific fluorescent markers that have been tagged so we can visualize the process in real time.

This is a great example of how confocal microscopy works in practice. The technology was so revolutionary that it became one of the most important tools in biological research. It's amazing to think that this entire concept — scanning one point at a time through a small aperture to reject out-of-focus light — came from an idea developed when its inventor was still a PhD student. It shows how a simple but brilliant idea can completely change the way we see the microscopic world.

Even though the confocal microscope may seem common today, the design is still considered incredibly clever. It's a perfect example of elegant engineering — combining optical precision with computational reconstruction to produce crisp, high-resolution images of living cells. Watching something as fundamental as a cell dividing, in real time, is truly fascinating. It reminds us how powerful and insightful optical imaging can be.

slide15:

Now, another important topic I want to mention is optical fiber.Here you can see how light travels through the fiber — the structure has a core made of a material with a high refractive index, surrounded by a cladding with a lower refractive index. This design is intentional because it causes the light traveling inside the core to undergo total internal reflection. Every time the light hits the boundary between the core and cladding, it reflects back in, staying trapped inside the fiber.

The result is that the light — and the energy it carries — can travel a long distance without significant loss. So, you can think of the optical fiber as a very efficient light guide. It delivers light, or an optical signal, from one point to another with minimal attenuation, even over long distances and through curved paths. It's a beautiful example of using the laws of physics — in this case, total internal reflection — for practical engineering.

Now, imagine you place this kind of fiber inside the body, say along the heart wall or into a living tissue. The optical fiber can both transmit light and collect light that's reflected or emitted back from the tissue. In the

schematic below, you see a pair of fibers — one for transmitting and one for receiving. The transmitted light excites or illuminates the target region, and the returning signal is collected and carried back through the receiving fiber.

This setup is the foundation for optical biopsy and miniaturized confocal microscopy. Using fiber optics, we can bring the principles of confocal microscopy directly inside a living organism — even inside a human patient. That means we can perform microscopic imaging in vivo, observing tissues at the cellular level without the need for large or invasive instruments. So optical fibers not only deliver signals; they open the door to new types of medical imaging that can look inside the body in ways that were impossible before.

slide16:

Now, let's talk about optical biopsy, which is another fascinating application of optical imaging.Traditionally, when we talk about a biopsy, we mean the physical process of removing a small piece of tissue from the body so it can be examined under a microscope. That's what doctors do tens of millions of times every year — they take a tissue sample, send it to the lab, and after several days or even weeks, you finally get the results. It's a reliable method, but it's slow, invasive, and can be uncomfortable for the patient.

Now imagine if you could see those same microscopic structures instantly — without ever cutting into the tissue. That's the idea behind optical biopsy. Instead of using mechanical tools to remove a sample, we use light. You simply insert a tiny, flexible optical fiber, shine light onto the tissue, and collect the reflected or emitted signal. The information is transmitted back through the fiber and processed in real time.

In other words, you can obtain the same kind of diagnostic information that you would from a traditional biopsy — but noninvasively, and in milliseconds rather than days or weeks. The optical fiber acts as both the light source and the detector, allowing you to see cellular and subcellular details immediately inside the living body.

This is truly an exciting step toward real-time, in vivo diagnosis. With optical biopsy, you can visualize tissue structure clearly and evaluate whether it's normal or diseased right at the point of care. And as we'll see in the next part of this lecture, the next major development in this direction is optical coherence tomography (OCT) — a powerful technique that extends this same concept of optical imaging into three-dimensional, high-resolution tissue visualization.

slide17:

Now, we've just learned about ultrasound, and we've also talked about confocal microscopy — so let's put them into context.As you can see in this comparison plot, confocal microscopy gives us excellent resolution — down to about one micron — but the penetration depth of visible light is quite limited. Optical light simply can't penetrate very deeply into tissue. For example, if I shine a light on my hand, you can't see it clearly through the other side. Human tissue isn't optically transparent, so the light is mostly scattered or absorbed before it can travel far.

On the other hand, ultrasound penetrates much deeper — several centimeters, even up to tens of meters in other contexts. That's why it's such a powerful clinical imaging tool for viewing inside the body. But its resolution is much lower than that of optical methods — typically in the range of hundreds of micrometers.

So, what about optical coherence tomography (OCT)? OCT sits right in between these two techniques. It offers resolution close to that of confocal microscopy — on the order of a few micrometers — but with a much greater imaging depth, typically a few millimeters. That's what makes OCT so useful: it bridges the gap between high-resolution optical imaging and deep-penetrating ultrasound.

These three — ultrasound, confocal microscopy, and OCT — form the most clinically and biologically relevant imaging modalities for soft tissue. They're not just theoretical; they're widely used, well-established, and extremely powerful in both research and clinical applications.

slide18:

That's why next, we're going to focus on Optical Coherence Tomography, or OCT — the green section in our outline.OCT is one of the most important and practically useful techniques in optical imaging. If you've been following the flow of this lecture — from optical microscopy to confocal microscopy — OCT is the next logical step. The last two topics in this lecture are more research-oriented, but OCT is something you'll see everywhere, from biology labs to hospitals.

So what makes OCT different from the previous techniques? Well, if you think about it physically, traditional or confocal microscopy uses light to illuminate a sample and then collects the reflected or emitted light. But conceptually, it treats light as particles, or photons — bouncing around, scattering, and being absorbed.

In OCT, we take a different approach. Here, we treat light as a wave — a coherent electromagnetic wave, like a laser beam. This shift in perspective is key. Because once you treat light as a wave, you can use interference — the way waves combine when they meet — to extract incredibly fine details about the structure of the tissue.

This coherence-based approach allows OCT to perform precision measurements of distance and structure on the order of micrometers — even smaller than the wavelength of light itself. So, if confocal microscopy gives us sharp 2D optical sections, OCT goes further. It lets us measure depth and reconstruct 3D images of tissue layers, much like an optical version of ultrasound.

slide19:

Alright, now let's go back a little bit. This part is some trigonometry stuff, okay? So just follow me step by step. Suppose you have one wave — I'll call it y one equals a cosine omega t. That's your first wave. Then you have another wave, same amplitude, but traveling a slightly different optical path. That one is y two equals a cosine omega t plus phi.

Now, this phi is the phase shift — it tells us how much the second wave is delayed compared to the first one. When these two waves come back together and recombine, they won't be perfectly in sync anymore. Because the optical paths are different, there's a small phase difference between them.

Now, when you add these two waves together — that's what happens inside the OCT interferometer — the total wave is the sum of those two cosine terms. So, mathematically, y equals y one plus y two equals a cosine omega t plus a cosine omega t plus phi.

Now, we can simplify that using a trigonometric identity. You might remember this one: cos A plus cos B equals two times cosine of A plus B divided by two, multiplied by cosine of A minus B divided by two. So if

we apply that to our case, A is omega t, and B is omega t plus phi. Plugging that in, we get this: y equals two a cosine phi over two, times cosine omega t plus phi over two. Let me say that clearly — y equals two a cosine phi over two times cosine omega t plus phi over two.

Now look — the amplitude of this combined wave depends on phi, the phase difference. That means as phi changes, the amplitude changes. So, a phase change directly changes the amplitude we measure.

Now here's another important point. In real optical systems, we can't measure the electric field itself — the light oscillates way too fast, hundreds of trillions of times per second. What we can measure is intensity. And intensity is proportional to the square of the amplitude.

So let's take that amplitude — two a cosine phi over two — and square it. That gives us intensity proportional to two times a cosine phi over two squared. Or, in simpler words, I equals four a squared cosine squared phi over two. Let me repeat that slowly so you can hear it clearly — I equals four a squared cosine squared phi over two.

Now, a squared term is just the reference intensity. We usually call that I naught. So when we replace a squared with I naught, we get the final formula: I equals four I naught cosine squared phi over two. That's the key relationship. I equals four I naught cosine squared phi over two.

Now, let's think about what that really means. Even a tiny change in phi — a tiny phase shift — can cause a big change in intensity because of that cosine squared term. If phi equals zero, meaning the two waves are perfectly in phase, the cosine of zero is one, and the intensity is maximum — four times I naught. If phi equals one hundred eighty degrees, meaning the waves are half a wavelength out of phase, the cosine of ninety degrees is zero, and the intensity drops all the way to zero.

So, with just a small phase change — half a wavelength — the brightness goes from full to completely dark. That's interference. That's how powerful this relationship is. Now, remember, the wavelength of light is extremely small — around four hundred to one thousand nanometers. That means this interference can detect changes in optical path length on the order of nanometers. That's unbelievably precise.

So, the big idea here is that a small change in phase produces a measurable change in amplitude — or in the brightness that we detect. That's the fundamental principle of OCT. And this intensity, I, is modulated by cosine squared phi over two. So as phi changes — from zero to three hundred sixty degrees — the intensity goes from bright to dark and back again. That creates an interference fringe pattern.

Those bright and dark fringes are what we analyze in OCT. Each fringe corresponds to a small change in path length, and by counting or analyzing those fringes, we can calculate the depth with extremely high accuracy — down to micrometers or even nanometers.

So, remember this equation — I equals four I naught cosine squared phi over two. That's the one you want to keep in mind. It's the bridge between phase change and amplitude change. And that's the beauty of optical coherence — we can't see the phase directly, but we can measure the amplitude. The change in brightness tells us how the phase has shifted.

That's why OCT is such a precise imaging technique. Two beams — one reference, one sample — recombine, interfere, and the resulting intensity pattern gives us incredibly fine depth information. So again, the simple rule: a small change in phase leads to a measurable change in amplitude. That's the essence of optical coherence tomography.

slide20:

Now, I mentioned optical path several times, so let's take a moment to understand what I actually mean by that. Suppose you have coherent light — right light that is coherent means all the waves maintain a fixed phase relationship with one another. This light can include multiple colors or wavelengths — red, orange, green, purple, and so on — each with its own wavelength, lambda one, lambda two, lambda three, and so on. So, when you look at this broadband light, it's really a mix of many slightly different frequencies that are all still coherent enough to interfere with each other.

Now, imagine this light entering an interferometer setup. The light from the source first hits a beam splitter. This beam splitter is partially transparent, meaning part of the light passes straight through while the rest is reflected upward. That splits the beam into two optical paths. The first beam passes through the splitter, hits a fixed mirror — which is a perfect reflector — bounces back, passes through the beam splitter again, and goes down to the detector. That's path one. The second beam, the reflected one, travels upward to a movable mirror, also a perfect reflector. It bounces back down, and because the beam splitter is semi-transparent, some of that returning light passes through to the same photodetector. That's path two.

So, now you have two optical paths — one fixed, one adjustable — and the light from both paths recombines at the detector. If those two paths are the same length, the two beams arrive perfectly in phase, meaning their peaks and valleys line up. When that happens, they interfere constructively, and the total signal at the detector is at its maximum. But if you move that movable mirror by half a wavelength, the light in one arm travels a little farther, and the two waves arrive completely out of phase — one is at a peak while the other is at a trough. They cancel each other out, and the detector sees zero intensity. So, by moving that mirror just a tiny bit — even by half a wavelength — you can make the signal go from bright to dark. This is the basic principle of interferometric measurement.

Mathematically, this is described by the formula we derived earlier, which follows that same "one plus cosine" relationship. The cosine term represents the phase modulation, and the argument of that cosine depends on the difference in optical path length multiplied by the frequency of the light. That product — the path difference times the frequency — gives you the phase angle. So, for a single wavelength, say the red light, the signal can be written as I equals two I naught times one plus cosine of two pi delta L times nu. Now, if you use several wavelengths — orange, green, purple — each will have its own cosine term because each frequency is slightly different.

Since light interference is a linear process, you can add all these contributions together. When you do that, you get a sum of cosine terms, each oscillating at slightly different frequencies. The overall signal, therefore, looks like a modulated waveform — a burst of interference fringes that fade in and out. If you had only one wavelength, you'd see a clean, continuous sine wave pattern of bright and dark fringes. But when you combine multiple wavelengths, the slightly different frequencies cause the waves to occasionally line up and then drift out of phase again. This creates the clustered or "envelope-shaped" interference pattern we often see in broadband interferometry.

So, the total intensity can be written as gamma of delta L equals two times the sum of I naught i times cosine of two pi delta L times nu i, where the sum is taken over all wavelengths. Each component is modulated by its own cosine term, and when all are added together, they form this envelope pattern. When the mirror is positioned such that all the wavelengths line up, you get strong constructive interference. As you move the mirror further, the wavelengths gradually fall out of phase, and the interference fades away.

This is the essential idea behind interference of coherent or partially coherent light. Two optical paths — one fixed and one movable — produce an interference signal that varies with the optical path difference. For multiple wavelengths, the pattern becomes modulated and forms an interference envelope. And by analyzing that envelope, you can precisely determine the optical path length difference. This concept forms the foundation for optical coherence tomography, where we use these interference patterns to extract depth information with extremely high precision.

slide21:

So now, let's think about what happens if we have not just a few components, but many, many wavelength components all mixed together — just like a rainbow of light. Imagine several sine-wave components of different colors, each slightly different in frequency. When we add them together, the result no longer looks like a clean, single-frequency wave.

Instead, you start to see an envelope pattern — a burst that looks like a Gaussian-shaped profile, or what we often call a coherence envelope. The more components you add — red, orange, green, blue, all slightly shifted — the narrower and sharper that envelope becomes. This figure visually shows exactly that. With two components, you can still see the oscillations clearly. Add three components, and the interference gets tighter. Add seven components, and now the oscillations begin to cluster and fade away at the edges, forming a smooth, localized shape.

So, if you keep increasing the number of wavelengths, the interference pattern becomes confined in space — like a short pulse — and its overall shape starts resembling a Gaussian curve. That's the idea behind partial coherence. The more wavelengths you combine, the shorter the coherence length becomes, and the interference fringes appear only within that small region.

slide22:

Now, if you keep adding more and more wavelength components — so many that they form a continuous spectrum — the math also changes. Instead of summing over discrete frequency components, we move into a limit where the summation becomes an integral. Each small intensity component contributes a little bit of energy within a narrow frequency band, and when you take the limit, the total becomes an integral over the full frequency range.

Mathematically, this looks like gamma of delta L equals two times I naught times the integral of S of nu times cosine of two pi delta L nu, d nu. Here, S of nu represents the power spectrum of the light source, describing how much energy each frequency contributes.

Now, if this looks familiar, it should — because this integral is nothing but a Fourier transform. You might remember that in a standard Fourier transform, we use an exponential kernel, e to the power of j two pi something, but since cosine is an even function, this is essentially the real part of that transform. So, using only the cosine kernel still gives us the same kind of frequency-to-space relationship. If you were to include sine terms too, you'd have the full complex exponential, but for our optical system, the cosine component is enough.

This means that what the interferometer is really doing — physically — is performing a Fourier transform of the source spectrum. The detected interference pattern is the cosine transform of the spectral density. So,

when you move the mirror and record how the intensity changes, you're effectively mapping out the coherence function of the light. When the optical path difference matches within the coherence length of the source, you get visible fringes. As soon as you move beyond that range, the interference fades away because the different frequency components go out of sync. That's why interference fringes are only observed when the two optical paths are matched within that coherence length.

Now, think about how we use that in imaging. As the movable mirror scans back and forth, you collect this varying signal over time. Some parts of the sample are closer, some are deeper, so their reflected signals arrive at slightly different delays. When you record and reconstruct all those depth-dependent interference signals, you essentially form an image — one line of depth at a time. And by stacking those lines together, you can build a full cross-sectional image. That's the basic principle of how optical coherence tomography captures structure in depth using interference.

slide23:

This same principle is used in many important applications. One of the most famous examples is the Michelson Interferometer, which was not just used in imaging but also in fundamental physics experiments. In fact, it played a crucial role in shaping modern physics. The Michelson setup is just what we described — a beam splitter, a fixed mirror, and a movable mirror. Light travels along two paths, reflects back, and the interference pattern tells us how the path lengths differ. The optical path length difference is given by delta L equals two times the difference between L two and L one. The phase difference between the two beams is phi equals two pi delta L divided by lambda. And the detected intensity at the photodetector is given by I equals I one plus I two plus two times the square root of I one times I two times cosine phi.

Now, if the two beams arrive in phase, meaning the optical path difference is an integer multiple of the wavelength, that gives constructive interference — maximum brightness. Mathematically, that happens when two pi delta L divided by lambda equals two m pi, or simply delta L equals m lambda, where m is an integer — zero, one, two, and so on. But if the path difference equals an odd multiple of half wavelengths — like delta L equals m plus one-half times lambda — then the waves cancel out, producing destructive interference, or darkness at the detector.

What's fascinating is that this same interferometer was used by Michelson and Morley in one of the most famous experiments in physics history. They wanted to test whether the speed of light changes depending on the motion of the Earth — whether it moves through something called the "ether." The experiment was performed at Case Western University, and at that time, people expected to see a shift in the interference pattern as the apparatus rotated, meaning that light would travel slightly faster in one direction than another. But no shift was observed. The result was completely null. That means the speed of light was constant in all directions, independent of the motion of the source or the observer.

That single observation changed the world of physics. It showed that light doesn't behave the way classical mechanics predicted. A few years later, Albert Einstein took that result and built his theory of Special Relativity around it — declaring the speed of light to be constant for all observers. So this simple optical instrument — based on the same interference principles we've been discussing — essentially marked the birth of modern physics. The Michelson Interferometer not only explains how OCT works, but also represents the bridge between classical and modern physics. It's a truly beautiful piece of science — elegant, powerful, and revolutionary.

slide24:

Now let's talk about the principle of OCT. It's really the same idea as what we just discussed with the Michelson interferometer, but now applied to imaging. You can see here, we have a light beam that comes from a low-coherence light source — so not a pure single wavelength, but a broadband light with many wavelengths mixed together. This beam first passes through a collimation lens to make the light travel in a parallel direction. Then it reaches the beam splitter, where the light is divided into two paths. One part of the beam goes upward toward the reference mirror, and the other part travels sideways toward the sample under test.

The light going to the reference mirror bounces back after reflection and returns to the beam splitter. Meanwhile, the light that goes toward the sample also gets reflected — but from within the tissue itself. Think of this like sending the beam into the sample; it penetrates a little, hits various microstructures at different depths, and each layer reflects part of the light back. Those reflected signals travel back toward the beam splitter and are then combined with the reference beam.

Now, depending on the phase difference between the reference signal and the reflected signal from the sample, the two waves interfere either constructively or destructively. This means that at each position — at each depth — the detector receives a slightly different intensity. You measure those variations point by point, and as the reference mirror scans along the axial direction, you record all the interference signals that come from different depths of the sample. Each point corresponds to a small reflection site inside the tissue.

So what happens next is that these interference signals are detected by the photodetector, digitized, and sent to a computer for processing. One depth scan gives you a one-dimensional line — we call that an A-scan. When you move the beam laterally across the sample — for example, in the x or y direction — you collect many A-scans side by side, building up a full cross-sectional image. That's the essential idea of optical coherence tomography — measuring optical interference point by point, and reconstructing the internal microstructure of a sample in depth.

slide25:

Now, let's look at one of the most common and important clinical uses of OCT — eye imaging. OCT is widely used for retinal examination. Here, you can see a cross-sectional OCT image of a normal human retina. The bright layers represent strong reflections — usually the interfaces between different tissue layers — while the darker regions show areas where light penetrates deeper before scattering. You can see here, the scale bar indicates about 250 micrometers — that's just one-quarter of a millimeter — so this technique provides incredibly fine resolution.

Each layer of the retina — from the nerve fiber layer at the top to the photoreceptor layer at the bottom — has its own distinct optical properties. Because of this, the reflected interference pattern from each layer is slightly different. By analyzing these interference patterns, OCT can reveal the microstructure of the retina in exquisite detail.

What I've explained so far is the basic principle, but in practice, it's a bit more sophisticated. You need laser scanning systems to move the beam precisely, and you use Fourier analysis to extract the depth information from the interference signals. By doing this repeatedly and precisely, you can reconstruct full three-dimensional images of the retina. OCT has become a standard tool in ophthalmology, helping doctors detect and monitor diseases like glaucoma, macular degeneration, and diabetic retinopathy.

So, from the same principle we learned earlier — optical interference — you can create a detailed, noninvasive image of the living eye. And not only that, this same technology can be adapted for other parts of the body. For example, you can even make a small optical biopsy probe or a catheter to use inside other organs.

slide26:

Now let's take that idea one step further. Suppose instead of imaging the eye, we use a tiny flexible optical fiber — the same kind of fiber used in telecommunications — and we put it inside the body. That's exactly what's done in intravascular OCT for cardiac studies. A miniature OCT probe is built right into a catheter — that's the small tube used in heart procedures. This fiber can be guided through the blood vessels, all the way to the coronary arteries.

Once the probe is in place, OCT can scan the inside of the blood vessel, creating high-resolution images of the artery wall. With this, doctors can see if there's a buildup of plaque, a blood clot, or any kind of blockage. They can even distinguish between different types of plaque — for example, whether it's hard and calcified or soft and lipid-rich. This distinction is extremely important because soft, unstable plaques are the ones most likely to rupture and cause a heart attack.

So, using OCT, clinicians can visualize the inside of the heart's arteries in real time, with microscopic precision — something that traditional imaging methods like X-ray angiography can't show. It's minimally invasive, it's fast, and it provides information that's both structural and functional.

You can think of this as doing an "optical biopsy" inside the blood vessel — without removing any tissue. This is why OCT has become such a powerful tool, not only in ophthalmology but also in cardiology and other clinical applications. It's an elegant combination of physics, optics, and medicine — a simple idea, but a very powerful and precise technology that continues to transform how we see inside the human body.

slide27:

Okay, how are we doing on time? We've wrapped up the second part, and now we'll continue into the third and fourth parts. In the first part I walked you through the key ideas. We talked about optical microscopy and the confocal idea — a wonderful idea. No Nobel Prize there, unlike OCT, but still pretty amazing. Then we looked at how interference lets you do imaging — not only optical coherence tomography, but interferometric imaging in general. The principle itself is profound, even when there isn't a Nobel attached to it.

You may have read the news about gravitational wave detection. How do they detect those waves? The waves create tiny vibrations — unbelievably small — that you would think are impossible to measure. The trick is interferometry. Two long arms, miles long, two optical paths that pick up a minute difference, and that difference is magnified by interference. There are setups in the United States and Europe, and even programs in China. Some people are pushing to use satellites to make the arms even longer. Interferometry is incredibly helpful whenever you have wave behavior and you want extremely precise measurements.

So the two ideas — confocal microscopy and optical coherence tomography — sit inside this larger family of interferometric imaging. It's not only optical; X-rays and even gravitational wave observatories use the same core idea. In the end, it all goes back to adding waves: cosine omega t plus cosine omega t plus phi. That

simple concept has powered multiple Nobel-level advances. Now, let's turn to something different: diffuse optical imaging. You may not know these methods yet, but you will.

slide28:

Light diffusion is the heart of the third part. These third and fourth parts aren't in common clinical use yet, but in research settings, they're very active. With X-ray imaging, rays go essentially straight. Think of that purple straight line: the signal is a line integral along the path. Everything contributes additively, and you can explain it with the linear attenuation coefficient along that ray. Any change in the measurement must come from features located along that ray — that's a strong localization claim, and it's very clear.

Optical imaging is different. Biological tissue scatters light strongly. I showed you that little cartoon of scattering — Rayleigh at small scales, Mie at larger ones — and you've all seen this in real life. In a dark room, shine a bright laser pointer and you see a diffuse glow. Put the laser behind your finger: can you see details behind the finger? No. With X-ray, the projection already looks like a picture. With optical light, even if you send a parallel laser beam, what you see is a cloudy smear. And yet, diffuse optical imaging says, despite that strong scattering, we still want to reconstruct an image.

So put it in simple words: we want to make a good image behind the finger. That sounds almost impossible, and it is hard. That's why diffuse optical imaging is not as practical as X-ray and some other modalities. But optical interactions carry rich biological information, so we study them. Our job as engineers is to turn "impossible" into "feasible." That's the challenge — and that's why we dig into diffusion models and inverse problems to see what can be recovered.

slide29:

First step: be quantitative about how light propagates in tissue that is highly and strongly scattering. We decompose a piece of heterogeneous tissue into many small elements — in three dimensions, these are typically tetrahedra — and we assume each small element is uniform. Then we launch photons into the mesh and watch what happens.

When a photon hits a boundary where optical properties change, it can reflect or transmit, with angles set by the physics. Once inside a region, it travels some distance, then scatters; all of these events are probabilistic. We use random number generators — this is the dice-rolling part — to decide whether a step reflects or transmits, how far the photon goes before the next interaction, and in what direction it scatters. We trace one photon step by step: maybe it reflects, maybe it refracts, maybe it travels, scatters, and then gets absorbed — good news and bad news, if it's absorbed, you're done with that photon. Or it might exit the tissue and reach a detector.

That's the story for a single photon. Then you launch another photon. And another. You keep going — millions, even billions of photons — to build up statistics. This Monte Carlo approach lets you model refraction, reflection, scattering, absorption, and escape to the detector, all according to the measured or assumed light–tissue interaction properties. From those simulated detections, you can predict measurements and eventually tackle the inverse problem of forming an image in highly scattering tissue.

slide30:

So once you run that Monte Carlo simulation, you can actually form a picture like this. In our lab, we once developed one of the fastest numerical simulators of its kind — it's called a tetrahedron-based, inhomogeneous Monte Carlo optical simulator. Quite a mouthful, right? But what it really means is that we used a mesh made of small tetrahedral elements to simulate photon transport in realistic biological tissue.

Here, you see a simulated mouse model. It's not just a cartoon — it's a detailed anatomical mesh. You can see the heart, lungs, liver, stomach, spleen, and kidneys, all represented as separate organs with different optical properties. We place a small bioluminescent light source — say, something glowing inside the mouse — and then we let it emit photons. Those photons scatter through the tissue, and our simulator tracks where they go and how much light reaches the surface.

On the right, you can see the simulated fluence maps, color-coded to show photon density. The color bar on the side tells you how strong the signal is — from deep blue, meaning very few photons, up through green, yellow, and bright red, meaning a high photon density. It's quite beautiful when you think about it: millions of photons bouncing, scattering, being absorbed — and yet out of all that chaos, you can reconstruct where the light came from inside the body.

This simulator was developed in our group, and I'm proud to say this work has become foundational. In fact, my former student, Qing-Sung, who worked on this, recently got a fantastic offer from Google to continue research in machine learning. So you can see how the skills we learn here — simulation, modeling, and data interpretation — connect to cutting-edge fields. It's not just physics, it's computation, it's innovation. This is the kind of work that combines optical science, biology, and computing into something truly powerful.

slide31:

Now let's talk about what happens at a quantitative level — how photons actually move through scattering media like tissue. When light enters tissue, it doesn't travel far before it gets scattered, because the medium is highly scattering. Statistically, we describe this using something called the mean free path. The mean free path is the average distance a photon travels before being scattered. For biological tissue, that's typically on the order of one-tenth of a millimeter — a very short distance.

Below that distance, we say the photons are in the ballistic region — they travel almost straight, like X-rays, without scattering much. But once they start scattering, things get complicated. The direction changes slightly with each scattering event, and after several scatterings, the photon is completely randomized.

That's why X-ray imaging works so well — the X-rays go straight, so you know exactly which direction they came from. But in optical imaging, photons quickly lose that directional information. You detect a photon, but you can't tell whether it came directly from your target or if it was scattered off several other structures first. It's like hearing an echo in a cave — you know sound reached you, but you don't know exactly from where.

To describe this loss of directionality, we use another parameter called the transport mean free path. After traveling one transport mean free path, the photon's direction has essentially been randomized — it has lost all memory of where it came from. In typical biological tissue, this distance is about one millimeter.

That means that within the first millimeter, you can still recover some useful information — that's where diffuse optical imaging and near-infrared spectroscopy operate. But beyond that depth, it becomes extremely challenging. You lose the ability to form high-resolution images because photons have wandered

too much. It's not totally impossible, but it's very hard. That's one of the fundamental physical limits of optical imaging.

So to summarize: photons in tissue move only about a tenth of a millimeter before scattering, and after about a millimeter, their paths are completely randomized. That's why optical imaging is powerful for shallow structures — like the cortex, or the retina — but much more difficult for deep tissue imaging. Later, we'll see how researchers combine optics with X-rays to overcome that limitation.

slide32:

Now let's move to another branch — diffuse optical spectroscopy, or DOS. And yes, I know what some of you are thinking — DOS sounds like the old computer operating system. If you remember that, you're not that young! But in this case, DOS stands for something different — diffuse optical spectroscopy.

Here, we're not trying to make an image. Instead, we're doing spectral analysis — we shine light in, we collect what comes out, and we analyze the overall absorption and scattering to learn about the tissue's composition.

Imagine my finger here. If I shine a broadband light — that means light with many colors — through my finger, different wavelengths will be absorbed differently. Blue light, for instance, is absorbed more strongly than red, which is why your finger looks red when you shine a flashlight through it.

Now, by measuring how much light gets through at two or more wavelengths, we can solve for unknowns in the tissue. Usually, the main absorber in tissue is blood, and in blood, the key molecules are hemoglobin and oxyhemoglobin — that's hemoglobin with oxygen bound to it. The ratio between those two forms gives us oxygen saturation, which is one of the most important physiological parameters for life.

Mathematically, it's quite straightforward. The transmitted intensity follows Beer's law — I equals I naught times e to the power of minus mu-a times L, where mu-a is the absorption coefficient and L is the path length. The absorption coefficient depends on how much hemoglobin and oxyhemoglobin are present and how strongly each absorbs light at that wavelength. By measuring intensity at two wavelengths, we can solve for two unknowns — the concentrations of Hb and HbO_2.

In fact, this is exactly the principle used in a pulse oximeter — that little clip you put on your finger in the hospital. It shines red and infrared light through your finger, measures how much of each gets through, and calculates oxygen saturation using this model.

So, this is diffuse optical spectroscopy. It doesn't form an image, but it gives us valuable quantitative information — blood oxygen levels, water content, and even tissue metabolism. In a sense, it's the simplest and most practical use of optical interaction with tissue — turning light absorption into a window on physiology. And the best part? It's completely noninvasive — just light in, light out, and a bit of math in between.

slide33:

Alright, now the second idea here is what I call the tiredly idea — it's really a continuation of the first part. We want to do tomography. We really want to do tomography — just like we do X-ray tomography. So, let's think about how we actually do X-ray or gamma-ray tomography.

In those systems, you send a pencil beam — a narrow, collimated beam — straight through the object. What you measure is basically a line integral. For nuclear imaging, it's sometimes a quasi-line integral because the photons are affected by an attenuation kernel, but fundamentally, the information you get comes from along that single line.

So, you get one line-integral measurement, and that becomes one linear system equation. In that equation, all the unknowns — those pixel or voxel values — are weighted according to how much of the beam passes through each small region. That weighting is the fraction of the path that cuts through each voxel. So, mathematically, it's like this: one unknown times its weighting factor, plus another unknown times its weighting factor, plus another, and so on — all along that beam line. That gives you one linear equation.

Now you can see how this builds up. You have many, many linear equations — one for each ray, one for each angle. The weighting factors are all known because they come from the imaging geometry — the way you arrange your source and detector. The only unknowns are all those mu values — the attenuation coefficients inside each pixel or voxel — which we call mu one, mu two, mu three, and so on.

Now, how does this relate to the Beer–Lambert law? Let's go back to that. You have an incoming intensity, called I naught, and you measure an outgoing intensity, I. According to Beer's law, the transmitted intensity equals the input intensity multiplied by e to the power of negative mu times delta, where mu is the attenuation coefficient, and delta is the distance through that small region.

So, suppose the beam passes through one region with attenuation mu one and thickness delta one, then into a second region with mu two and thickness delta two, and so on. After the first region, the output of that layer becomes the input for the next. Multiply all those together, and inside the exponential, you get the total, negative the sum of mu times delta over all the layers.

So mathematically, we say I equals I naught times e to the power of minus the sum of mu times delta. That's what we measure. Now, if we take the natural logarithm of both sides, it becomes linear. Log of I naught over I equals the sum of mu times delta. That's the key point: each measurement gives you one linear equation in the unknown mu values.

All the geometry — those delta distances — are known, and the intensities I and I naught are measured. So that's how we form the reconstruction problem for tomography.

Now, this all works nicely for X-ray imaging, where photons travel in straight lines. But in optical tomography, things are very different. When we send optical light into biological tissue, we still have all those local attenuation coefficients, all those mu's — but the light doesn't travel straight. It scatters everywhere. So we can't just send one clean laser beam, measure on the other side, and call that a line integral. The light spreads out in all directions, taking many different paths.

Each measurement still has weighting factors, but now those weighting factors are far more complicated because the light is diffusing rather than traveling straight. And when we add up all those equations, they still look linear in the unknown mu's, but all the equations start to look very similar to one another. That's the main problem.

It's like this: suppose you have three unknowns — x, y, and z. And all your equations look like x plus y plus z equals one. Then the next one says x plus y plus z equals one point zero zero one. Then another says x plus y plus z equals zero point nine nine nine. You have a lot of measurements, yes, but they're almost identical. You can't solve for x, y, and z because the equations don't give you independent information.

In tomography, we call that an ill-posed problem. You have a lot of data, but not enough unique or independent information to get a stable solution.

So the analogy here is that in optical tomography, all the light measurements look very similar because scattering smears out the information. Each detector sees a blurred, mixed-up version of everything inside. The equations are still linear — yes, still linear — but they're all nearly the same, so it becomes very difficult to separate one unknown from another.

And that's the essence of the challenge. Too much diffusion means too little unique information. That's why optical tomography is so difficult. It's mathematically ill-posed, physically diffusive, and highly sensitive to noise. The principle is beautiful, but the reconstruction is tricky.

slide34:

But researchers don't give up easily, right? They keep trying. And they've managed to make progress using infrared light, because infrared wavelengths are longer and scatter less than visible light. That gives you better penetration into tissue. So you can still get a measurable signal, even from deeper layers. And then, using some clever mathematical tricks, you can try to recover the internal absorption and scattering distribution — the actual optical properties inside the tissue.

Now, think about how useful that could be. Suppose you have a tumor. In that region, the absorption and scattering characteristics will be different from the surrounding healthy tissue. Water content, fat content, and hemoglobin concentration — all these change when tissue becomes cancerous. For example, cancerous tissue tends to have more hemoglobin, because tumors grow new blood vessels to feed themselves — we call that angiogenesis. So, in optical terms, the tumor region is "hemoglobin-rich."

So, people have written many papers using this principle. Under certain experimental conditions, the results do make sense — the reconstructed maps of hemoglobin concentration and oxygenation often agree with what's seen in pathology. You might see a region with higher hemoglobin absorption — that's where the tumor is. It's exciting work. But the question remains: has it reached clinical use? Not really — at least, not yet.

After many years of research, diffuse optical tomography is still not mature. The principle is solid, the physics is beautiful, but the system is what we call ill-posed — the solutions are unstable, sensitive to noise, and dependent on assumptions. It's just not as robust or reliable as X-ray imaging.

And that's an important comparison. We're all very proud of our X-ray imaging technology. Every hospital depends on it — CT, fluoroscopy, mammography — they're all part of daily clinical life. If you took X-ray imaging away from a hospital, they'd practically have to close their doors! It's that critical. So when you ask a doctor, "If you could only keep one imaging modality in your clinic, which would it be?" — none of them will pick optical imaging.

They'll all say, "We'll keep X-ray." That's the truth.

So, yes — optical tomography is very interesting. It's innovative, and it provides functional and molecular information that X-rays can't. But it's still a research tool. X-rays, on the other hand, are robust, quantitative, and universally trusted in clinical medicine. That's where we stand today.

slide35:

So anyway, one way to improve the localization capability — meaning how precisely we can tell where the light comes from — is through something called temporal gating. The idea is very clever. You see, when photons travel through tissue, they go all over the place — scattering in every direction. But if we can measure when each photon arrives, and only collect the ones that arrive early, we can get a much cleaner signal.

For optical signals, the temporal resolution — that is, the time resolution — of modern detectors is extremely good. So we can actually do this kind of time gating. Temporal gating simply means we only collect the photons that arrive first, before they have time to bounce around and scatter too much.

In other words, those early-arriving photons are the ones that have traveled almost straight through, or maybe with just a few scatterings — not too many. So by keeping only those early photons, we're kind of reconstructing a straighter path through the tissue. This gives us a better localized image and sharper information. That's the essence of temporal gating — a brilliant idea to improve spatial resolution in optical imaging.

Now, you can think of this as a kind of collimation mechanism. Earlier, we talked about mechanical collimation, where you physically restrict the beam using a tube — that's what we use in X-ray systems. I jokingly said, It's my 'technical way' of collimation — you just put a tube in front of the detector so only straight rays go through.

And then we also mentioned electronic collimation, which you might remember from the pattern-imaging systems or coincidence detection in nuclear imaging. There, you don't use a physical tube — you use timing electronics, or coincidence circuits, to select only the photons coming from the right direction.

Now here, temporal gating plays a similar role — but for light. Instead of a physical or electronic filter, we use time as the filter. We only accept photons that arrive early enough, meaning those that haven't wandered too much. That's why I call it a "temporal collimation" mechanism.

And remember, I also mentioned, a while ago, that you can do something similar with polarized radiation? In that case, we used a magnetic collimation trick — flipping the spin of particles back and forth to control their direction, kind of "cheating" to get the same effect.

Anyway, the big picture is this: optical imaging is not easy. It's still a very active research topic. Even now — well, I think maybe early this year, or perhaps late last year — my group published a paper about improving optical tomography using advanced computational techniques. So we're making progress. But I must say, it's still far from routine clinical use. It's fascinating physics and engineering, but not yet something you'd see in a hospital every day.

slide36:

Now, let's move on and talk about fluorescence molecular tomography — often abbreviated as FMT. Fluorescence and bioluminescence imaging are both very important tools, because they allow us to use fluorescent or luminescent probes to label biological molecules — like proteins, genes, or even drugs — so we can track them in real time.

Most of the time, we use small animal models for this. Small animals — mice, rats — are used to model almost all kinds of human diseases. We don't want to experiment directly on humans, of course, so we test

our ideas on these animals first. We give them certain diseases — cancer, bone degeneration, inflammation, all sorts of things — and then we test how treatments work.

Once the biological principles look promising, we test drugs on these animal models to see if they behave as expected before we move on to human trials.

So, for example, suppose we label cancer cells with a red fluorescent protein. Then we label a drug molecule — maybe an antibody or some therapeutic compound — with a green fluorescent marker. Now, if the green-labeled drug binds to the red-labeled cancer cells, that means the drug has successfully targeted the tumor. In that case, the red signal would decrease or disappear, and we could monitor that process in vivo, meaning inside the living animal, in real time.

That's the beauty of fluorescence molecular tomography — we can visualize biological processes directly, without opening the body.

Now, if you look at the illustration here, you see the setup. The animal is standing upright, which is not really a physiological position — ideally, the animal should be lying down to minimize stress. But this figure is just to show the principle of the system.

Inside the animal, suppose both the cancer cells and the drug molecules are labeled with different fluorescent proteins — maybe blue, green, yellow, or red, depending on what we want to track. At first, you might not see any fluorescence signal on the surface. But then, we shine in a femtosecond laser beam — that's an extremely short, high-intensity pulse of light — at a specific wavelength. That excites the fluorescent molecules inside the body.

Those molecules then emit light at longer wavelengths — that's the fluorescence emission. We collect those emission signals using an optical detector, usually a highly sensitive ICCD camera — that's an intensified charge-coupled device.

Now, the emission filter plays a very important role. It blocks the excitation light — the laser light that we sent in — and only allows the emitted fluorescent light to pass through. This way, we measure only the signal that comes from the fluorescence, not the original illumination beam.

So, in the end, you get a set of projection views — one from each rotation angle — just like a CT scan. You rotate the animal, collect data from multiple views, and reconstruct a 3D distribution of fluorescent activity inside the body.

Of course, it's not as straightforward as CT, because light in tissue is highly scattered, so the reconstruction problem is much harder — not very stable, not very unique, and sensitive to noise. But still, you can get useful information by combining advanced algorithms and multiple measurements.

So this system — using a femtosecond laser, a rotating stage, and an ICCD detector — is an example of fluorescence molecular tomography in small animals. It's a powerful tool for biological and pharmaceutical research, allowing us to visualize how drugs move, how genes express, and how diseases progress — all in a living system.

slide37:

Alright, so here we can see an example of what's called FMT reconstruction, meaning fluorescence molecular tomography reconstruction. You can see the image of the animal here — this is usually a mouse,

and we've done the imaging after injecting fluorescent probes. So, you collect the signals from different views, and then you reconstruct a 3D distribution of those fluorescent signals inside the body.

Now, some research groups — and I've seen quite a few papers like this — claim that they can do very high-quality reconstruction. And yes, to a certain degree, they can. These results look pretty good, and they do show useful biological information. But, to be honest, the results are still not very reliable yet. The reconstruction quality is limited by noise, by the scattering of light, and by how well the model matches the real tissue.

If you look at this picture, you can see the 3D reconstructed image with color bars on the right side. The colors correspond to different levels of fluorescent intensity. So, when you look at the image, you get a sense of where the fluorescent probe has accumulated inside the animal's body. This helps identify regions of biological activity — for example, where a particular enzyme is expressed, or where a drug is localized.

This example shows a three-dimensional reconstruction of fluorescence molecular tomography — FMT. So, yes, it's a practical imaging modality, and it really works, but it's not for human use. There are a few reasons for that.

First, most of the fluorescent probes used in small-animal imaging are not safe for human use — they're either toxic or not approved for clinical injection. Second, there's a size issue — humans are much bigger than mice, and optical light simply cannot penetrate that deeply into human tissue. The penetration depth of laser light is very limited.

So, for now, FMT is mainly a preclinical imaging technique — something that's used in research, mostly for small animals. It's great for drug development, molecular biology, and early disease studies, but not yet practical for clinical imaging in humans.

slide38:

Now, before I joined RPI, I was actually on the faculty at Virginia Tech and Wake Forest University, both at the medical school and the engineering school. I had laboratories on both campuses. And one of the major projects we worked on — a multi-million-dollar project — was about using optical molecular tomography to monitor blood vessel growth.

Let me explain the motivation behind it. When someone has cardiovascular disease, part of a blood vessel can get blocked — completely closed off by plaque buildup. Right now, the standard treatment is to insert a stent to reopen the vessel and restore blood flow. But as you know, stents come with their own problems — they can cause inflammation, thrombosis, and sometimes even re-narrowing of the artery over time.

So, the future solution that people are exploring is regenerative medicine — to actually grow a new biological blood vessel from your own stem cells. The idea is that we can take a small segment of your blood vessel, or even some of your cells, and grow a new vessel in the lab — a fully biological, living blood vessel that can eventually be implanted back into your body.

To do this, we use something called a bioreactor — a special device that maintains controlled conditions for growing tissue. Inside the bioreactor, you have this new blood vessel growing while nutrient-rich liquid continuously circulates it, simulating what happens inside the body.

Now, the question is — how do you monitor this growth process without destroying the tissue? That's where our imaging technology comes in.

We use OCT, or Optical Coherence Tomography, to capture the structural information — basically, the microscopic anatomy of the growing vessel. You can actually see the layers, the texture, and the thickness of the vessel wall.

Then we combine that with fluorescence imaging to observe what's happening at the cellular and molecular level. For example, using fluorescence probes, we can track whether endothelial cells — the cells that form the inner lining of blood vessels — are growing properly. That's critical because without a strong endothelial layer, the vessel won't be functional or stable.

So, in our project, we integrated all these optical methods — OCT for structure, diffuse optical tomography for optical property recovery, and fluorescence tomography for molecular activity — all working together.

The advantage is that once you know both the structure and the optical properties of the tissue, you can use that information as a prior to improve the accuracy of fluorescence tomography. That means the reconstruction becomes more stable, more accurate, and better localized.

And since these vessels are much smaller than an entire small animal — just a few millimeters in diameter — the light can penetrate through them more easily. So we can actually achieve better resolution and depth information.

This whole system, combining multi-probe and multi-modal optical molecular tomography, was designed for regenerative medicine — specifically, to visualize and monitor bioengineered blood vessels both in the bioreactor and after implantation into living animals.

So that's what we did a number of years ago — a collaboration between my lab and my colleague, Dr. Shay Soker. It was supported by the NIH R01 BRP grant HL098912, running from 2010 to 2014. And that's a really good example of how optical molecular imaging can move beyond small-animal studies toward real applications in regenerative medicine — imaging living tissue as it grows, heals, and integrates into the body.

slide39:

Now, this next imaging modality is what we call bioluminescence tomography. It's one of the most fascinating — but also one of the most challenging — tomography problems we can deal with.

Let me explain why.

In fluorescence tomography, you can send in a laser beam — you control where it goes, you can illuminate the tissue from different directions, and that makes it what we call an active imaging modality. You actively excite the fluorescent molecules, and they emit light back, which you can detect and reconstruct into an image.

But bioluminescence tomography is totally different. Here, the light comes from inside the animal itself — not from any external laser. You don't get to shine a laser and choose where to excite. The light is produced biologically, inside the tissue, by a chemical reaction, typically involving something like luciferase, the same enzyme that makes a firefly glow.

So, imagine you genetically modify an animal — usually a mouse — to express this bioluminescent probe inside its cells. The animal literally becomes like a little living firefly. When the probe reacts inside its body, it emits light that escapes through the tissue, and you can capture that light on the surface.

That's beautiful to watch — I still remember when we did this experiment back at the University of Iowa. We worked in a completely dark room, and when the mice began to glow faintly, you could actually see the light coming through their skin. It was really amazing.

But scientifically, this is very challenging, because you can't control the direction of illumination — there is no external laser input. All you can do is measure the light that escapes to the surface. That makes it a purely passive imaging problem, and therefore much harder than fluorescence tomography.

So, the question is: can we reconstruct where the bioluminescent light is coming from inside the animal, based only on what we observe on the surface? That's the big challenge of bioluminescence tomography.

When I was at Iowa, we actually got a multi-million-dollar NIH grant to tackle exactly this problem. My idea was to use a multi-modality approach — to combine optical imaging with micro-CT or micro-MRI so we could obtain both the anatomical structure and the optical measurements together.

Here's how it works. First, we perform a micro-CT or micro-MRI scan of the animal. That gives us very detailed anatomical information — we can segment the bones, the organs, the tissue layers. We then take that structure and build a finite element mesh, so now we have a 3D computational model of the animal's anatomy.

Next, we perform optical measurements — diffuse optical tomography — to estimate the optical properties of the tissues, like how much they scatter and absorb light. Once we have both the anatomical map and the optical map, we put them together. Now we have a very detailed, voxel-by-voxel model — a so-called prior model — of the animal's body.

With that model, we can perform Monte Carlo simulations. This is similar to the photon simulation I showed you earlier. You put a light source at a certain location inside the model, and then you let thousands — or even millions — of photons propagate randomly in all directions. Each photon scatters, reflects, refracts, or gets absorbed, following the physical laws of light propagation in tissue. Eventually, some photons reach the surface, where they are detected.

This process tells us, for any assumed source position, what the light distribution on the surface would look like. We call that the forward problem — going from a known source to the expected surface measurement.

Then, we flip the problem around. This is the inverse problem — we already know the surface light pattern (that's what we measured), and we want to find the internal source distribution that could have produced it.

So we try different possible source locations — say, one in the middle of the body, one near the liver, one near the kidneys — and simulate each case. We compare the simulated surface pattern with the measured one. If they don't match, we adjust the model, move the source slightly, and try again.

After many iterations, when the simulated pattern matches the measured pattern well, we can say with confidence that the actual bioluminescent source is located in that region inside the animal.

I remember one particular case — we found two bioluminescent spots, one near the upper right and one near the left kidney. When we later sacrificed the mouse and examined it, we indeed found two small tumor nodules exactly in those locations. It was like magic — a scientific magic trick!

We even published that work in Optics Express, showing this successful reconstruction. It was a very satisfying result.

However, as with all research, the story isn't all perfect. The reconstruction is extremely sensitive to errors in the model. If your optical properties are slightly off, or your anatomical registration isn't perfect, the reconstructed source location can be wrong. That's why, although it's feasible, it's not as robust or as straightforward as, say, X-ray CT.

So, to summarize — bioluminescence tomography is feasible, it works, but it's still very much a research topic, not yet a clinical tool. It's beautiful, it's challenging, and it's fun to work on. And along the way, we also contributed several new ideas to this field — such as combing tomography, helical scanning, interior tomography, and other concepts that pushed the boundaries of imaging research.

So yes — this was, and still is, a very enjoyable area of scientific discovery.

slide40:

Now, let's talk about another idea — a really interesting one — that tries to make fluorescence tomography and bioluminescence tomography even better. The idea here is to combine optical imaging with X-ray imaging — a kind of hybrid approach.

You see, the limitation with traditional optical tomography is that it doesn't give you good structural information. You can get molecular or functional data — yes, you can tell where the light comes from or which cells are active — but the anatomical detail is missing. On the other hand, X-ray CT gives you beautiful structure, very precise geometry, but not the molecular activity.

So, what if we could merge the two? What if we could combine the molecular sensitivity of optical imaging with the structural precision of X-ray imaging? That's the motivation behind this work — we call it X-ray Optical Fusion.

Now, traditional micro-CT — the standard kind — has a limitation. It's great for bone, great for dense structures, but not so good for soft tissue. The contrast inside soft organs is very weak because the attenuation difference is small. So, you can't see the fine internal structures clearly.

That's why researchers have started exploring X-ray Phase-Contrast CT, or PCCT. This is a newer type of X-ray imaging that doesn't just rely on attenuation — it actually detects the phase shift of the X-rays as they pass through tissue. This phase information is much more sensitive to subtle density changes, especially in soft tissue.

So, the hybrid idea is: combine Fluorescence Molecular Tomography (FMT) with X-ray Phase-Contrast CT. Together, they can provide both the functional molecular information from FMT and the high-resolution anatomical detail from PCCT.

There's a paper on this — you can see it here — titled "FMT–PCCT: Hybrid Fluorescence Molecular Tomography and X-ray Phase-Contrast CT Imaging of Mouse Models." It was published in IEEE Transactions on Medical Imaging in 2014 by researchers from the Helmholtz Center in Munich and collaborators in Germany. This kind of work really represents the next step — hybrid optical-X-ray imaging for small animal studies.

slide41:

Now, some of the results reported in this area are quite interesting. In these experiments, the researchers scanned small animals — typically mice — sometimes for an entire day, to acquire very high-quality data.

Instead of using conventional micro-CT, they used a grating-based phase-contrast tomography system. I actually mentioned this technique before, when we talked about my visit to Japan — the grating-based interferometric system. The setup uses a series of gratings — a source grating, a phase grating, and an analyzer grating — to measure very tiny changes in the X-ray phase as the beam passes through the sample.

So, what they did was essentially scan the animal — or tissue samples — using this grating-based X-ray phase-contrast CT. And they achieved beautiful soft-tissue visualization, something you simply can't get from normal X-ray attenuation imaging.

These results, published by several European research groups, clearly demonstrated that grating-based phase-contrast tomography can reveal fine internal details — such as blood vessels, organ microstructures, and even tumor margins — without the need for any contrast agents.

It's a big step forward for optical-X-ray fusion, because now the X-ray side of the hybrid system provides much richer anatomical detail to guide and constrain the optical reconstruction.

slide42:

At RPI, we've been taking this concept even further. We're working on developing a tightly integrated system for in vivo optical and X-ray imaging — that means imaging live animals, not just tissue samples.

In the earlier work, what we call ex vivo imaging, the studies were done on isolated or preserved samples, which is great for testing physics, but not for observing biological dynamics. So, our goal is to move toward in vivo imaging, where we can study living processes in real time.

Our approach is to build an orthogonal imaging chain — basically, an X-ray phase-contrast system aligned at a right angle to an optical imaging system. One part of the setup handles the phase-contrast X-ray imaging, while the other captures fluorescence or bioluminescence signals. By merging these two datasets, we can combine the strengths of both modalities — structural accuracy from X-rays and molecular sensitivity from optics.

The system includes components like a CCD camera, filter wheel, laser stage, isoflurane anesthesia line for live animal support, and a rotating gantry for tomographic data collection. The optical part uses mirrors, spectrometers, and digital micromirror devices for detecting and filtering the light.

So, what we're working toward at RPI is a fully integrated hybrid imaging system — one that can acquire X-ray phase-contrast data and optical molecular data simultaneously.

Now, this is still an active research topic. Achieving truly precise, stable 3D tomography in such a hybrid setup is not easy. There are still open questions — how to synchronize the modalities, how to register the datasets, how to compensate for motion, and how to achieve stable reconstructions.

But it's an exciting direction. When you can successfully merge X-ray and optical data — one showing the anatomy, the other revealing the function — that's when you get a truly powerful multimodal imaging system.

So, this is what we're actively working on at RPI — a tighter, more integrated system that pushes the boundary of X-ray–optical fusion imaging. It's challenging, but also very rewarding work.

slide43:

Alright, in this last part, let's talk about something quite new and exciting — X-ray optical coupling. If you look at the slide, you can see the last two names listed here: XLCT and XMLT. These stand for X-ray Luminescence Computed Tomography and X-ray Molecular Luminescence Tomography.

Now, what we're doing here is merging X-ray imaging and optical imaging — two worlds that traditionally don't overlap much. But when you think about it, it makes perfect sense. X-rays penetrate deeply and travel straight, while optical imaging provides rich molecular and functional information. So, if we can somehow couple these two, we can get deep, high-resolution, molecularly sensitive images — the best of both worlds.

So this is what we mean by X-ray Optical Coupling — using X-rays to stimulate optical signals inside tissue, and then using those emitted optical photons to form an image. It's a very elegant idea, and it's starting to gain serious research attention.

slide44:

Okay, so let's go a bit deeper into the physics and materials behind this idea. This concept really reminds me of how fluorescent and bioluminescent proteins work. In fluorescence imaging, you have a fluorophore — say, a protein that absorbs one color of light and emits another. In bioluminescence, you have a chemical reaction that produces light on its own, like fireflies.

Now, what researchers have done is introduce a new type of material, called a nanophosphor — it's not a protein, not biological by nature, but rather an engineered material. Nanophosphors are nanoscale particles that have unique luminescent properties.

Here's how it works. When you shine an X-ray beam onto these nanophosphors, they emit visible or near-infrared light — in other words, they glow. You can even make different kinds of nanophosphors that emit different colors — green, red, near-infrared, far-infrared — depending on the specific material and dopant you use.

So, it's very much in parallel with fluorescence imaging, except for one big difference: in fluorescence, you excite the molecules using laser light, which is itself light, and therefore it scatters heavily inside tissue. But in this case, the excitation comes from X-rays, which travel straight through the body. So, the spatial resolution is much better because X-rays are not affected by scattering the way light is.

One example of this can be seen in the study shown here — Photostimulated Near-Infrared Persistent Luminescence as a New Optical Readout from Chromium-Doped Lithium Gallium Oxide. This material, when irradiated by X-rays, stores the energy and then later emits near-infrared light when stimulated by low-energy light.

This concept — using X-ray-excitable nanophosphors — opens the door to a completely new kind of optical imaging. Instead of laser excitation, we use X-rays, and the resulting emitted light can be collected optically. That's the core idea behind X-ray Optical Coupling.

slide45:

Now let's look at this next example — this is an article introducing X-ray Luminescence Computed Tomography, or XLCT. This is a very creative and forward-looking concept.

The essential idea is shown right here. You introduce X-ray-excitable nanophosphors into a small animal — usually a mouse. These nanophosphors can be functionalized — meaning you can coat their surface with polymers or peptides to make them target specific biomarkers. For example, if you want to target cancer cells, you attach a peptide that binds to receptors that are overexpressed in tumors.

So now, when these nanophosphors are injected into the body, they accumulate more densely in the tumor region than in normal tissue. Then, when you shine an X-ray beam, it goes straight through the body. Wherever it hits these nanophosphors, they emit optical light. And that light — the X-ray-induced luminescence — is detected by sensitive optical cameras around the animal.

So think about it: by scanning the X-ray beam and recording the emitted light, you can reconstruct the 3D distribution of these nanophosphors inside the body. That gives you a molecular image that is excited by X-rays and detected optically.

This is a very promising approach because it potentially makes X-rays a molecular imaging modality — something traditionally reserved for optical, nuclear, or MR imaging. You can now use X-rays not just for structure, but also to identify specific molecular targets.

Of course, there are still challenges. When you collimate the X-ray beam — that means narrowing it to a small pencil beam — you face physical limits. Mechanical collimators can only get so thin, and even a narrow beam will spread out slightly due to scattering and beam divergence. So the current spatial resolution of XLCT is about 1 to 2 millimeters.

Now, the question becomes: how can we improve that? How can we go from millimeter resolution down to, say, 100 microns — an order of magnitude better? That's still an open research problem. It will require new hardware designs, smarter reconstruction algorithms, and probably better nanophosphor materials.

But the vision is clear — this approach could bridge the gap between anatomical imaging and molecular imaging, between optical precision and X-ray penetration. It's one of those rare ideas that brings together two worlds — physics and biology — into a single, unified imaging platform.

So that's the direction of X-ray Luminescence CT and X-ray Optical Coupling — still developing, but with tremendous potential for the future of biomedical imaging.

slide46:

Alright, now let's move on to something very exciting — X-ray micro-modulated luminescence tomography, or simply XMLT. This is one of the latest developments that builds upon everything we've discussed so far.

If you're really interested, I recommend reading the original paper. But let me give you the idea in simple terms.

In XMLT, we treat the X-ray beam as a wave, not just as a stream of particles. That means we can actually use certain X-ray optical components — things that already exist on the market, like X-ray polycapillary lenses. These are special lenses that can focus X-rays, although not quite as tightly as optical lenses can focus visible light. Still, they can concentrate the X-ray beam to a small spot — around 100 microns, or even smaller in some setups.

And because X-rays penetrate tissue so well, this focused beam can reach deep inside the object. So now, imagine we focus the X-ray into the tissue — we're creating a tiny, well-defined excitation region.

When we introduce nanophosphors into the sample — the same luminescent particles we talked about earlier — these particles light up only along the path of the focused X-ray beam. Everywhere else remains dark. That's very important because it means the optical signal we collect corresponds precisely to the X-ray excitation path — no background interference from outside that region.

So the spatial resolution of this imaging method is determined entirely by the X-ray focal spot size, by how small that beam is. Think of it like drawing with a very fine pen instead of a broad brush.

And geometrically, this is not a pencil beam or a fan beam like in standard CT. It's more like a double-cone beam — a cone that converges to a focus point and then diverges again. So you get a cone of excitation in both directions.

In short, we perform tomography by collecting optical signals generated from these small, localized regions excited by the focused X-ray. We have published several papers on this topic, exploring the physics, the reconstruction algorithms, and the system design. It's a fascinating direction for hybrid imaging research.

slide47:

Now, we've been collaborating with General Electric — GE — to take this idea further.

With a micro-focused X-ray beam, we can excite nanophosphors deeply inside small animals, such as in the brain. You can imagine a setup like this: the X-ray beam enters an integrating sphere — that's a spherical chamber lined with a highly reflective surface. It has just one small opening where the X-ray beam enters and where the emitted photons exit.

Here's how it works. The X-ray beam excites the nanophosphors in a very small region. Those nanophosphors emit optical photons — visible or near-infrared light. The integrating sphere reflects these photons multiple times — bouncing them around the inner wall — so that, eventually, almost every photon finds its way to the detector, usually a PMT, or photomultiplier tube. This design allows us to collect nearly all the emitted light with minimal loss.

That's how we can do small-animal imaging efficiently.

Now, if we scale up this concept, we can also envision doing human imaging, especially for the brain. The human cerebral cortex has about six layers, and together they're roughly six millimeters thick. The interesting part — the one related to higher-level brain function and intelligence — lies not just on the surface, but deeper in those inner cortical layers.

So, how can we noninvasively reach those deeper layers? Well, one potential path is to use nanophosphors that can cross the blood–brain barrier and label specific neural structures. Then, we use a micro-focused X-ray beam to excite those particles deep inside the brain.

Here's where it gets really exciting — and, yes, a little bit like science fiction. When neurons fire, they produce action potentials — tiny electrical currents. These currents can interact with the nanophosphors, changing their quantum emission properties — for example, the color or intensity of the light they emit.

So in theory, by monitoring how the emitted light changes, we could infer neural activity — we could literally watch the brain think. Imagine detecting color changes corresponding to neural firing patterns!

Now, I'm not saying we're there yet — far from it. This is still in the conceptual stage. But it's a very cool direction to pursue — combining X-ray excitation, optical emission, and nanophosphor sensing to probe deep brain function noninvasively. It's futuristic, but scientifically grounded.

slide48:

Let's shift gears and talk about another exciting technology — optogenetics.

Optogenetics is a biological technique where light is used to control cells — typically neurons — that have been genetically modified to express light-sensitive ion channels. Depending on the wavelength of light you use, you can either activate or inhibit neural activity.

For example, blue light — around 470 nanometers — can open certain channels, allowing positively charged ions like sodium or calcium to enter the neuron, thereby triggering an electrical signal. On the other hand, yellow or red light — around 570 or 590 nanometers — can open other types of channels that inhibit activity, stopping the neuron from firing.

So, by simply changing the color of light, you can make a neuron fire or stop firing. It's almost like using a remote control for the brain!

Researchers have used this technique to study complex neurological conditions — things like depression, Parkinson's disease, and addiction — where precise control of neural circuits is essential. You can shine light on specific regions of the brain and observe behavioral changes in animals — say, a mouse moving left when you use one color, or right when you use another.

However, there's a limitation. Optical light doesn't penetrate deeply — it only reaches about one millimeter into tissue. If you want to reach deeper layers, you have to insert an optical fiber directly into the brain. And, of course, that's invasive — it causes tissue damage, and you can't easily use it for large-scale or long-term experiments.

So, while optogenetics is an incredibly powerful tool, its reach is still limited by the physics of light scattering.

slide49:

Now, that brings us to our next idea — something we call X-optogenetics.

My students and I proposed this concept a few years ago, and we even published an article titled "X-Optogenetics and U-Optogenetics: Feasibility and Possibilities."

The idea is simple but powerful. We know that X-rays can penetrate deeply into tissue, and we also know that nanophosphors can convert X-rays into visible light. So, what if we could use X-rays as a remote light source to activate optogenetic proteins deep inside the brain?

Here's how it would work. We inject nanophosphors small enough to pass through the blood–brain barrier — so they can distribute evenly within the brain tissue. Then, we shine a focused X-ray beam into a specific region. The nanophosphors there absorb the X-rays and emit light locally — right where the neurons are.

That emitted light can then activate the optogenetic channels — just like in traditional optogenetics, but now without inserting any optical fibers. Everything is done noninvasively.

In our paper, we discussed this as a possible future imaging and neurostimulation technique. One of my undergraduate students — who is now working full-time on this — helped develop the idea. We're currently characterizing different nanophosphor materials, testing their emission spectra, decay times, and compatibility with biological tissue.

This concept could open the door to deep-brain stimulation without surgery, combining the precision of optogenetics with the penetration of X-rays. It's a very exciting frontier — truly where physics meets neuroscience.

slide50:

Alright, that brings us to the end of today's lecture.

For homework, I'd like you to do three things:

First, review this lecture carefully and summarize the key ideas — not just a quick summary, but really try to capture the main concepts and how they connect.

Second, transcribe one section of this lecture — you can pick the first part, the second part, or the last two parts — and write it down clearly. This will help reinforce your understanding.

And third, here's a fun one — a creative thinking question. Imagine you could make a smartphone that can send and receive light in any way you want. What kind of medical imaging applications could you create?

Think about it — a smartphone already has a light source, a detector, computing power, and connectivity. What if it could emit polarized light, or detect fluorescence, or measure oxygen levels through tissue? Could it monitor glucose levels? Could it perform optical tomography?

You can be as creative as you want — just stay within the limits of physics.

So, that's your assignment: review, transcribe, and imagine. I'm really looking forward to seeing your ideas. And that's all for today — thank you!