slide1:
Welcome to this lecture on medical imaging. In our earlier sessions, we introduced some general concepts in medical imaging and provided a brief overview of essential tools, including MATLAB programming. Today, we begin what I consider the foundation section of this course. This foundation will help you build the knowledge needed to understand the major medical imaging modalities we'll explore throughout the semester.

slide2:
We are right on schedule in our journey through this material.
If you've had a chance to look over the reading materials for today's lecture, that's great — it will make it easier to connect the ideas we cover. If not, that's fine too. I encourage you to follow along closely and take time afterward to review the main points. Consistently reinforcing concepts as you go will help you develop a stronger and more intuitive grasp, especially as the material becomes more mathematically detailed

slide3:
Today's lecture is straightforward. We're going to talk about the concept of a system—a core idea that underpins everything we'll cover in this course. If you've downloaded the draft of my textbook, you'll see that this topic forms the first chapter. Reading just the first ten pages will give you a solid grasp of the key ideas.
In each chapter, I aim for a balanced structure: we start with three main sections, followed by a concluding section with remarks and points for deeper thought. In this lecture, we'll begin with general systems, then narrow our focus to linear systems—which, as the name suggests, behave in a straightforward, predictable way, much like linear functions in mathematics.
To qualify as a linear system, two properties must be satisfied: additivity and homogeneity. Together, these form what we call the superposition principle.
We'll break these ideas down, explore when one property implies the other, and consider whether they are truly independent requirements. Lastly, we'll touch on nonlinear systems—which are simply systems that don't meet these linearity conditions. This will give you a broad overview of the system concepts essential for medical imaging.

slide4:
Let's build on that by comparing systems to something you're already familiar with—functions. A function is a mathematical rule that maps each input from one set, called the domain, to an output in another set, called the range. Systems work in a similar way. They take an input, process it through interconnected elements, and produce an output.
These elements could represent anything: parts of a social system, components of a physiological system, or pieces of a mechanical system. The important point is that both functions and systems describe relationships between inputs and outputs. Our goal is to define these relationships clearly so that we can apply them to engineering and medical imaging problems.

slide5:
Here's a simple illustration of what a function does. Imagine selecting a value from the domain—say, a number—and applying the function to find its corresponding value in the range. This basic idea of mapping inputs to outputs might seem simple, but it's fundamental to everything we're discussing.
In my slides, I use symbols to highlight key points: a red diamond marks something essential—concepts you should remember clearly. A green circle points out interesting but optional material, things that go beyond what you strictly need for engineering purposes. While we don't aim to become mathematicians in this course, understanding these key ideas deeply will strengthen your grasp of medical imaging systems.

slide6:
Functions come in many forms—like sinusoidal functions, linear functions, and even more complex structures such as those defined recursively in fractal geometry. If you're unfamiliar with fractals, they're fascinating patterns that

repeat at every scale, and I encourage you to explore them further if you're curious. But for our purposes, it's enough to be comfortable with regular functions and their properties.

slide7:
Another powerful tool in mathematics is the Taylor expansion. This allows us to approximate any smooth, well-behaved function using a sum of terms: starting with a constant (the function's value at a point), then adding linear terms based on the first derivative, and quadratic terms based on the second derivative, and so on.
Each added term improves the approximation, as long as certain conditions are met to ensure the series converges. In engineering practice, though, we often stop at the simpler terms—constant or linear—because they offer a good balance of accuracy and simplicity. This idea of building approximations is central to how we model systems efficiently.

slide8:
In fact, for most practical purposes, we're satisfied with linear functions. In one dimension, this means a straight line; in two dimensions, it's a plane; and in higher dimensions, we extend the idea further. These simple, linear relationships are easy to work with and provide the foundation for much of engineering analysis—including in medical imaging.

slide9:
As engineers, we often think of the world in terms of systems. We use mathematical operators—like L or O—to describe how a system transforms an input into an output. If you feed an input signal, V, into a system, the output, W, represents how that signal is modified.
These inputs and outputs can take many forms: time-dependent signals, images, tensors, or even discrete data sets like color values in an image. The concept is general: a system processes an input and produces an output, just as a function maps domain to range.
In medical imaging, we apply this thinking to devices like CT scanners, where the entire system transforms physical signals into diagnostic images. And as we analyze these systems, we rely on mathematical relationships to describe their behavior in precise terms. This distinction between thinking as engineers about systems and as mathematicians about abstract functions is subtle but important as we move forward.

slide10:
As engineers, much of our work revolves around systems, and one of the first things we need to do is measurement. Without measurement, we can't really understand or control anything. In fact, measurement is at the heart of both engineering and physics. For example, in quantum mechanics, measurement plays a critical role—until we measure, we don't truly know a system's state. It's like the famous thought experiment with Schrödinger's cat: is the cat alive or dead? Only measurement collapses the uncertainty into reality.
When we perform measurements, we typically have input variables and sometimes a reference. Consider weighing an object: we place a known weight on one side of a balance and the object on the other. If the balance is level, we know the object's weight equals the standard. Similarly, in electrical measurements, we might compare an unknown resistance or voltage against a reference using an ohmmeter or a multimeter. These devices let us measure things like current, resistance, and capacitance with precision.
More broadly, think of an instructor as a system: when the calendar says it's time for class, that's the input. The output is the lecture. But if something changes—say, the instructor is unwell or called away—the output changes too. This is a reminder that systems can have modifiers that influence their behavior.

slide11:
An imaging system is a specialized type of measurement system. Consider an MRI scanner: its purpose is to measure and produce an image of a cross-section of the body's anatomy. Ideally, this image is a faithful representation—a sharp,

clear point corresponds exactly to a sharp point in the body.
But not all systems are created equal. A high-quality system will produce sharp, high-contrast images, while a less precise system might blur details, much like how a low-resolution camera produces softer, less defined pictures. This blurring is a key characteristic of the system and tells us a lot about its performance. In essence, imaging systems, sensing systems, and measurement systems all belong to the same family—they process inputs to produce meaningful outputs through measurement.

slide12:
Another type of system is the control system. Unlike simple measurement systems, control systems don't just generate an output—they use that output as feedback to adjust the input or the system's behavior.
For instance, think of a radar tracking an airplane. The radar detects the plane's position and adjusts itself to keep the target within view. This feedback loop helps the system achieve its goal, even as conditions change. Control systems are essential in engineering whenever we want systems to adapt or correct themselves in real time.

slide13:
Today, robotic systems represent a hot area of control system development. From robots that provide assistance in hotels or nursing homes to those that help care for infants, these systems combine control with a degree of machine intelligence. Their adaptability and responsiveness open up incredible possibilities in both industry and healthcare.

slide14:
Let's not forget neurological systems, like the human brain. The brain is a complex system of billions of neurons that process inputs from the environment and generate intelligent actions. Neurons decide whether to fire based on inputs, passing signals that enable everything from basic reflexes to complex reasoning. Together, these neural networks form the foundation of our physiological systems.

slide15:
The human system as a whole is incredibly intricate. It spans multiple levels—from genes and cells up to organs and entire physiological systems like the circulatory or digestive systems. When disease strikes, understanding these interconnections is crucial. This is the realm of systems biology or systems medicine—a field dedicated to studying the body as a network of interacting parts, so we can better understand, diagnose, and treat disease.

slide16:
Looking ahead, we can expect increasing convergence between robotic systems and human systems. Technologies like brain-machine interfaces are already emerging, with the potential to enhance memory, reasoning, and other cognitive functions. This merging of neuroscience and artificial intelligence represents an exciting frontier. As I once heard in a talk: physics helps us understand the external world, while neuroscience helps us understand ourselves. Bringing these fields together may fundamentally reshape what it means to be human.

slide17:
Given the vast complexity of all these systems, we need to narrow our focus. That's why in this course, we concentrate on linear systems—the simplest and most important class of systems, especially in medical imaging. Linear systems are manageable to study, yet powerful in their applications. Once we master them, we'll be ready to tackle the greater complexity of nonlinear systems.
In a linear system, two key properties must hold: additivity and homogeneity. Together, these form the superposition principle. We'll explore these in detail, starting with simple cases and building up to more rigorous requirements.

slide18:
Let's look at additivity first. If you know how a system responds to input F1, producing output K1, and separately how it responds to input F2, producing K2,

then additivity tells us the system's response to F1 + F2 is simply K1 + K2. This property allows us to analyze complex inputs by breaking them into simpler parts—a key advantage in engineering.

slide19:
Now consider homogeneity, or the scaling property. If a system produces output K1 for input F1, then scaling the input—multiplying it by some factor—should scale the output by the same factor. For example, if we double the brightness of an image input, the output image should also be twice as bright.
These two properties—additivity and homogeneity—are at the heart of linear systems. Understanding them clearly is essential, because they allow us to predict how systems behave, simplify complex problems, and build reliable imaging technologies.

slide20:
Let's test these ideas with a few examples. Suppose we have a system that maps positive numbers, x, to the square root of x. Is this system linear? Well, while the output is zero when the input is zero, that alone doesn't make a system linear. Remember, true linearity requires both additivity and homogeneity. The square root function doesn't satisfy either—if you take two inputs and sum their square roots, that doesn't equal the square root of the sum, and scaling the input doesn't scale the output proportionally.
On the other hand, consider a linear transformation in three-dimensional space, like multiplying a vector by a 3×3 matrix. This kind of operation satisfies both additivity and homogeneity, making it a linear system. Similarly, mathematical operations like integration and differentiation also qualify as linear operators. They obey the superposition principle, which combines both key properties of linearity in a compact form.

slide21:
Over the years, some have argued that the definition of a linear system might be redundant—perhaps additivity alone implies homogeneity, or vice versa. This is an intriguing idea, and it reminds us that crafting precise definitions, like those in mathematics or geometry, is both an art and a science. Just as Euclid defined geometry with five foundational postulates, and mathematicians later debated whether all were independent, we can ask similar questions about system properties.
In the next few slides, we'll look at when additivity and homogeneity are truly independent and when they might be connected. This kind of deep thinking helps us build stronger foundations for engineering practice.

slide22:
For continuous functions, additivity and homogeneity can indeed be shown to be equivalent. For example, if a system satisfies additivity, you can use that property repeatedly to demonstrate homogeneity for integer scalars, rational scalars, and ultimately for all real scalars, thanks to continuity. Conversely, if a system satisfies homogeneity, you can reason your way back to additivity.
But this also points to a larger truth: when building formal systems—whether in mathematics or engineering—we must ensure they are logically consistent. This is no trivial task, and it's part of what makes formal reasoning so challenging and rewarding.

slide23:
However, equivalence doesn't always hold. Consider the complex conjugate operation. This system satisfies additivity: the conjugate of a sum equals the sum of the conjugates. But it fails homogeneity when the scalar is complex, because conjugating a scalar times a complex number doesn't give the same result as scaling the conjugate by that scalar. This shows that additivity doesn't always imply homogeneity—you really do need to check both.

slide24:
Similarly, it's possible to have homogeneity without additivity. Imagine a function that applies one slope if the input is rational and another slope if it's irrational. Scaling works consistently because multiplying by a rational

scalar preserves the input’s type (rational or irrational), so homogeneity holds. But if you add two irrational inputs that sum to a rational number, the slopes change—and additivity breaks down.
These examples remind us why linear systems require both properties. We can’t shortcut the definition by assuming one implies the other in all cases.

slide25:
Together, additivity and homogeneity form what we call the superposition principle. This principle allows us to break down complex inputs and predict how the system will respond by summing simpler outputs.
Linear systems are easier to work with than nonlinear ones, and tools like Fourier analysis and convolution give us powerful ways to study them.
Importantly, even nonlinear systems can often be approximated as a collection of piecewise linear systems—so mastering linear systems lays the groundwork for handling more complex situations.

slide26:
One subtle idea is relative linearity. A system’s behavior might depend on the chosen coordinate system or reference point. For example, a linear function with an intercept (like $y = mx + b$) is still linear in its relative changes. If you redefine the origin to subtract out b, it behaves like a linear function passing through zero.
Similarly, for systems, we often focus on relative changes in input and output rather than absolute values. This helps us apply linear system theory even when the system’s baseline isn’t perfectly aligned with the origin.

slide27:
Let’s consider classic electrical components. A resistor follows Ohm’s law: $V = IR$. This is a linear relationship—double the current, and the voltage doubles. Whether we write $V = IR$ or $I = V/R$, the linearity remains clear. This is a simple, familiar example of a linear system.

slide28:
A capacitor introduces a bit more complexity. Its voltage depends on the integral of the current over time, plus an initial condition. If we focus on the relative change—the change from that initial state—the system behaves linearly. This shows how initial conditions affect our view of linearity, and why considering relative changes can be so useful.

slide29:
An inductor works similarly, but with current depending on the integral of voltage over time, or voltage depending on the derivative of current. Again, if we account for initial conditions, we can apply linear system theory to analyze these components effectively.

slide30:
In medical imaging and other engineering applications, we often focus not just on linear systems, but on shift-invariant linear systems. This means that if the input shifts—say, you move an object from one location to another—the output shifts in the same way.
For example, if I take your picture in this room, then take it again in my office, the image of you should look the same. The system’s response stays consistent despite the shift in location. This stability, or shift invariance, is a key property of many imaging systems.

slide31:
Think of ripples in a pond. If you drop a stone at one point, the ripples spread out. If you drop it at another point, the same pattern emerges—just shifted in space. This is an example of spatial invariance in a system.

slide32:
Similarly, a music player is temporally invariant. Whether you play your favorite song today or next week, it sounds the same. This consistency over time reflects shift invariance with respect to the time domain.

slide33:
Mathematically, shift invariance means that if the input shifts by some amount—say a—the output shifts by the same amount. The system's operation stays consistent, no matter where or when the input occurs.
Shift invariance, combined with linearity, gives us powerful tools for analyzing systems. Together, they allow us to model and predict system behavior with confidence, whether in medical imaging, signal processing, or other fields of engineering.

slide34:
Let's finish this part of our discussion by touching briefly on nonlinear systems. This is a topic marked with a green button—so it's something nice to know, not required at the core level of this course. Nonlinear science is a relatively new and complex area. Unlike linear systems, nonlinear equations are often difficult, or even impossible, to solve exactly. But thanks to modern computing, we can discretize these equations and compute solutions numerically—even if we have to do so through brute force.
Beyond practical applications, nonlinear systems are fascinating because they produce surprising, often counterintuitive phenomena. And anything that defies our intuition naturally grabs our interest.

slide35:
One famous example of a nonlinear system is the logistic map, a simple but powerful mathematical model. It describes how a population—say, of rabbits—evolves over time. The equation looks straightforward: next year's population depends on this year's population and a growth rate factor, r. But because of the nonlinear term involving the population itself, the system's behavior changes dramatically depending on the value of r.
When the population is small, growth is nearly linear. But as the population approaches the environment's limits, growth slows—resources become scarce, and the population stabilizes or even declines. This simple model reflects how nature regulates itself.

slide36:
What's remarkable about the logistic map is how its behavior changes as the growth rate increases. At low rates, the population settles to a single stable value. As the rate rises, the population starts oscillating between two values, then four, then many—and eventually behaves unpredictably, in what's called chaotic behavior.
This is chaos in a mathematical sense: a deterministic system, governed by precise rules, but with outputs that seem random because they're so sensitive to initial conditions. This is why long-term predictions, like weather forecasts, become unreliable—small differences in starting conditions can lead to dramatically different outcomes. This idea ties back to quantum mechanics, where we see that at a fundamental level, nature behaves probabilistically.

slide37:
Another example of a nonlinear system is the biological neuron. A neuron accumulates inputs over time—small signals have no effect, but once inputs surpass a certain threshold, the neuron fires, sending an electrical impulse. This threshold mechanism helps the nervous system filter out noise and only respond to meaningful stimuli.
We can model this mathematically as an artificial neuron, where inputs are combined through a weighted sum (the linear part), and this sum is passed through a nonlinear function that decides whether or not to fire. This is the basic building block of artificial neural networks.

slide38:
When we connect these artificial neurons in layers, we create a neural network. With many layers, we have what's called a deep neural network—the foundation of modern deep learning.
In these networks, the lower layers detect simple features, like edges or corners, while higher layers learn more abstract patterns, like faces or traffic

signs. Deep neural networks now drive advances in fields from facial recognition to self-driving cars. This is a powerful example of combining linear and nonlinear elements to solve complex problems.

slide39:
As we wrap up, let me remind you that alongside the exams and homework, those of you who are particularly motivated are welcome to pursue a class project. I recommend looking into machine learning—it’s an exciting and fast-growing area. If you’re interested, you can read my perspective article on this topic, and we can discuss potential directions for your project.

slide40:
If a project doesn’t interest you, that’s perfectly fine—you can focus on the regular course structure of lectures, exams, and homework. Speaking of which, I’ll upload today’s homework assignment for you to download. Please submit your solutions by midnight next Tuesday. After the deadline, we’ll post the answers for your review.
Thank you for your attention today!