slide1:
Hi everyone, and welcome to this online course on Medical Imaging. I'm Professor Ge Wang, and I'll be guiding you through this exciting journey.
Medical imaging is a fascinating and powerful field. It brings together science, engineering, and medicine to help us see what's happening inside the human body—without making a single cut. Whether your background is in engineering, physics, computer science, or biology, you'll find that medical imaging offers something valuable and fascinating.
In the weeks ahead, we'll build a strong foundation in the core principles of imaging and describe a variety of technologies that are transforming modern healthcare. I'm excited to get started—and I hope you are too.

slide2:
We'll begin with the fundamentals — systems, Fourier analysis, signal processing — and hands-on MATLAB sessions. Then we'll explore major imaging techniques: X-ray and radiography, CT, PET, SPECT, MRI, ultrasound, optical, and deep imaging. Along the way, we'll have three exams to help you review and reinforce what you learn
Now, just a quick note: in a traditional classroom, I'd mention office hours; however, since this is an online course, I encourage you to participate actively in our discussion forums. Don't hesitate to post questions or thoughts as they come up. Learning is always more effective when it's interactive.

slide3:
So, to kick things off, let me give you a quick overview of what we'll cover today.
There are four main questions we'll explore:
First, what is medical imaging?
Second, why should we study it?
Third, who are involved in this course?
And finally, how are we going to approach the material?
Now, when we talk about medical imaging—sometimes we use the terms biomedical imaging or bio-instrumentation—we're referring to a wide range of methods, technologies and systems that help us see inside the human body, animals, or biological samples. These methods are grounded in engineering, physics, and biomedicine.
And this isn't just about theory. You'll also get hands-on experience with tools like MATLAB, and you'll work with selected reading materials and online resources. By the end of the course, you'll not only understand how these imaging systems work but also be able to apply that knowledge in a meaningful way.

slide4:
Let's continue explaining our idea of medical imaging as a form of inner vision. You know, as humans, our natural vision is powerful—but it's limited. We can see the world around us, we can observe surfaces, recognize faces, navigate through our environment. But our vision is tied to visible light, and that means we can't see through most objects. We don't know what's behind a wall, inside a box, or within our own bodies.
That's where medical imaging comes in. It gives us the ability to look beyond the surface—to see inside the human body without making a single incision. This is what makes medical imaging so extraordinary. It extends our natural vision, allowing us to explore what's hidden, to understand both the structure and function of our organs, and to help guide medical care.
Think about it—when someone doesn't feel well, when there's discomfort or concern, what do we do? We turn to imaging. It becomes the eyes of the physician, revealing what's happening beneath the surface. Through imaging, doctors can measure physiological function, detect disease, and ultimately provide treatment that can ease suffering or even save lives.
That's the true power of medical imaging—and why we call it the inner vision of modern medicine.

slide5:
Let me ask you this—what if we could see through the human body, almost like

having X-ray vision? Imagine being able to look at tissues, organs, even down to the cellular level. Sounds like science fiction, right?
Well, that’s exactly the dream of medical imaging. This field gives us a kind of "super vision"—an inner vision—that lets us explore the human body from the inside out, completely noninvasively.
And the truth is, what once seemed futuristic is now part of everyday medical practice. With the tools we’ll talk about in this course, clinicians can detect problems earlier, treat patients more precisely, and monitor progress more closely. Many of the advances in imaging came from the needs that once seemed impossible—until technology caught up. You’ll see how these breakthroughs happened, and why they matter.

slide6:
Now, let’s go back to where it all began. In 1895, a German physicist named Wilhelm Conrad Roentgen was experimenting with cathode rays when he discovered something totally unexpected—an invisible form of radiation that could pass through solid objects. He called them X-rays because he didn’t quite know what they were.
To test his discovery, he took the first X-ray image—a picture of his wife’s hand. You could clearly see the bones, and even her wedding ring. It was the first time humans saw inside the body without surgery. That moment changed medicine forever.
X-rays could depict dense structures like bones and metal with incredible clarity. This breakthrough was so significant that Roentgen received a Nobel Prize in 1901.
That single discovery launched the field of radiology—and over a century later, we’re still building on it.

slide7:
Fast forward to today, and X-ray imaging—specifically X-ray radiography—is one of the most common tools used in hospitals and clinics around the world. It’s fast, affordable, and incredibly useful.
Let’s say someone falls and hurts their arm. An X-ray can quickly show whether there's a fracture. Or in breast cancer screening, mammography—a specialized type of X-ray radiography—can reveal tiny calcifications or early tumors that might not be detectable by touch. Early detection like this can save lives.
But there’s a limitation with X-ray radiographs. They’re 2D projections, so everything in the body along the X-ray path gets stacked into one image. That makes it hard to see soft tissues clearly or to isolate specific structures.
So, the next logical step was to ask—what if we could look from multiple directions and reconstruct a full 3D view of what’s inside? That question led to the development of Computed Tomography, or CT, which we’ll get into next.

slide8:
Alright, so earlier we talked about how regular X-rays give us a flat, two-dimensional view of the body. But there’s a catch—everything in that path gets layered on top of each other. It’s like looking at shadows overlapped together.
Now imagine if we could take many X-ray radiographs around the body—virtually, of course—and figure out what’s happening on transverse sections. That’s exactly what computed tomography, or CT, allows us to do.
The word "tomography" comes from Greek—"tomos" means slice, and "graphy" means writing or drawing. So, we're basically drawing slices of the body. With CT, we collect X-ray projections from many angles, then use a computer to reconstruct a cross-sectional image.
The result? A crystal-clear view of internal anatomy, without overlapping tissues. Today’s CT scanners are incredibly precise—some can even resolve features smaller than a millimeter. Think of it like slicing a watermelon and looking inside but doing it without a knife—just scanning. It’s an amazing tool, and one that has completely transformed modern medicine.

slide9:
Now, if you want to get a fun and visual explanation of how X-rays work, there’s a great animated video I did for TED-Ed. It’s called “How X-rays See Through Your Skin.”

In just a few minutes, it walks you through how X-rays are produced, how they interact with the body, and how detected signals are turned into a cross-sectional image. It’s a great refresher, especially if you’re a visual learner. You’ll also see how this technology became the foundation for other types of imaging we’ll cover—like PET, SPECT, and MRI. The video’s optional, but I really recommend checking it out when you have a moment.

slide10:
So far, we've focused on X-rays—but medical imaging goes way beyond that. In fact, it covers almost the entire electromagnetic spectrum.
Depending on what part of the spectrum you use, you get different kinds of information. For example:
Gamma rays are used in nuclear imaging, like PET and SPECT
Radio waves are instrumental in MRI
Visible and infrared light are for optical imaging
And even ultrasound, though not part of the EM spectrum, gives us mechanical wave-based imaging, referred to as ultrasound imaging.
Each imaging modality has its own strengths and weaknesses. Some are great for structure, others for function. Some are good at seeing bones, others at detecting molecular processes like metabolism or gene expression.
In fact, no single technique gives us the whole picture. That’s why modern medicine often combines them—what we call multimodal imaging. When we link them together, we get a rather rich view of what’s going on inside the body.

slide11:
Let’s talk a bit more about what makes tomography special.
When you take a regular photo with your phone, you get a picture—that’s it. It’s a direct result of what the camera sees. But tomographic imaging is different. You don’t directly capture the image you want. Instead, you collect data from many different angles, and that data is a kind of indirect measurement—a projection of what's inside.
Now here’s the key part: to get a cross-sectional image, we need to do something called inversion. In other words, we take the measurements and reverse-engineer the internal structure that produced them.
This is an inverse problem. It’s different from the more familiar forward problem, where you know all the inputs, and a forward model, and calculate the outcome. Like in physics: you know the force and the mass, so you calculate how something moves.
In tomography, it’s the opposite. You see the outcome—the measurement—and you ask: what could have caused this? That’s a lot harder.
Here’s a fun analogy: think of boiling an egg. That’s a straightforward process. But what if someone gave you a boiled egg and said, “Unboil it”? That’s the inverse problem. Sounds impossible—but researchers have found ways to reverse that process under specific conditions. It’s challenging, but not hopeless. And that’s what makes inversion in imaging both difficult and exciting.

slide12:
Here’s another example that might feel a bit more familiar.
Let’s say you take a photo with your phone, but it turns out a little blurry—maybe the camera shook, or the focus was off, or the camera is cheap and of poor quality. What you get is not the true image, but a version that’s been blurred in several ways.
Now, how do we describe that blurring? In image processing, we model it using something called a point spread function—or PSF for short. This describes how a single point of light spreads out in the image.
The process that causes blurring is called convolution. It’s a mathematical operation that combines the original image with the PSF. Convolution is sort of like multiplication—but in the so-called Fourier space that we will explain later in this course. Indeed, in the Fourier domain, convolution becomes regular multiplication. The Fourier transform is often performed using a fast algorithm, FFT in short. That simplifies a lot of computations.
But now comes the hard part: what if all you have is the blurry image? Can you work backward and recover the original, sharp version?
That’s another inverse problem. And just like with tomography, it’s tricky.

You're trying to undo the effect of the blur, based on your knowledge of the system. It takes clever algorithms and good models—but it can be done. And that’s the kind of thinking we’ll develop throughout this course.

slide13:
Let’s take a step back for a moment and think about how mathematics builds over time.
When you were in school, you learned the basics: addition, subtraction, multiplication, division. These are the building blocks. But once you reach higher levels—like where we are now—those simple operations take on more advanced forms.
Take addition, for example. At this level, we’re not just adding a few numbers—we might be summing over an infinite number of values. That’s what integration is: it's a continuous version of addition, used when we add up contributions from every point in a region.
What about subtraction? In calculus, that becomes differentiation. Instead of just finding a difference between two numbers, we’re looking at how fast something is changing at an exact point by evaluating a very small difference of a smooth function between two rather close points.
Now let’s look at multiplication. In imaging, we move from regular multiplication to something called convolution—a fundamental concept in systems and signal processing. Convolution helps us model how a spatially-invariant linear system modifies a signal or an image. You’ll see this again and again in medical imaging. Again, a convolution in a spatial or time domain is a multiplication in the Fourier domain.
And finally, what about division? Here’s where things get interesting. The advanced version of division—where you’re trying to reverse a process—is what we call an inverse solution; the deblurring problem is a good example. You’re given a blurry image, and your job is to figure out what is the underlying ground truth image. And that’s essentially what we do in tomographic imaging.

slide14:
Now, here’s the point: inversion is hard. Much harder than the forward process. In a forward problem, you start with all the information, and you calculate the outcome. That’s usually pretty straightforward. But an inverse problem is the opposite: you’re starting with the outcome, and trying to infer what caused it. And this is where things get tricky. There might be multiple possible answers, or the data might be noisy or incomplete. You need to bring in clever mathematics, algorithms, physical models, and prior knowledge to make sense of it all.
That’s why this course isn’t just about memorizing facts. It’s about developing your quantitative thinking and really understanding what’s going on underneath the surface. It’s challenging—but if you put in the effort, it’s also incredibly rewarding.

slide15:
Now let’s take a moment to appreciate just how interdisciplinary medical imaging really is.
We’re combining mathematics, physics, chemistry, biology, and engineering—all in one field. Think about that. Whether we’re modeling how X-rays interact with tissue, analyzing ultrasound signals, or designing contrast agents for optical imaging—we’re pulling from all corners of science and technology.
This means that medical imaging isn’t just about one skill set. It’s about thinking broadly and connecting the dots between different fields. And if you enjoy solving complex problems, this is one of the best places to be.
So by now, you should have a pretty good sense of what this course is about—and what makes it so special. We’ll focus on modern biomedical imaging technologies, and we’ll build the knowledge and tools needed to understand how they work, and how they’re used to help people.
And honestly, it’s really cool stuff.

slide16:
So by now, you might be thinking—this is all very cool, but why do I really need to learn medical imaging?

Well, I'd say there are at least two solid reasons.
First—and perhaps most importantly—it gives you essential health-related knowledge. Whether you pursue a career in engineering, medicine, or something entirely different, understanding how imaging works can help you make sense of your own health and the health of those around you, during your hospital visits and afterwards.
The second reason is career-related. We'll get to that more. But for now, just keep in mind—this is cool knowledge for a niche group of people. It's personal and extremely useful in several ways; to say the least, it is a required course for your college education.

slide17:
Let me share something personal with you.
A couple of years ago, I had a kidney stone. Not fun. I ended up in the hospital, they injected a contrast agent, ran a CT scan, and just like that—they could see the size and location of the stone.
Now, because I understood what was going on, I wasn't worried or confused. I could follow the diagnosis, and I knew what the doctors were looking for.
And honestly, sooner or later, most of us will have a medical scan—whether it's an X-ray, a CT, or an MRI. And perhaps, all of them and multiple times. It could be you, or someone you love. When that happens, wouldn't it be helpful to understand what those images mean?
That's why this knowledge matters—not just in theory, but in real life.

slide18:
Looking ahead, many of you will have families of your own. When a family member gets a scan or a test result, wouldn't you want to understand what's happening? That's where this course really empowers you.
Even if you don't go into imaging research, even if you don't become a radiologist or an imaging scientist, just being able to interpret results and ask the right questions can make a big difference.
So that's the first reason: personal and practical value. Now let's move on to the second one—your career.

slide19:
Let's zoom out for a moment and look at biomedical engineering as a field.
It's one of the newest branches of engineering. Engineering started with disciplines like civil and mechanical, then moved into electrical, chemical, and aerospace. More recently, biomedical engineering emerged to bridge the gap between technology and medicine.
And it's grown fast—because it's relevant, it's interdisciplinary, and it's impactful.
In many universities, biomedical engineering includes at least two major focus areas. One is imaging, which is what we're covering here. The other is tissue engineering and regenerative medicine, often called TERM.
Whether or not you plan to specialize in imaging, this course plays a foundational role. It builds your technical knowledge, strengthens your systems thinking, and helps you understand one of the most important technologies in healthcare today.
So if you're here as part of your degree requirements—great. But don't just treat it like a checkbox. Learn it well. It will pay off, academically, professionally, and personally.

slide20:
Now, for those of you considering a concentration in imaging—or even just taking this course seriously—it can really pay off in your career.
Medical imaging is a technical field, but it's also full of opportunities. If you understand how imaging systems work, and you can apply that knowledge to real-world problems, you'll bring serious value to any team—whether in research, healthcare, or industry.
And in case you're still wondering: yes, biomedical engineering is a great profession. It sits at the intersection of life sciences, medicine, and engineering—and that makes it both exciting and employable.

slide21:
In fact, if you look at trends in job growth, biomedical engineering consistently ranks among the fastest-growing careers. Just take a glance at some recent employment reports—you'll see biomedical engineers right near the top.
Why is that? Well, part of the reason is that healthcare is becoming more tech-driven. Devices, sensors, data analysis, machine learning—it all depends on people who can bridge the gap between biology and engineering.
So not only is this field intellectually exciting—it also comes with strong job prospects and competitive salaries.

slide22:
Now let's zoom out to the global scale.
Medical imaging isn't just a field—it's a major global industry. And it's growing fast. Countries around the world are investing heavily in imaging infrastructure, service, and innovation.
If you look at the data, regions like Asia and Europe are rapidly expanding their market share. Asia, in particular, has seen a surge in R&D spending and adoption of advanced imaging technologies. That means more jobs, more research funding, and more innovation coming from across the globe.
So, whether you're aiming for industry, academia, or entrepreneurship—this is a great space to be in.

slide23:
When it comes to imaging modalities, CT and MRI continue to lead in market share and clinical use.
That's no surprise—both are incredibly versatile and widely used across medical specialties. CT scans are fast and great for detailed anatomical views. MRI provides outstanding soft tissue contrast, which is essential in areas like cardiology and neurology.
You might be surprised, though, that nuclear imaging—like PET and SPECT—is also a major player. These techniques are critical for functional imaging and cancer imaging. They are becoming more popular as we move toward personalized and molecular medicine.
Each year, investment in these technologies grows. The message is clear: the field is fast evolving, and there's a strong momentum behind it.

slide24:
So now that we've seen how vital CT, MRI, and other imaging technologies are in practice and in the global market—here's something interesting to reflect on.
A few years ago, a survey asked physicians and scientists to name the most important technological innovation of the entire 20th century. Think about that—this included everything: the internet, smartphones, even space travel.
And what came out on top? Magnetic Resonance Imaging (MRI) and Computed Tomography (CT). Not even close. These tomographic imaging technologies were voted number one by a considerable margin.
Why? Because they fundamentally changed how we diagnose and treat diseases. They let us see inside the human body in ways that were never possible before—accurately, noninvasively, and in real time.
As you go deeper into this course, you'll learn how these systems work—so not only will you understand their value, you'll also be able to explain it confidently in, say, a job interview or a professional conversation. That's powerful knowledge.

slide25:
Speaking of innovation—let's talk about one of the major players in the field: General Electric, or GE, in the neighborhood of Rensselaer Polytechnic Institute.
This is a company with a deep history of technological contributions, not just in medical imaging, but across many industries. And interestingly, "GE" also happens to be my initials—but theirs is with a capital E. So, no confusion!
GE was founded back in 1892, not long after X-rays were discovered. And Rensselaer Polytechnic Institute was founded even earlier. So these two organizations have grown side by side, both focused on pushing the boundaries of science and technology.

GE has been part of some most transformative developments in imaging—from early X-ray machines to advanced MR scanners. They've really helped shape this field.

slide26:
Now, within GE, one of the most important innovation engines is their Global Research Center. This is where long-term, high-impact R&D happens.
Their record speaks for itself—two Nobel Prizes, major contributions to physics, chemistry, materials, energy, and of course, medical imaging.
In 1984, they played a critical role in the development of MRI systems. And in 1999, digital X-ray systems came out of their labs. They've consistently been at the forefront of translating ideas into real-world technologies. My lab and Global Research Center has been closely collaborating on medical imaging over decades.
So when we say medical imaging sits at the cutting edge of innovation—we mean it.

slide27:
And what makes GE even more impressive is the global reach of their research. They have major centers across the U.S., Europe, Asia, and South America. For example, the Global Research Center in Niskayuna, New York—right near RPI—is home to over a hundred labs and nearly two thousand technologists.
For a student interested in imaging, that's a golden opportunity. You can collaborate, intern, or even launch a career through the GE-RPI partnership. Some of you may want to go into academia, others into entrepreneurship or industry. Wherever you go, connecting with innovation hubs like GE gives you a head start. Especially in imaging, they are a top-tier industrial partner.

slide28:
Before we wrap up this introduction, I'll also give you a simple task—just a brief introductory slide you can prepare. It's not graded, but it helps me get a sense of who you are, what your background is, your GPA, research experience, and what you hope to learn. That way, I can adjust the flow of this course to better support your journey.

slide29:
As we continue our journey into medical imaging, let me take a moment to introduce myself in a little more depth.
I'm Professor Ge Wang, and in addition to teaching this course, I'm deeply involved in research on medical imaging. My work spans both the theoretical and applied sides of the field. I've dedicated much of my career to advancing medical imaging science so that we can better diagnose and treat disease.
If you're interested in exploring more about my work, feel free to check out my home page and lab's website, where we share updates on our latest projects.

slide30:
Now let me share a bit more about the imaging program at RPI.
My main interests are in X-ray imaging and optical imaging. I direct our Biomedical Imaging Center, where we work on advancing these technologies. One exciting direction is how artificial intelligence can help improve imaging—making it faster, more accurate, and more insightful.
This work doesn't happen in isolation. I collaborate with outstanding colleagues like Professor Xavier Intes, a leader in optical molecular tomography, Professor Pingkun Yan, a leader in AI-based image analysis and healthcare analytics, and many others. Together, we team up with world-class researchers from GE, Yale, Wake Forest, MGH, Stanford, FDA, and other top institutions. Our shared goal is to create imaging solutions that are innovative, influential, and that truly make differences in patient care.

slide31:
These collaborations all come together in a fantastic building—the Center for Biotechnology and Interdisciplinary Studies, or CBIS.
CBIS is more than just a building. It's a hub where experts from biology, engineering, and computing come together to tackle big problems in science and healthcare. It's where our imaging center is rooted, and where much of the work

that I’ve been describing happens.
Over the past decade, CBIS has been at the forefront of interdisciplinary research, and it continues to grow in both scope and impact.

slide32:
Now, I often like to say: if CBIS were a coordinate system, my office would be right at the origin!
Of course, for this online course, we’re not meeting in person. But the idea is the same—I’m here as a resource for you. The lectures, materials, and any additional resources I share are all designed to help you succeed and make the most of this learning experience.

slide33:
We’ve designed this course with two major phases in mind: a foundation phase and a modality-specific phase. In the foundation phase—highlighted in light blue—you’ll build the critical background knowledge necessary to understand modern biomedical imaging systems. This includes essential concepts in linear systems, convolution, and Fourier analysis. These are the mathematical and conceptual tools at the heart of almost every imaging modality we’ll cover later.
Why this structure? Well, my previous experience with diverse students tells me that not everyone enters this course with the same preparation. Some of you may have strong backgrounds in physics, applied math, and electrical engineering—and that’s great. You might already be familiar with system theory and Fourier transform. But many others come from biology, chemistry, or material science backgrounds, where these topics are not emphasized. That’s completely okay. Medical imaging is interdisciplinary—and complex. So my job is to guide all of you, regardless of background, to the point where you understand these core ideas clearly and intuitively. It won't be enough to just listen and watch these lectures. You’ll need to invest time in previewing and reviewing the materials and doing the assignments. But rest assured—your effort will pay off. These foundational topics will resurface again and again, especially in modalities like CT and MRI, where the Fourier transform is a key player.
Also, linear systems might look simple on the surface—after all, they obey additivity and scaling rules—but they come with nuances. One particularly subtle concept is convolution, which plays a crucial role in both systems analysis and image reconstruction. We’ll make sure you understand not just how to compute it, but also why it matters and how it shows up in real-world scanners.
So, we begin with the foundation. Later, we’ll shift to the modality phase, shown in dark green—where we delve into imaging techniques like X-ray CT, nuclear imaging, MRI, ultrasound, and optical imaging. Each one has unique physics, engineering principles, and clinical applications.
Think of it this way: the foundation is your universal imaging 'language'. Once you master it, learning each new modality becomes much easier, and far more meaningful.
Let me walk you through our tentative timeline for the course—though as with any plan, we may make some adjustments along the way.
This lecture kicks off the course. In our next session, you’ll be introduced to MATLAB. If you're unfamiliar with it, don’t worry—our teaching assistant will help you get started. MATLAB is a vital tool in imaging research, especially when modeling systems or reconstructing images from tomographic data. After the introduction, you’ll have a short MATLAB-related assignment to sharpen your skills.
Then we move into the foundation module—the light blue section. This is where we cover the theoretical concepts that apply across imaging technologies. After that, we transition to the green section—where we examine individual imaging modalities.
You’ll also be evaluated through three closed-book exams spaced across the course: one after the foundation module, one after the first set of modalities (CT and nuclear imaging), and one to cover everything else. Your final grade will be based on a combination of these exams, your homework, and participation.
Now, this is an online course, so we’ve adapted all the content accordingly. Pre-recorded lectures, interactive demos, and exercises will all be made available. If you have questions, we encourage you to reach out via the discussion forum or email. Quick responses are part of our commitment to help

you succeed.
A quick tip: spend quality time understanding the Fourier transform early on. It's a recurring theme. And if you've ever wondered why exponential functions like
$e^{i\theta}$ show up everywhere, this course will make that connection crystal clear. We'll discuss how they lead to elegant mathematical representations and how these expressions play into image formation.
Finally, one last point. This course is not just about passing exams—it's about equipping you with a deep understanding that you can carry forward into research, engineering design, and healthcare work. Let's work together and learn from each other.

slide34:
For the modality-specific portion of the course—the dark green section—we'll be using the green textbook: Introduction to Biomedical Imaging by Andrew Webb. It's available online or through our university bookstore.
This book isn't very thick, but it's well-structured and written with clarity. Even though it was published a while ago, the fundamental principles it covers haven't changed. Physics is timeless in that way. We'll supplement the book with more recent findings, particularly in areas like deep learning for image reconstruction, low-dose imaging, and multi-modal fusion—advances that have significantly shaped the field over past several years.
The textbook will guide your learning, especially in the second part of the course. For the foundation topics like linear systems and Fourier theory, we'll use tailored lecture notes and curated readings, because unfortunately, there's no single book that explains everything in a concise and consistent manner for students from all backgrounds.
So, keep this in mind: while you may rely on the textbook for the modality section, the theoretical part of the course will require close attention to lecture materials and exercises. That's where most students need the extra help—and that's why we're here to support you.

slide35:
In addition to our course materials, I'd like to mention a particularly helpful supplementary resource that aligns well with the foundation we're covering—especially Fourier transforms and linear systems.
Stanford University offers an excellent course titled 'EE261: The Fourier Transform and Its Applications'. It's taught by Professor Brad Osgood, and the best part is that it is freely available online through Stanford's SEE platform. You can access the full textbook there and watch high-quality lecture videos.
Now, a quick word of caution: the Stanford course spans an entire semester. The book is over 400 pages long. Don't feel overwhelmed. You absolutely do not need to cover all of it. Think of it as a reference companion to our course. In fact, for our purposes, reading even 80–100 pages from selected sections will be more than enough to reinforce your understanding.
What makes this resource valuable is the clarity and consistency of its explanations and mathematical treatment. It builds the theory from the ground up, with elegant examples and a style that's both rigorous and readable. If you ever find yourself puzzled by a concept I introduce—say, the derivation of the Fourier series, or the intuition behind convolution—this book can offer a second viewpoint. Sometimes, a slightly different presentation is all you need for clarity.
However, please note that we will follow our own structured path in this course. For example, I introduce linear systems before Fourier transforms, because in imaging, systems thinking provides the conceptual grounding. The Stanford course takes the reverse order. Both are valid—but here, we start with linear systems to better prepare for their applications in imaging modalities like CT and MRI.
Please use this Stanford material as a valuable supplement, not a replacement. You'll benefit most by following our course flow carefully and referring to the Stanford book only when you want to go deeper or review challenging topics.
Lastly, although I have reached out for formal permission, the material seems publicly accessible and openly shared for educational use. So, feel free to explore it from your own dorm room, your library, or anywhere you study.
This is an exciting time to learn these topics—tools like ChatGPT make it easier

than ever to gain a deeper understanding. Let’s use them wisely.

slide36:
In this course, you’ll have access to a complete set of digital resources to support your learning.
We’ve included professionally recorded video lectures, so you can review key concepts anytime. All PowerPoint slides used in the course are also available for download and review.
For the modality sections, the textbook is available as hardcopy and PDF, making it easy to follow along even to search specific topics digitally. The course schedule follows a well-structured plan from previous offerings, with minor updates to reflect some latest progress in imaging science.
These materials are here to help you succeed—use them actively as we explore the foundation materials and modalities of medical imaging.

slide37:
Continuing from our previous discussions, let’s dive into the physics of X-ray imaging and how we can collect and process data to reconstruct an image.
In X-ray imaging, we start with the concept of a line integral. Essentially, the X-rays pass through the object, and what we measure is the energy of penetrated x-ray photons and then compute the line integral along the x-ray path. Such line integrals together form the Radon transform. To make it simpler, imagine you have an object, and you put all line integrals from a specific direction to form an x-ray projection. This projection is a one-dimensional profile, which is the information you can get along all the X-ray paths passing through the object in that direction.
Now, let’s take this idea into the realm of mathematics. We start with a 2D image, denoted as
$f(x,y)$. When we have all x-ray projections or the Radon transform, we get a new 2D function $p(\theta,t)$, where t is the coordinate along the 1D detector array and θ is the projection orientation. By changing the angle θ, we rotate the data acquisition system to gather x-ray projections.
Now, the magic happens when we flip the process around. Given this collected projection dataset, our goal is to reconstruct the original 2D image $f(x,y)$. This is the inverse process, and here’s where things get interesting: by perform the 1D Fourier trans of each and every of the 1D projections, we move into the Fourier domain.
To put it simply, the 1D projections we collect represent profiles in the frequency domain. By rotating the data acquisition system, we eventually capture the entire Fourier space. And once we have all the data in the Fourier space, we can apply the inverse Fourier transform to reconstruct the original image. This process connects the data from the physical object to the Fourier spectrum of the underlying image, and from the Fourier spectrum, we can reconstruct the original image.
While this overview gives you the fundamental idea, later in the course, we’ll dive deeper into the Fourier transform, why they can be inverted, and how this mathematical framework ties directly to the imaging process. But for now, keep in mind that the forward process (how we collect data) and the inverse process (how we reconstruct the image) are two sides of the same coin, closely linked by Fourier analysis.

slide38:
Continuing with our exploration of imaging modalities, let’s now discuss Positron Emission Tomography, or PET. While the previous imaging techniques focused on capturing the anatomical structure, PET offers a powerful way to visualize the functional activity in the human body.
So, how does it work? Well, the key to PET is the introduction of radioactive tracers into the body. These tracers are typically radioactive chemicals that participate in your body’s metabolic processes. For example, tumors tend to consume a lot of glucose. If we couple glucose molecules with a radioactive tracer, we can track how much glucose the tumor is using, which directly correlates with its activity. Tumors, which are often highly metabolic, will naturally absorb more of this tracer.
Once the PET tracer is in the body, it begins to emit gamma ray photons in

pairs, and two photons in a pair will travel in opposite directions. Gamma rays, much like X-rays, can penetrate through tissues and be detected externally. When we receive gamma ray photons from two different detector locations, we know they originated from the same event somewhere on the line connecting the two detectors.
This process is similar to the measurement of line integrals used for CT. Then, we can apply Fourier analysis to reconstruct the images. In PET, Fourier transforms help us create detailed, functional images of the body’s processes. The goal here is to understand how much activities are happening at specific locations. So, while we’ve seen how CT gives us anatomical details, PET allows us to observe the biological functions—like metabolism—within those structures. This modality is especially useful for identifying cancerous tissues, studying brain function, and even monitoring heart conditions. Later, we’ll explain how the gamma ray photons are detected and processed to reconstruct these functional images. But for now, remember that PET imaging provides us with a way to look inside the body not just for its structure, but also for how physiology is actively working.

slide39:
Now, let’s explore a closely related imaging technique in nuclear medicine called Single Photon Emission Computed Tomography, or SPECT. While Positron Emission Tomography (PET) relies on the detection of paired gamma ray photons emitted from the tracer, SPECT works differently.
In SPECT, the radio tracer emits gamma ray photons randomly and individually, without being paired. Essentially, instead of detecting two photons coming from opposite directions (as we do in PET), we measure single gamma ray photons emitted from a radioactive tracer within the body.
When a photon is detected through a collimator, we know it came from a particular path through the patient body, but since the photons are not paired, we need to use different data preprocessing and image reconstruction methods. SPECT provides valuable information, but it’s a bit different from PET. In SPECT, we capture these single photons from various angles around the body and use dedicated tomographic reconstruction techniques to create an image.
We’ll discuss the details of SPECT imaging later in the course, but for now, keep in mind that while PET relies on paired photons to provide highly precise spatial information, SPECT uses single photons to gather functional information in a different way. Both PET and SPECT produce tomographic images of radio-tracer activities.

slide40:
Now, let’s turn our attention to the third major imaging modality: Magnetic Resonance Imaging, or MRI. MRI uses a completely different approach compared to X-ray and nuclear imaging techniques including PET and SPECT. In MRI, we utilize radio frequency signals to gather detailed images of the body.
Here’s how it works: First, we subject the patient to a very strong magnetic field. This magnetic field aligns the hydrogen nuclei—which are abundant in water molecules in your body. Then, we introduce a radio frequency (RF) signal that briefly disrupts this alignment, causing the hydrogen nuclei to emit their own radio frequency signals as they return to their original stable state.
The key advantage of MRI is that it is sensitive to the water content in different tissues and interactions of hydrogen nuclei with their micro-environments, which makes MRI extremely useful for imaging soft tissues like the brain, muscles, and organs. The signal we collect can give us a lot of valuable information, such as chemical shifts in tissues, which help us distinguish between different types of tissues and even detect abnormalities like tumors or inflammation.
MRI is also great for studying brain activity, as the technique can be used to observe changes in blood flow over time. This makes MRI not only a structural imaging tool but also a functional one, adding a layer of depth to our understanding of the body.
Together with CT and nuclear imaging techniques, MRI forms the trio of most important imaging modalities. While CT focuses on capturing detailed anatomical structures using X-rays, and nuclear imaging provides insight into physiological functions, MRI offers unmatched soft tissue contrast and the ability to examine

both structure and function in a non-invasive way, even into our thinking processes.

slide41:
Now, let's look at Ultrasound Imaging. While it might not be as widely discussed or as advanced as the big three—CT, nuclear imaging, and MRI—ultrasound remains a powerful, affordable, and safe imaging tool in medical practice, and seems playing an increasingly important role for point of care imaging.
So, how does ultrasound work? It operates based on the principle of mechanical vibrations. In ultrasound imaging, we use a device called a transducer, which contains piezoelectric materials. These materials can generate high-frequency sound waves when subjected to an electric current. These sound waves are then sent into the body.
As the sound waves travel through tissues, they interact with various structures inside the body. Some waves are reflected, and others continue through the tissues. The reflected waves are detected by the transducer, and from this data, we can reconstruct an image of the internal structures. Essentially, we're mapping the body's anatomy based on how sound waves bounce off tissues.
One of the greatest benefits of ultrasound is that it's non-invasive and completely safe—unlike X-rays or CT scans, ultrasound doesn't involve ionizing radiation. It's also incredibly cost-effective and high-speed, which makes it a go-to tool for many real-time imaging tasks. You'll often see it used during fetal imaging in pregnancy, as well as in blood flow monitoring and guiding biopsies.
While it may not provide the same level of detail as MRI or CT, ultrasound's affordability, speed, and safety make it a valuable tool for many clinical applications. Coupled with AI methods, ultrasound imaging has yet more to offer.

slide42:
Now, let's talk about Optical Imaging, another fascinating imaging modality that uses visible and infrared light. Unlike the other imaging techniques we've discussed, optical imaging focuses on how light interacts with cells, allowing us to reveal molecular and cellular interactions in the body.
One unique aspect of optical imaging is its ability to extract signatures from these biological interactions, providing valuable information about what's happening at a microscopic level. This is where things get exciting! For example, we can use luminescence probes to tag specific proteins or gene expressions. Once tagged, these proteins will emit a luminescent light and make them visible to our imaging system.
Think of it like the glow of a firefly in a summer night. When certain proteins or cells are tagged, they emit bioluminescent light, which we can detect. This is a form of passive imaging, meaning the light we need to detect is emitted by the body itself, rather than steered by any external light source.
For example, in a typical experiment, an animal might be placed in a dark room, and we would observe the light emissions coming from specific areas on its body. We want to perform tomographic reconstruction to find the internal distribution of the bioluminescent sources. For that purpose, we combine an optical imager with a separate tomographic scanner, which helps us get both the anatomy and the optical properties of organs and structures to build a forward imaging model. After that, we can invert this forward model for reconstruction of a 3D distribution of bioluminescent sources.
The real challenge in optical imaging isn't in the forward process—where we predict the external signal based on the known internal structures. The difficulty is in the inverse process, where we have this inverse problem: Given the external light signals we measure, how do we reconstruct the distribution of bioluminescent sources inside the body?
To do this, we use multiple mirrors to collect all the emitted light, creating a complete set of external views. We then use this data to reconstruct the internal light source distribution. This method has been known as bioluminescence tomography that we pioneered and has become quite popular in the literature.
Note that optical imaging includes a good number of techniques. Bioluminescence tomography is only one of them. We have fluorescence tomography as well, with fluorescence probes excitable by an external laser source. More clinically

important is optical coherence tomography or OCT, which is used to check your retina in eye clinic.
In summary, optical imaging, like other modalities, allows us to visualize and understand structures, probes or tracers inside the body. It’s a unique and powerful way to study biology from the inside out.

slide43:
Now that we’ve explored the different imaging modalities and how they allow us to visualize features inside the body, the next logical step is to talk about what we can do with the images we’ve acquired. Once you have an image, you don’t just look at it—you analyze it to gain task-specific information. This is where the domain of image processing grows.
In medicine, image analysis is crucial for extracting actionable insights. There are various techniques we use, such as image segmentation, where we identify specific regions or structures within the image—like tumors, organs, or blood vessels. By doing this, we can focus on areas of interest, making it easier to understand what’s happening inside the body.
Another important tool is image classification, which allows us to categorize different structures or tissues based on their visual features. This helps doctors make diagnosis, track changes over time, or even guide interventional procedures.
We also use techniques like image enhancement, which improves the quality of the image by making certain features more visible, such as image deblurring, which sharpens images that might have been blurry due to motion or technical issues.
Some textbooks on medical imaging even include image visualization. This is where new technologies like augmented reality (AR) become relevant. I recently read an article, where the use of AR is discussed for medical imaging-guided surgery. With this innovative method, surgeons wear specialized glasses that overlay imaging data directly onto the patient’s body during surgery, guiding their actions in real time. While this is an exciting application of medical imaging, for the scope of this course, we’ll focus more on bio-imaging and bio-instrumentation, with an emphasis on tomographic imaging.
In this course, we’ll briefly cover some key aspects of image analysis, but we won’t delve too deeply into every aspect of processing. By the end, you’ll have a solid understanding of how image analysis ties into the broader field of medical imaging.

slide44:
To put everything in perspective, let’s take a step back and look at the big picture. In understanding the human body, we often talk about two key concepts: phenotype and genotype.
Genotype refers to your genetic makeup—the DNA that gives rise to your traits. Meanwhile, phenotype refers to the observable characteristics of an individual, such as physical traits, biological functions, and disease states. Biomedical imaging plays a crucial role in studying the phenotype. Through imaging, we can visualize and analyze the biological structures and functions that define an individual’s phenotype.
Now, bioinformatics and genetic profiling give us a lot of information about the genotype—things like your genetic predispositions and variations. However, these genetic insights need to be connected to the physical manifestation of the genes—the phenotype. We can use biomedical imaging to observe and understand the effects of genetic factors on the body.
While we’re entering the world of imaging in this course, it’s important to recognize that we’re focusing primarily on phenotype—how imaging allows us to visualize and analyze the body's structures and functions. However, linking genotype (your genetic information) with phenotype (the observable traits and conditions) is an essential goal for personalized medicine. This connection between the two will allow for more accurate diagnosis, more effective treatment, and unprecedented prediction in the future.

slide45:
Looking ahead, the next lecture will focus on building a solid foundation in MATLAB, which will be essential as we move into topics like Fourier transforms and linear systems. Rather than just listening my lecture passively, you will be

using MATLAB directly, as this hands-on approach will help you understand key concepts like convolution and the Fourier transform in a deeper, more practical way.
MATLAB is an invaluable tool used across many disciplines; for example, medical imaging and biomedical engineering in general. It will be an important part of your learning journey. If you don't have it installed yet, you can easily download it from the official MATLAB website—this is freely available to you. If you run into any installation issues, support is readily available through your university’s resources, such as the IT support team.
To get started, I recommend completing the MATLAB On-ramp tutorial, which is a two-hour interactive course that will walk you through the basics. I've personally found this tutorial very effective in building MATLAB skills, and it’s designed for learners of all levels. Rather than only spending time in the next lecture on MatLab, I encourage you to complete this on your own time before the next lecture. This will give you the background you need and prepare you for the upcoming lessons, where we’ll apply these concepts together.
Let me make it your homework: going through the MATLAB On-ramp. By the end of that session, you should be comfortable with MATLAB’s interface and basic functions. Then, you’ll be ready to learn more.

slide46:
You’ll notice that this course is not just about the theory—it’s also about putting those theories into practice. For example, you can use MATLAB to simulate various types of data processing. It’s not about learning complicated tricks; it’s about using basic MATLAB functions to analyze data, visualize results, and perform tasks that will help with your assignments.
Remember, the goal here is not to get bogged down in fancy techniques. Rather, we want to ensure that you have a solid foundation in MATLAB so that when it’s time to do your homework or apply concepts like the Fourier transform or linear systems, you’ll have the tool you need to execute them effectively. Now, ChatGPT is a great tool to help you develop MatLab codes. Please keep this in mind and become familiar with ChatGPT in this context.
With hands-on practice, you'll get familiar with the interface and the functionality, making your learning process or workflow more effective and more efficient. It’s all about building that confidence through practical experience.
So, let’s keep learning with MATLAB. By engaging with it now, you’ll set yourself up for success in the course.

slide47:
Now, I want to say something that’s transforming the world of medical imaging: the huge wave of machine learning, artificial intelligence (AI), and robotics. These technologies hold huge potential, and they are revolutionizing how we analyze and interpret medical images, to say the least.

slide48:
As we continue to explore the exciting advancements in medical imaging, it's clear that this field is still young. Through a recent survey, I found that medical imaging is one of the hottest areas in research today, with a significant increase in the number of published papers every year. However, what’s really catching attention is machine learning—it's quickly becoming an even hotter topic than medical imaging itself.
What does this mean? The combination of medical imaging with machine learning, especially deep learning, holds enormous potential for the future of diagnostics, treatment planning, and personalized medicine. This intersection is where we can expect to see transformative changes in how we approach medical imaging. The ability to reconstruct, analyze and interpret medical images with AI-empowered algorithms opens doors to endless possibilities.
In 2016 I wrote the first perspective article titled Perspective on Deep Imaging, where I share my thoughts on how deep learning and AI are poised to revolutionize medical imaging in the coming years. If you’re curious about my views on this, feel free to search for it and dive deeper into the future of medical imaging.
However, in this course, we’re going to stick with the traditional classroom teaching, but for those of you who are highly motivated and eager to push beyond

the syllabus, there’s an exciting option. You can choose to work on a project that integrates machine learning with medical imaging. This would involve hands-on research, where your work could contribute directly to a meaningful outcome in the field.
While it’s not a requirement, this is an opportunity for those of you who want to challenge yourselves and explore advanced concepts further. For students who choose this option, 30% of your grade could come from the research project, with 70% from other course elements. Many students find that they don’t have the time to take such a project, but if you’re passionate about it, the opportunity is there. In fact, some of my former students, even high school students, have worked with me on research projects, and have gone on to get their work published in conferences and journals, which can be a huge boost to their career.
So, if you feel like the course is a bit easy and you’re ready to go beyond, this project option could be a perfect fit. The office door is always open, and if you choose this path, you'll get real-world experience and potentially contribute to cutting-edge research.

slide49:
Building on what we have discussed so far about the intersection of medical imaging and machine learning, let’s take a moment to look further ahead.
I truly believe the future of medical imaging—and medicine in general—is full of incredible possibilities. As artificial intelligence continues to advance, we’re already seeing systems that can match or even surpass human performance in interpreting medical images. It’s very possible that, in the not-so-distant future, many of the tasks traditionally done by radiologists will be basically handled by AI.
And it goes beyond that. Imagine robotic surgeons, guided by multimodal imaging and machine learning, performing procedures with precision that’s hard for even the best human surgeon to match. Over time, we might see AI and robotics take on more and more roles in healthcare, not to replace professionals, but to enhance what we can achieve—and to make healthcare safer, faster, and more accessible.
This is the direction the field is moving in, and it’s one of the reasons why understanding imaging technologies today is so valuable and timely. You're not just learning how things work now—you’re preparing to shape the future.

slide50:
As we think about how artificial intelligence and robotics might transform medical imaging and healthcare, it’s hard not to be reminded of how often science fiction points the way to future innovations.
For example, have you seen the movie Elysium? In that film, there’s this futuristic machine—a kind of all-in-one medical scanner. It can detect and treat any disease, regenerate or even replace organs, and remarkably, reverse the aging process. It’s the ultimate vision of technology and medicine coming together to provide perfect healthcare.
Of course, today we aren’t there yet. But the ideas behind that machine—integrated imaging, precise diagnostics, personalized treatment—are exactly what we’re working toward with advanced medical imaging, machine learning, and robotics. It’s a powerful reminder of what could be possible.
I hope this inspires you as we move forward in the course. By learning the foundations of medical imaging today, you’re taking the first step toward contributing to the invention of our future.