slide1:

Hello my students!

Today we'll wrap up our MRI part by explaining pulse sequences. Although the title highlights pulse sequences, we'll also review the overall architecture of an MRI scanner and a few related concepts that tie everything together.

Toward the end of the lecture, I'll introduce some recent brain-initiative projects and a few emerging ideas in MRI research.The main emphasis today follows the green book chapter, so I strongly encourage you to read that section carefully. If you fully understand the ideas we cover in this lecture, you'll find the final exam easier.

Of course, I'll also share a few extra details and insights that go beyond the textbook—these will help you develop a deeper and more practical understanding of MRI as a whole.

slide2:

Let's take a quick look at our course schedule to see where we stand. After the MRI part, we have two more major imaging modalities: ultrasound and optical imaging, which are quite different but have their unique utilities. Finally, I will give you a very brief overall of the emerging area called "deep imaging" or "deep reconstruction".

slide3:

Last time, I asked you to read the textbook chapter on MRI. By now, many of you have probably gone through the green book and have some familiarity with the hardware components of an MRI system. Let's review these ideas together so you can build a stronger, more intuitive understanding.

MRI imaging relies on both hardware and software. The hardware provides the physical foundation—magnets, gradient coils, and radio-frequency coils—while the software defines the pulse sequences, the inversion schemes, and the reconstruction algorithms that convert raw signals into images. Together, all these parts form what we call an MRI scanner.

At the heart of the scanner is a strong, stationary magnetic field, which we denote as B-zero. Inside this B-zero field, all the tiny magnetic moments in the body—what we call the magnetization vector, or M-vector—naturally align along the direction of B-zero, which we usually define as the z-axis or G-direction.

Next, we introduce a radio-frequency field, called B-one. This is a rotating magnetic field applied perpendicular to B-zero. Because of this rotation, the B-one field exerts a torque on the M-vector, causing it to flip from its resting position toward the x-y plane. In other words, it moves the system from a stable, low-energy state to a higher-energy state.

Once magnetization has been flipped into the x-y plane, the spins begin to process, producing an alternating electromagnetic field. If we place detection coils near the sample—or the patient—this changing field induces an electrical signal known as the free-induction decay, or FID. That is the fundamental signal we detect in MRI.

After excitation, the M-vector gradually relaxes back to its original, low-energy position aligned with the B-zero field. This relaxation process releases energy, and we can monitor it to extract physical and biological information.

Now, this M-vector is not a single magnet; it represents the collective effect of countless microscopic spins—some pointing slightly up, some slightly down. The small imbalance between them adds up to produce a measurable net magnetization.

If we only used B-zero and B-one, we would get information about the entire sample as a whole—but no spatial localization. To achieve tomographic imaging, we need gradient coils that generate linearly varying magnetic fields, denoted G-x, G-y, and G-z. By controlling these gradients, we can selectively encode spatial information: G-z is typically used for slice selection, G-x for frequency encoding, and G-y for phase encoding. Together, these allow us to map the detected signals into k-space—a spatial-frequency domain that follows the Fourier-slice theorem. From there, we apply Fourier transforms and inverse Fourier transforms to reconstruct the final tomographic image.

So, the hardware provides the data, and the mathematics—mainly Fourier analysis—turns that data into images. This is a high-level overview of how MRI works. If you've already reviewed the textbook section on MRI scanner architecture, you should now recognize the key components we're talking about.

Let's move on and look at these elements in more detail.

slide4:

Now let's take a closer look at the signal model for a ninety-degree pulse. When the magnetization vector — the M vector — is flipped by ninety degrees, it moves entirely into the X-Y plane. This gives us the maximum signal, because the transverse magnetization is now at its full strength.

Once in the X-Y plane, the M vector begins to rotate, or process, around the main magnetic field. This motion induces an electric current in the nearby quadrature coils, and that's the MRI signal we detect. Because the magnetization both rotates and decays, the signal has two parts — an amplitude and a phase — so we describe it as a phasor.

Mathematically, we can write the two components like this:S-x of t equals M-naught times e to the power of negative t over T-two, times cosine of omega-naught times t.S-x of t equals M-naught times e to the power of negative t over T-two, times sine of omega-naught times t.

If we combine these two into one complex expression, we get: S of t equals S-x of t plus i times S-y of t, which is equal to M-naught times e to the power of negative t over T-two times e to the power of negative i omega-naught t. This compact form represents both the magnitude and the phase of the signal together. If you prefer real numbers, you could keep the sine and cosine forms separately, but the complex expression makes the math much easier—especially once we introduce phase encoding and frequency encoding later.

In the rotating reference frame, we remove the fast rotation caused by the main field B-zero, so the signal simplifies to: S of t equals M-naught times e to the power of negative t over T-two. This describes the exponential decay of the transverse magnetization—what we call T-two relaxation.

When the sequence is repeated after a repetition time, or T-r, the magnetization recovers along the Z-axis following this relation: M-z of T-r equals M-naught times the quantity one minus e to the power of negative T-r over T-one.

Now if we combine both effects—T-one recovery and T-two decay—the total signal becomes: S of t equals M-naught times the quantity one minus e to the power of negative T-r over T-one, all multiplied by e to the power of negative t over T-two. This signal depends on two groups of factors: biological parameters — proton density rho, T-one, and T-two — and technical parameters such as B-zero, T-r, T-e, and time t.

What's elegant here is how physics and mathematics come together. The sine and cosine functions we use are direct solutions of Maxwell's equations, which describe electromagnetic fields. And when we add phase and frequency encoding, those same sinusoidal functions naturally form a Fourier representation—the mathematical tool we use to reconstruct MRI images.

So, Maxwell's equations describe the physical interaction, and Fourier analysis provides the mathematical framework. Together, they create what we call the K-space theorem—a beautiful harmony between physics and mathematics that makes MRI imaging possible.

slide5:

Now that's really fascinating. Let's take a look at the overall MRI system diagram to understand how all these components work together.

At the very top level, the MRI scanner begins with a strong, static magnetic field, denoted as B-zero. This field provides the main alignment for all magnetic moments inside the body.

To this, we add a rotating magnetic field, called B-one, which is generated by the radio-frequency coil, or R-F coil for short. "R-F" stands for radio frequency, which corresponds to the Larmor frequency—the processional frequency of the spins.

When the system applies this B-one field, the spins are excited, and the resulting signal is induced in the same coil placed nearby. Because the coil operates at the same resonant frequency, it can both transmit and receive the signal effectively.

Now, let's trace the signal path through the system. Everything begins at the host computer, which contains the control software. The computer sends pulse sequence commands to a pulse programmer or software interface. These commands are then passed to a waveform generator and timing boards, which create precise control signals for the gradient amplifiers and the frequency synthesizer.

The frequency synthesizer defines the exact frequency of the R-F pulse, and the R-F amplifier boosts that signal to a level strong enough to drive the R-F coil. Through the transmit–receive switch, the system alternates between sending energy into the body and listening for the returning signal.

On the receiving side, the weak magnetic resonance signal first passes through a preamplifier, which increases its strength without distorting it. Then it goes through receiver blanking—a brief muting step to prevent overload when the transmitter is active. Next, the signal passes through another R-F amplifier, a demodulator, and a quadrature mixer, which separate the real and imaginary components of the complex signal.

This analog signal is then converted into digital form using an A–D converter, meaning analog to digital converter, and the data are stored in the system's memory, or RAM. From there, the computer performs a two-dimensional inverse fast Fourier transform, often abbreviated as two-D I-F-F-T, to reconstruct the spatial image from the frequency-domain data, also called k-space.

Finally, the reconstructed image is displayed on the computer screen.

So, to summarize — you have the B-zero field providing magnetic alignment, the B-one field for excitation and signal detection, and the gradient fields for spatial encoding. All these are coordinated by computer software that carefully synchronizes every pulse and acquisition step.

This is a high-level overview of how an MRI system operates—from the initial excitation pulse all the way to the final image reconstruction.

Next, we'll dive a bit deeper into some of the hardware details behind these components.

slide6:

Now let's talk about how the MRI system actually detects the signal. You don't need to worry too much about the mathematical details here—think of this as a conceptual overview. The signal we receive from the patient contains both a real part and an imaginary part, and we often combine them into a complex form to make the analysis simpler.

When we first learned about Fourier series, we started with real functions—things like sine and cosine. Later, we introduced complex notation, and suddenly the whole expression became much more compact, taking only about one-third of the space. That's exactly what we're doing here—it's purely a mathematical convenience.

So, we have two components: the real signal, written as S-R of t, and the imaginary signal, written as S-I of t. When we combine them, the total signal can be expressed as: S of omega equals the integral from negative infinity to positive infinity of the quantity S-R of t plus j times S-I of t, multiplied by e to the power of negative j omega t, integrated with respect to t. This simply represents the Fourier transform of our signal, showing how it behaves across different frequencies.

Now, in practical terms, the signal detected by the MRI coil is an oscillating voltage. For a typical 1.5-tesla scanner, this voltage oscillates at around 63.9 megahertz—that's an extremely high frequency—and its strength is only a few microvolts. Because it's so high and so weak, it's very difficult to digitize the signal directly.

To make digitization possible, we first have to demodulate the signal—in other words, we shift it from that very high Larmor frequency down to a much lower intermediate frequency, usually around 10.7 megahertz.

Here's how that works: The analog signal coming from the preamplifier is sent into a demodulator, then through a quadrature mixer. The quadrature mixer separates the signal into two parts—often labeled "I" for the in-phase component, and "Q" for the quadrature, or ninety-degree phase-shifted, component. These two parts correspond to the real and imaginary signals we discussed earlier.

After that, the signals pass through filters and variable-gain amplifiers to remove noise and adjust their strength. Finally, they go into two analog-to-digital converters, often called A–D converters, where the continuous analog signal is converted into discrete digital samples.

Once this conversion is done, the signal enters the digital domain—from there, the computer can perform all kinds of mathematical operations, such as Fourier transforms, image reconstruction, and signal averaging.

So, to summarize: the goal of this detection process is to convert a weak, high-frequency, analog signal into a lower-frequency, digital signal that can be handled by a computer. This step bridges the physical world of magnetic resonance with the digital world of image formation.

These are primarily electrical engineering concepts, but understanding them gives you a complete picture of how MRI signal acquisition really works.

slide7:

Now, let's talk about the background magnetic field, often referred to as B-zero. There are three main ways to generate this magnetic field inside an MRI system.

The mainstream approach, used in nearly all modern MRI scanners, is to use a superconducting magnet. This magnet operates in an extremely cold environment—close to absolute zero—so that the wire inside becomes superconductive, meaning it has zero electrical resistance. That's quite remarkable, because once current starts flowing in a superconducting loop, it can continue indefinitely without any energy loss.

If a superconducting magnet is not used, the second option is a permanent magnet. Permanent magnets can be very effective and relatively inexpensive when the total field strength is low—typically below about zero point three five tesla. These are often used in open MRI systems or low-end scanners, where high field strength is not essential.

The third method is to use electromagnets made from copper coils, where electric current is continuously injected to generate the magnetic field. However, this method has major drawbacks: running current through the coils produces heat, which leads to noise and energy loss, making it inefficient and less stable for imaging.

So, among the three methods, the superconducting magnet remains the primary choice for modern MRI scanners. It provides a strong, stable, and uniform B-zero field, which is essential for producing high-quality images.

If you look at the cross-sectional illustration in your textbook, you'll see how the magnet is constructed. At the center, we have the room-temperature magnet bore, where the patient lies. Surrounding it are layers containing the superconducting windings, which are immersed in liquid helium to keep them cold enough to maintain superconductivity. Outside that, there are layers of vacuum insulation and liquid nitrogen to minimize heat transfer from the environment.

The resulting magnetic field, B-zero, is oriented along the Z direction in our coordinate system, with X and Y lying perpendicular to it.

So, to summarize:Superconducting magnets are the standard in modern MRI systems,Permanent magnets are used for low-field or open designs, andResistive magnets using copper coils are mostly historical or experimental.

These are the key ways we generate the powerful and stable B-zero field that makes MRI imaging possible.

slide8:

If you look closely, you can see that the entire system is built like a layered cylinder. At the center is the main winding — a set of coils made from a special metal alloy. When this material is cooled to an extremely low temperature, its electrical resistance suddenly disappears. That's what we call superconductivity.

Normally, in any electrical circuit, a conductor always has some resistance — no matter how small. But once the temperature drops low enough, the resistance becomes exactly zero. That's an amazing physical phenomenon. In fact, a professor from our own physics department many years ago did pioneering work on superfluidity and superconductivity, and his research earned a Nobel Prize. So, this topic has deep scientific roots right here on campus.

In this system, liquid helium is filled into the helium vessel to cool the main windings to superconducting temperatures. The outer layers include vacuum insulation and cold shields — both inner and outer — to prevent heat from entering the chamber. There are also shim coils used to fine-tune the uniformity of the magnetic field inside the bore. The entire setup is housed inside a metal casing to maintain stability and safety.

When everything is working properly, a large current circulates continuously through the superconducting wire without any energy loss. That's the key advantage — it's energy efficient and can generate an extremely strong and stable magnetic field.

Now, if you look at the first link provided on the slide, it explains more about this superconductive design in detail. The second link, written partly in English and partly in Chinese, was authored by one of my former postdoctoral fellows, Dr. Wei. He proposed a theoretical interpretation of superconductivity and superfluidity from a quantum-mechanical perspective.

As you may recall, in quantum mechanics, matter is treated as both a wave and a particle. According to the standard theory, the probability wave that describes where a particle might be found is always non-negative — it can be zero or positive, but never negative. Dr. Wei's extension of the theory suggests that, under certain quantum conditions, the probability wave could actually take on negative values. That would mean a particle might seem to disappear in one region and reappear elsewhere instantly — a phenomenon sometimes described as quantum teleportation.

Of course, whether this interpretation is physically real is still open to debate, but it offers a fascinating theoretical connection between quantum behavior and macroscopic phenomena like superconductivity and superfluidity.

So, to summarize: in a superconducting magnet, the absence of resistance allows a persistent current to flow indefinitely, producing a strong and stable B-zero field essential for MRI. And at the same time, the underlying physics continues to inspire new ideas that bridge quantum mechanics and medical imaging technology.

slide9:

Now, I'm not an expert in electromagnet design, but let me share some essential information so you can understand the concept clearly.

The large superconducting magnet provides the strong, uniform B-zero field—that's our background magnetic field. But to create spatial encoding for imaging, we also need additional gradient fields.

So, how do we generate these gradient fields? A magnetic gradient means that the magnetic field changes linearly with position. In other words, the field strength increases or decreases in a straight-line fashion along one direction.

To achieve this, the MRI scanner uses three separate gradient coils, which produce fields we call B-X, B-Y, and B-Z. Each of these gradients corresponds to one of the three spatial directions: X, Y, and Z.

If you look at the illustration, you can see how the coils are wound differently to produce different kinds of gradients. The middle portion of each field—the region near the center of the magnet—is approximately linear. That's the part we actually use during imaging. Outside the center, the field is not perfectly linear, but that's okay, because we only rely on this central, uniform region for accurate encoding.

By wiring the coils in specific patterns, we can make the magnetic field vary linearly along any chosen direction. So, for example, one coil arrangement gives you a gradient along the Z direction, another gives a gradient along Y, and another along X.

We don't need to go deep into Maxwell's equations here—that would require advanced electromagnetic theory. For now, it's enough to have a heuristic understanding: when we wind the coils in a certain configuration, the resulting magnetic field changes linearly in space, and that's exactly what we need to distinguish signals from different locations inside the body.

So, remember — the B-zero field provides the uniform background alignment, and the gradient coils introduce controlled, linear variations along X, Y, and Z, which allow the MRI system to encode spatial information for image reconstruction.

slide10:

When you read the textbook chapter on MRI hardware, pay special attention to why we need to generate a linear magnetic field. The basic idea is actually very simple, and this diagram illustrates it nicely.

You might remember from high school physics that whenever an electric current flows through a wire, it produces a magnetic field that circles around the wire. If we take that wire and bend it into a loop, the small circular magnetic fields from each segment combine together. The result is a field that looks like the one produced by a tiny bar magnet, pointing in the direction given by the right-hand rule — curl your fingers along the current, and your thumb points toward the magnetic field direction.

Now, if you use a single loop, the magnetic field strength is small. But if you wind many loops together — for example, a coil with a hundred turns — the magnetic field becomes roughly a hundred times stronger. That's the basic principle behind electromagnets and, in our case, gradient coils.

On the left side of the figure, you see a loop with current flowing upward, generating a magnetic field that points to the right. On the right side, the current direction is reversed, producing a magnetic field that points to the left.

When we place these two coils symmetrically and run equal but opposite currents, something interesting happens. At the center between the coils, the magnetic fields cancel each other out, giving a field strength near zero. But if you move slightly to one side, one field becomes stronger while the other becomes weaker. That means the magnetic field increases linearly as you move away from the center — exactly the gradient we need.

So, in the central region, the total field varies in a nearly linear fashion, even though the actual fields from each coil are curved. This linear section is the useful part of the gradient field for imaging.

In summary:A single loop of current generates a circular magnetic field.Multiple loops amplify the field strength.Two symmetric, oppositely oriented coils produce a linear gradient in the central region.

That's the simple physical idea behind how gradient coils work — elegant, practical, and directly derived from the basic laws of electromagnetism you learned in school.

slide11:

Now let's connect this idea to the actual gradient coils used in MRI scanners. The linear field that we just discussed is formed between two opposing magnetic fields created by the left and right coil pairs. That configuration generates the Z-gradient field.

However, for the X and Y gradients, we can't arrange the coils in the same way—because the patient needs to pass through the tunnel of the magnet. So, the coil design must be adjusted while keeping the same physical principle.

For the Y-gradient, the setup uses a top coil and a bottom coil. Each coil forms what's called a saddle-shaped loop, but it still maintains a closed current path. The magnetic field from the top coil points downward, while the field from the bottom coil points upward. In the middle, these two fields cancel each other, creating a region of zero field. As you move upward, the downward field weakens and the upward field strengthens, forming a linear Y-gradient.

If we rotate this entire configuration by ninety degrees, we can generate the X-gradient using exactly the same concept.

When all these coil systems—the X, Y, and Z gradients—are combined, we can generate any arbitrary gradient field, often represented as G-X, G-Y, and G-Z. By carefully controlling the current in each set of coils, we can produce gradients in any direction or combination needed for imaging.

In practice, the gradient coils are designed in multiple layers, and additional shim coils are included to fine-tune the magnetic field so it remains as linear as possible within the imaging region.

Although the engineering details are complex, the fundamental idea is simple. Using Maxwell's equations—just four equations that describe how electric and magnetic fields behave—we can simulate the magnetic field using finite element computation. By adjusting coil shapes, adding new loops, and fine-tuning current paths, engineers can iteratively refine the field design until it achieves excellent linearity.

So even though the hardware looks complicated, the underlying physics is straightforward: we excite the coils in such a way that their combined fields produce the precise gradient patterns required for MRI imaging.

And one final point—during imaging, these gradient coils rapidly switch on and off, producing mechanical stress that makes the loud knocking sounds we hear during MRI scans.

That's the practical side of how gradient coils work—both physically and acoustically.

slide12:

Now, let’s talk about the B-one field, which is the missing component we haven’t discussed yet. This field is generated by the radio-frequency coil, or R-F coil.

The most common type of RF coil used for transmitting the B-one field is called a birdcage coil, shown on the left. You can understand this coil design quite intuitively. It looks like the frame of a cylindrical cage — a circular structure with evenly spaced conducting bars connected at both ends by rings.

When an alternating current flows through these bars, it creates a rotating magnetic field inside the bore. This rotating field is the B-one field, which is applied perpendicular to the main magnetic field B-zero. The B-one field is responsible for exciting the protons and flipping the magnetization vector, M, away from its alignment with B-zero.

On the right side of the slide, you can see other kinds of RF coils as well. The surface coil, shown in the middle, is used for imaging regions close to the surface of the body — for example, the shoulder or knee. Because it’s placed near the area of interest, it gives excellent signal-to-noise ratio locally, though it doesn’t cover deep tissues as well.

The third design, shown on the right, is a phased-array coil. This one combines multiple small coil elements arranged in an array. Each coil picks up signals from a different region, and the signals are combined by the system to produce a high-quality image. Phased arrays are very useful for spine imaging and other applications where we need both wide coverage and high sensitivity.

In summary:The birdcage coil is typically used for transmission — it creates a uniform B-one field, ideal for brain or whole-body imaging.The surface coil is best for localized, high-resolution scans near the body surface.The phased-array coil allows us to combine multiple receiver channels for stronger and faster imaging.

So, the RF coils are essential for both transmitting the B-one excitation field and receiving the MRI signal from the body, making them one of the most important components in the entire imaging chain.

slide13:

Now, let’s take a closer look at how the birdcage coil produces a uniform B-one field. If you consider the two cross-sections shown here, the idea is quite intuitive. When the current in one connecting wire flows in a particular direction, it generates a local magnetic field pointing downward. To produce a local field pointing upward, the current in the adjacent wire must flow in the opposite direction, forming a closed loop.

This basic configuration—currents flowing in opposite directions through adjacent loops—creates magnetic fields that point downward in some regions and upward in others. When multiple loops or cross-sections are added around the cylinder, these individual magnetic fields combine to form a uniform field across the entire region inside the coil.

Mathematically, this can be shown by modulating the current as a sinusoidal function of the angular position, which we call phi. If you plot the current intensity as a function of the phi angle, it varies like a sine wave—starting from zero at phi equals zero degrees, reaching a maximum at ninety degrees, and then becoming negative at one hundred eighty degrees. This sinusoidal modulation ensures that the magnetic field intensity remains uniform across the cross-section.

To make the B-one field rotate in space, the current itself is also modulated over time, so that the maximum and minimum of the field continuously move around the coil. As a result, the B-one field rotates smoothly, allowing it to flip the magnetization vector, M, in a controlled and consistent way during excitation.

In essence, this is how we achieve a uniform rotating magnetic field that interacts with the main B-zero field. Together, the B-zero field, B-one field, and the three gradient fields—G-sub-X, G-sub-Y, and G-sub-Z—form the complete electromagnetic environment of the MRI system.

What's truly fascinating is that all of these ideas—the generation of magnetic fields, their directions, and the way they add up—can actually be explained using high school physics: Ampère's law, the right-hand rule, and the principle of superposition.

Of course, for real MRI coil design, we need precise optimization to make sure the fields are uniform and stable. That's done through finite element analysis, a numerical technique that lets engineers compute the detailed magnetic field distribution and fine-tune the geometry of the coil.

So, while the physical concept is beautifully simple, the practical realization requires sophisticated computation and engineering.

slide14:

Now, let's look at what happens when we use a surface coil instead of a full birdcage coil that covers the entire patient. A surface coil is a smaller coil placed directly near the region of interest—for example, near the spinal cord or breast. Because it's positioned so close to the tissue, it produces a very strong B-one field right next to the coil surface. However, the field strength drops off rapidly with distance, as you can see from the plot on the left. So the signal intensity is highest near the coil and becomes much weaker for tissues deeper inside the body.

This makes surface coils particularly useful for localized imaging, where we only need high-quality signals from a specific region, such as the spine or breast. They are also very important for parallel MRI, where multiple small coils are used together to collect data simultaneously from different parts of the body. This parallel setup helps speed up imaging and improve the overall signal-to-noise ratio.

In the past, implementing parallel MRI was quite challenging because it required careful calibration and coordination among many coils to cover the entire patient. But with modern electronics and coil arrays, this technique has become highly effective and widely used.

So, with this slide, we've completed our overview of the MRI scanner hardware — the superconducting magnet for the B-zero field, the RF coils for generating and detecting the B-one field, and the gradient coils for spatial encoding along X, Y, and Z directions.

Understanding how these fields — B-zero, B-one, and the gradient fields — work together is enough to grasp the essential physics behind MRI. Once you're comfortable with these concepts, you'll have a solid foundation for understanding the imaging process itself.

In the next section, we'll shift our focus to pulse sequences — beginning with the spin echo sequence, which allows us to measure T-one-weighted, T-two-weighted, and proton-density images. Then, we'll move on to gradient echo sequences for faster imaging, and later discuss MRI angiography, diffusion-weighted imaging, spectroscopy, and MRI contrast agents, along with their safety considerations.

This marks the transition from the hardware foundation of MRI to the signal manipulation and image formation techniques that make this technology so powerful.

slide15:

We've now completed our overview of the MRI scanner hardware—including how the system is configured and how the three key magnetic fields are generated: the B-zero field from the main magnet, the B-one field from the radio-frequency coil, and the gradient fields, G-X, G-Y, and G-Z, which are produced by the gradient coils.

Understanding these fundamentals gives you a strong foundation for grasping how MRI actually works. You now know how each component contributes—how the hardware creates, manipulates, and detects magnetic signals that form the basis of MRI imaging.

Next, we'll move on to the pulse sequences, which define how we control these fields in time to generate different types of images. The first and most fundamental one we'll discuss is the spin echo sequence. This sequence is the key to producing T-one-weighted, T-two-weighted, and proton-density images—the three major image contrasts used in clinical MRI.

After that, we'll cover gradient echo sequences, which are designed for faster imaging and are commonly used in dynamic studies. Then we'll extend our discussion to MR angiography for blood flow visualization, diffusion-weighted imaging for tissue microstructure, and MR spectroscopy for biochemical analysis. Finally, we'll talk about MRI contrast agents—how they work, and the important safety considerations associated with their use.

So from this point onward, we'll move from the hardware foundation of MRI into the timing and control of signals, which is where much of the image contrast and diagnostic power of MRI truly comes from.

slide16:

Now we're entering one of the most important ideas in MRI — phase and frequency encoding. Essentially, MRI operates in what we call a Fourier imaging mode, so understanding the Fourier transform is absolutely critical to understanding how images are formed.

This slide is really the key to everything that follows. If you can grasp what's happening here, all the other pulse sequences — no matter how complicated they look — will make much more sense.

In the stationary reference frame, the overall signal that we detect — or equivalently, the magnetization vector M — is proportional to the proton density, written as rho of x, y. This function describes how many protons exist at each point in a given slice.

When we apply phase encoding and frequency encoding gradients, we introduce exponential phase factors into the signal. Mathematically, this means we multiply the signal by e to the power of negative j times gamma times G-x times x times t, and also by e to the power of negative j times gamma times G-y times y times tau pe, where gamma is the gyromagnetic ratio, G-x and G-y are the gradient field strengths, and t and tau pe are the times for which those gradients are applied.

These exponential terms introduce both phase and frequency variations across the image slice. The phase encoding adds a phase shift that depends on position in the Y direction, and the frequency encoding controls how the signal oscillates in time for each X position.

Together, these variations give each location — each small voxel — a unique combination of frequency and phase. That's what allows us to reconstruct a tomographic image. Without this encoding, all the information from the slice would be averaged together, and we'd lose spatial detail.

Now, by performing a few mathematical steps — specifically dividing and multiplying by two pi — we can rewrite these expressions in a form that matches the Fourier transform. We define two variables: k-x equals gamma times G-x times t divided by two pi, and k-y equals gamma times G-y times tau pe divided by two pi.

Using these definitions, the detected signal S of k-x, k-y becomes proportional to the double integral of rho of x, y times e to the power of negative j two pi times (k-x times x plus k-y times y) integrated over x and y. This is exactly the two-dimensional Fourier transform of the object function, rho of x, y.

So, the measured data that we collect in MRI — what we call k-space — is simply the Fourier transform of the spatial proton density. Once we've filled k-space by varying the gradient fields, we apply an inverse Fourier transform to reconstruct the image in real space.

Now, in practice, the process can go in either order — you can perform phase encoding first, followed by frequency encoding, or vice versa. For each phase-encoding step, we record one line of data in k-space, and by repeating this process many times, we gradually fill the entire k-space grid.

The green trajectory shown on the slide illustrates how we move through k-space: phase encoding determines the starting position, and frequency encoding defines the direction in which data are collected. Different pulse sequences change the way we traverse k-space — sometimes line by line, sometimes in spirals or zigzags — but the underlying principle is always the same.

Finally, because of symmetry, if the image function is purely real, we theoretically need to sample only half of k-space. However, in practice, we usually sample the full space to ensure higher accuracy and image stability.

So, to summarize: phase encoding and frequency encoding translate spatial information into distinct frequency and phase variations. The signals we collect form the Fourier representation of the object in k-space. By applying the inverse Fourier transform, we convert this frequency-domain data back into a real, tomographic image.

This is the essence of the k-space theorem, and it's the mathematical foundation of nearly all MRI imaging methods.

slide17:

Now that we understand how MRI encoding works through Fourier principles, let's see how this idea is applied in one of the most fundamental pulse sequences — the spin-echo imaging sequence.

In this diagram, the horizontal axis represents time, and you can see the major events in the sequence: the radio-frequency pulses, the gradients for slice selection, phase encoding, and frequency encoding, and finally, the signal readout by the analog-to-digital converter.

The process begins with a ninety-degree radio-frequency pulse, which flips the magnetization vector, M, from the z-axis into the x-y plane. Immediately after this, the spins start to dephase because of local field inhomogeneities — some spins precess a bit faster, others a bit slower. As a result, the transverse magnetization gradually fades away.

To correct this dephasing, a one-hundred-eighty-degree pulse is applied halfway through the time interval known as T E, or echo time. This pulse effectively flips all the spins — the ones that were ahead now fall behind, and those that were behind move ahead. After an equal amount of time, all the spins come back into phase, and an echo signal is formed. That's why we call it a spin echo — the system "rephases" itself to produce a measurable signal.

The T R, or repetition time, is the total time between successive ninety-degree pulses. By changing T E and T R, we can emphasize different types of image contrast — for example, T-one weighting, T-two weighting, or proton-density weighting.

Now, if you look at the lower part of the diagram, you'll see three gradient waveforms:G - slice selects the imaging slice,G - phase applies a brief pulse to encode phase differences along the Y direction, andG - freq, or the frequency-encoding gradient, is turned on during the echo readout to separate spatial information along the X direction.

The signal is then digitized by the A/D converter, forming one line of data in k-space. By repeating the process with different phase-encoding gradients, we fill up all the lines of k-space. After that, an inverse Fourier transform converts the data into the final image.

So, in summary: the spin-echo sequence uses a ninety-degree pulse to excite the spins, a one-hundred-eighty-degree pulse to rephase them, and a carefully timed set of gradients to encode spatial information in both frequency and phase. This sequence forms the foundation of most MRI imaging techniques, and every other sequence — including gradient echo, echo planar, and spiral imaging — can be understood as a variation of this same principle.

slide18:

Now, let's look at how multi-slice imaging works in MRI.

In MRI, slice selection allows us to excite only a specific region of the body instead of the entire volume. This is achieved by manipulating the resonance frequency of spins within a narrow slice using a gradient field. When we apply the gradient, the magnetic field strength changes linearly along the gradient direction. As a result, spins in different positions experience slightly different Larmor frequencies.

By sending in a radio-frequency pulse that matches the frequency range of one specific slice, we can flip the magnetization vector, M, only within that selected slice. The rest of the body remains unaffected.

However, because the field varies linearly across the slice, the spins acquire different phase shifts — spins in regions of stronger field accumulate larger phase factors, and those in weaker fields accumulate smaller ones. This variation is called dephasing.

To correct for this, we apply a refocusing, or rephasing pulse, that has the opposite polarity. This pulse restores the overall phase coherence. It doesn't bring every spin perfectly back to its original state but rephases them on average, so that the net signal becomes measurable again.

In the context of slice selection, we call this process defocusing and refocusing. The refocusing gradient usually has half the area of the initial defocusing gradient because we only need to correct the average phase dispersion, not the exact microscopic alignment.

A similar process occurs during frequency encoding. When the frequency-encoding gradient is applied, it introduces another phase variation. To recover the maximum signal strength, a negative gradient pulse with half the area of the first one is applied to refocus the spins again. Although the polarity of the gradient may appear the same in diagrams, the one-hundred-eighty-degree pulse inverts the spin directions, effectively flipping the polarity of the phase, so that the final echo appears consistent in direction.

The phase-encoding gradient then introduces controlled phase shifts across the slice, allowing each spatial location to contribute a unique signal component. These differences are what allow us to reconstruct distinct pixel values in the final image.

In the spin-echo sequence, the two key timing parameters are the echo time, T-E, and the repetition time, T-R. By carefully adjusting these parameters, we can control image contrast — and, importantly, the time between echoes can also be used for additional slice selection.

That's how we can acquire multiple slices within a single T-R period, as you see in the example here. Each slice is excited and read out in sequence using its own frequency offset, labeled as omega-one, omega-two, omega-three, and so on. This is called multi-slice imaging.

It's a highly efficient way to cover a larger region of the body without increasing the total scan time.

slide19:

In multi-slice imaging, we can further optimize how slices are acquired. For example, we can excite slice three, then slice four, and then slice two, alternating the order to maximize time efficiency. This alternating strategy helps reduce the blurring or cross-talk that can occur when slices are too close together in time or space. By allowing enough separation between successive excitations, we achieve cleaner and more distinct images.

Now, the spin-echo sequence we just discussed provides an excellent signal-to-noise ratio, because it effectively cancels out the dephasing caused by T-two-star effects. This gives us a pure T-two-weighted signal that's highly reliable for tissue contrast. However, the main drawback is that the spin-echo sequence takes a long time. A full imaging session can easily last half an hour or more, which is often too slow for many clinical applications.

To address this limitation, we use gradient-echo imaging as an alternative. In a gradient-echo sequence, there is no one-hundred-eighty-degree refocusing pulse. This omission greatly shortens both the echo time, or T-E, and the repetition time, T-R. Because there's no refocusing pulse, gradient-echo imaging does not compensate for magnetic field inhomogeneities, so the signal decays more quickly according to T-two-star relaxation rather than pure T-two.

To keep the signal strong and the scan fast, the echo time — T-E — is kept as short as possible, typically just a few milliseconds. In addition, the flip angle of the radio-frequency pulse is reduced to a small value, usually well below ninety degrees. This allows for rapid repetition of pulses and enables much faster imaging — often completing a full scan in less than a minute.

There are several variations of gradient-echo sequences, including: FLASH, which stands for Fast Low-Angle Shot; FISP, or Fast Imaging with Steady Precession; GRASS, which uses Gradient-Refocused Acquisition in the Steady State; and SSFP, or Steady-State Free Precession. Each of these methods is built on the same principle: using small flip angles and gradient rephasing instead of a 180-degree pulse to achieve high-speed imaging.

For any given repetition time, T-R, there exists an optimal flip angle that maximizes the signal-to-noise ratio. This is known as the Ernst angle, given by the formula: alpha Ernst equals cosine inverse of e to the power of negative T-R divided by T-one. By choosing the Ernst angle, we balance image contrast and signal strength for efficient imaging.

So, to summarize: spin echo sequences give you excellent image quality but require long scan times, gradient echo sacrifices some signal stability for dramatically faster acquisition, and together, these two techniques form the foundation of nearly all MRI pulse sequences used in practice today.

slide20:

Now let's look at three-dimensional gradient echo imaging. In this method, we extend the basic idea of gradient echo imaging into three spatial dimensions by using two separate phase-encoding gradients and one frequency-encoding gradient. Here's how it works: we begin with a radio-frequency pulse that flips the magnetization vector, M, by a small angle alpha, much less than ninety degrees. Then, the system applies the first phase-encoding gradient, labeled G phase one, followed by a frequency-encoding gradient, labeled G freq, and finally a second phase-encoding gradient, G phase two. This sequence allows us to encode spatial information along all three axes — x, y, and z — so the final data form a three-dimensional Fourier transform. Mathematically, the reconstructed signal can be expressed as a triple integral over kx, ky, and kz, that is, rho of x, y, z equals the integral of S of kx, ky, kz times e to the power of j two pi times (kx x plus ky y plus kz z), integrated over dkx, dky, and dkz.

Each phase-encoding step changes the gradient strength slightly, producing a different k-space plane. When we collect all these planes and perform a three-dimensional inverse Fourier transform, we obtain a volumetric image of the scanned region. Because the signal is read out immediately after excitation, it decays rather quickly due to T-two-star relaxation. However, this approach allows for very rapid data acquisition, enabling high-resolution three-dimensional imaging in a relatively short time. Compared with spin-echo sequences, the signal-to-noise ratio is somewhat lower because gradient echo does not compensate for field inhomogeneities. But the speed advantage is substantial — scans can be completed much faster.

To achieve this balance, we use a small flip angle rather than a large one. If the flip angle is too large, the magnetization along the z-axis doesn't have enough time to recover before the next excitation, and the signal from subsequent repetitions becomes weaker. By using a smaller flip angle and a short repetition time, or T-R, we can repeatedly excite and read out signals quickly, maintaining good image quality while minimizing scan duration. In short, three-dimensional gradient echo imaging sacrifices a bit of signal strength for a dramatic improvement in imaging speed. It is particularly useful for volumetric scans of the brain, heart, and other organs where time is critical.

And if we push the concept even further, we arrive at echo planar imaging, which can acquire an entire image — or even a full volume — in just one or a few rapid readouts. Echo planar imaging is the foundation for real-time MRI and functional MRI studies, where both speed and temporal resolution are essential.

slide21:

Now we come to one of the fastest MRI imaging techniques — Echo-Planar Imaging, or EPI. In this sequence, the magnetization vector, M, is first flipped into the x–y plane using a ninety-degree radio-frequency pulse. Next, slice selection is performed to isolate a specific plane — for example, the x–y plane of the imaging region. Once the slice is selected, the imaging process begins with a combination of phase encoding and frequency encoding that alternate in an oscillating pattern. This means that the frequency-encoding gradient repeatedly switches direction — first positive, then negative, then positive again — while data are continuously acquired.

As a result, the signal alternates between positive and negative readouts, forming a continuous train of echoes. Each echo corresponds to one horizontal line in k-space, and the entire pattern forms a zig-zag trajectory across k-space, as shown on the right side of the slide. Let's interpret this step by step: slice selection determines the imaging plane, typically the x–y plane. Phase encoding moves the position in k-space vertically — that is, along the k-y direction — shifting the data line by line. Frequency encoding determines the horizontal readout, or the k-x direction, as the data are collected. Each line of k-space is traversed in alternating directions: the first line moves from left to right, the next line from right to left, and so on. This continuous acquisition pattern allows us to collect a full image, or at least a major portion of it, in a single shot.

However, as data acquisition progresses, the signal amplitude gradually decreases due to T-two-star decay, which limits how long useful data can be collected. Because of this decay, echo planar imaging may sometimes cover only part of k-space during one acquisition, requiring multiple passes to fill in the remaining data. Even with that limitation, EPI is dramatically faster than conventional three-dimensional gradient-echo imaging. In fact, it can complete a full image in just a fraction of a second. That's why echo planar imaging has become the foundation for real-time MRI and functional MRI, where rapid temporal sampling is more important than perfect image contrast.

And when we want to go even faster, we can move beyond echo planar imaging to spiral imaging, which traces k-space in smooth circular trajectories, further reducing acquisition time.

slide22:

Now, let's move on to spiral imaging, a powerful and elegant approach that enables even faster data acquisition than echo-planar imaging. In spiral imaging, there is no separate phase encoding step. Instead, both gradient fields — G x and G y — continuously change their amplitudes over time. The relative strength of these two gradients determines the instantaneous slope of the trajectory in k-space.

At the very beginning, the system starts from the origin of k-space, meaning the initial point of the phase vector is zero. As time progresses, the amplitudes of G x and G y vary smoothly, causing the trajectory to move outward in a spiral pattern. The direction of this spiral is controlled by the relative phase and amplitude ratio between the two gradients. Initially, the trajectory moves along a specific direction for a short distance. Then, as the relative gradient amplitudes change, the slope and direction of the path gradually rotate. This continuous change allows the k-space trajectory to trace a smooth spiral curve, expanding outward from the center as higher spatial frequencies are collected.

The advantage of this approach is that it provides continuous control over both gradients, allowing highly efficient coverage of k-space. Depending on the design, the spiral can move from the center outward, from the periphery inward, or even use multiple interleaved spirals that fill k-space more densely and reduce artifacts. Because of this flexibility, spiral imaging offers excellent speed while still maintaining high spatial resolution. It captures low-frequency components first — which contain the overall image contrast — and then gradually collects higher-frequency components that define fine structural detail.

After data acquisition, the spiral k-space data must be interpolated onto a rectangular grid before applying the two-dimensional inverse Fourier transform to reconstruct the image. This reconstruction yields T-one-weighted, T-two-weighted, or proton-density-weighted images, depending on how the pulse sequence parameters — such as flip angle, T-E, and T-R — are chosen. Understanding how data are collected in k-space is crucial. If you fully grasp how the gradients define the k-space trajectory — whether linear, zigzag, or spiral — the entire process of MRI image formation becomes intuitive. Otherwise, it can seem abstract or confusing.

Finally, these flexible gradient-based trajectories also open the door to specialized MRI pulse sequences. For example, in MR angiography, we can tailor the sequence to highlight blood flow, just as CT angiography does for X-ray-based imaging. By manipulating the gradient and timing parameters, we can emphasize vascular structures and obtain clear, high-resolution images of the circulatory system. So, spiral imaging not only accelerates acquisition but also serves as a bridge to more advanced and application-specific MRI techniques, such as angiography, diffusion imaging, and functional MRI.

slide23:

Let’s now talk about Time-of-Flight Imaging, or TOF, which is a fundamental technique in Magnetic Resonance Angiography, or MRA. In a standard CT image, blood vessels and soft tissues often appear very similar in density. As a result, it’s difficult to distinguish the vasculature — you can’t easily tell whether a vessel is open, narrowed, or completely blocked. To solve this problem in CT, we use a contrast agent — typically iodine — a liquid metal with a very high linear attenuation coefficient. When injected into the bloodstream, iodine increases the attenuation of X-rays in the blood, making vessels appear much brighter than the surrounding soft tissue. If a vessel that should be wide suddenly appears narrow or faint in the CT image, that suggests a vascular blockage, which could require stenting or further cardiac intervention. This is the basic principle behind CT angiography.

Now, for MRI angiography, the principle is different, but the goal is the same — to make the blood vessels appear distinct from the surrounding tissue. Here’s how it works: we first select a single imaging slice, and then apply a radio-frequency pulse, typically using an echo-planar imaging sequence, to flip the spins — those small magnetic dipoles — within that slice. At that instant, both the blood within the vessel and the stationary tissue generate signals that are quite similar, so initially, the vessel is not clearly visible. But if we wait a little longer, something interesting happens. The blood that was originally in the imaging slice flows out of the plane, and new blood — containing unflipped spins — flows in from outside. This is what we call the through-plane mechanism.

Because the newly arrived blood has not been excited by the RF pulse, it produces no signal in that moment. Meanwhile, the surrounding stationary tissue continues to emit its signal, which gradually decays but does not vanish immediately. As a result, when we collect the data after a short delay, the moving blood appears

bright, while the stationary tissue looks relatively dark. This contrast difference allows us to visualize the vasculature clearly.

The Time-of-Flight effect is strongest when the blood flow is perpendicular to the imaging slice — that is, when blood moves directly through the plane of excitation. If the vessel runs within the plane of the slice, then the same spins remain in the imaging region, and the through-plane mechanism does not work effectively.

So, to summarize: in CT angiography, iodine contrast enhances X-ray attenuation. In MR angiography, the motion of blood itself provides the contrast, through the time-of-flight effect and the through-plane mechanism. This principle forms the foundation of non-contrast MR angiography, allowing clear visualization of blood flow without any injected contrast agents.

slide24:

Now let's move on to Phase-Contrast Imaging, which is another powerful method used in MR angiography. Unlike the time-of-flight technique that works mainly through the through-plane mechanism, phase-contrast imaging is an in-plane mechanism — meaning it can detect blood flow in any direction, not just perpendicular to the slice. The basic idea is similar to what we've seen in echo-planar imaging. We start with an alpha pulse that flips the magnetization vector, M, by an angle alpha for slice selection. All the spins within the selected slice then work together to produce a coherent signal. As usual, frequency encoding and phase encoding are applied to collect the Fourier information from that slice.

However, in a normal Fourier-based MRI scan, the pixel intensities for blood and stationary tissue — in terms of T-one, T-two, or proton density — often appear very similar. This makes it difficult to distinguish flowing blood from static tissue. To separate them, we introduce a pair of bipolar gradient pulses — one positive and one negative — before data acquisition. The first, or positive, gradient pulse adds a phase shift to all spins, and the second, or negative, gradient pulse adds an equal but opposite phase shift. For stationary spins — such as those in soft tissue — these two effects cancel perfectly, leaving no net phase change.

But if the spins are moving, as in flowing blood, something different happens. Between the two gradient pulses, the blood cells physically move to a new position, so they experience slightly different local magnetic fields. As a result, the positive and negative gradients do not cancel out completely. The motion introduces an extra phase term that accumulates over time, and this additional phase is directly proportional to the velocity of the moving spins. By measuring the difference in phase between two acquisitions — one with and one without motion encoding — we can isolate the component of the signal that depends purely on blood flow. This produces a phase-contrast image, where bright regions correspond to areas of flow and darker regions to stationary tissue.

In mathematical form, this accumulated phase, phi, is given by an integral of the gradient over time: phi equals the integral from zero to tau over two of gamma times G-x times x of t, d-t, plus the integral from tau over two to tau of gamma times negative G-x times x of t, d-t. The net phase difference reflects the motion of the spins under the bipolar gradients. Conceptually, you can think of time-of-flight imaging as a passive effect, where the signal difference arises simply because new, unflipped blood enters the slice. In contrast, phase-contrast imaging is an active mechanism, where we intentionally encode blood motion into the phase of the signal.

And unlike X-ray phase-contrast imaging — which measures refractive index variations — MRI phase-contrast imaging detects the phase shift of the magnetic vector caused by moving spins. It doesn't matter whether the blood flows up, down, or sideways — as long as the spins move through regions with slightly different magnetic fields, the phase difference can be detected and visualized. By exploiting these motion-induced phase shifts, we can generate clear vascular images and even quantify blood velocity or flow rate. Together with time-of-flight imaging, phase-contrast imaging provides a complete toolkit for visualizing the circulatory system without the need for injected contrast agents.

slide25:

Now, let's talk about diffusion-weighted imaging, or DWI for short. Remember, our body is made up mostly of water, and these water molecules are constantly moving around. This motion — or diffusion — becomes physiologically and pathologically important, especially when something goes wrong at the cellular level. For example, when a cell membrane becomes swollen or damaged, water can leak out of the cell, altering the normal diffusion and perfusion patterns in that tissue. That's why studying how water moves can tell us a lot about what's happening inside the body — particularly in the brain — in cases like stroke, tumor, or inflammation.

Now, conceptually, diffusion-weighted imaging is similar to angiography imaging, because both deal with motion — but in DWI, we're focusing on the microscopic motion of water molecules rather than blood flow. The pulse sequence looks a little different, but the underlying idea remains very similar. Here, we use a spin-echo sequence with a 90-degree excitation pulse followed by a 180-degree refocusing pulse. You'll notice two gradient pulses, placed symmetrically on both sides of the 180-degree pulse — these are called diffusion-sensitizing gradients.

We already know that a 180-degree pulse flips the magnetization vector upside down. So, from the perspective of the first gradient pulse, the second one has the opposite polarity. If nothing in the tissue is moving — meaning the water molecules are stationary — then whatever phase shift is induced by the first gradient will be perfectly canceled out by the second one. In that case, there is no net signal change. But if the water molecules move between the two gradient pulses — for example, diffusing from one position to another — they experience slightly different magnetic field strengths at those two locations. As a result, the phase shifts from the two gradients no longer cancel out perfectly. This causes a measurable signal loss, which tells us that motion or diffusion has occurred.

The signal attenuation due to diffusion depends on three main factors: first, the gradient strength, represented by capital G; second, the duration of each gradient pulse, denoted by the Greek letter delta; and third, the time separation between the two gradients, represented by capital Delta. Together, these parameters determine how sensitive the sequence is to molecular motion. The diffusion-related signal intensity can be expressed as I of G n equals I zero times e to the power of negative D gamma squared G n squared delta squared times (Delta minus delta over three). This equation shows how the measured signal decreases with stronger diffusion.

Practically, we don't need to go through the derivation — the key takeaway is that diffusion efficiency can be quantified and used as a diagnostic marker. In a typical spin-echo sequence, the gradients can be applied along the x, y, or z direction, depending on which axis of diffusion we want to measure. For example, diffusion in the through-plane direction — perpendicular to the imaging slice — can reveal how water molecules move between layers of tissue. When diffusion appears unusually high, it may indicate disrupted

or damaged cellular structure, such as in ischemic tissue. So, by analyzing diffusion in this way, we gain valuable insight into both normal physiology and pathological conditions.

slide26:

Now, the diffusion rate we just talked about is not constant. It varies from place to place — and more importantly, it also depends on direction. If you measure diffusion along one direction, you'll get a certain value. Measure it along another direction, and you'll likely get a different one. And along yet another direction, it changes again. So, unlike uniform motion that can be described by a simple vector — where the magnitude represents the speed and the direction shows the path — here, the diffusion rate itself depends on orientation. To describe such directional dependence, we need a more general concept than a vector — and that's where the tensor comes in.

Mathematically, we call this the diffusion tensor, often denoted by capital D. It's a three-by-three symmetric matrix that characterizes how diffusion behaves along different spatial directions. Each element of this matrix describes the relationship between diffusion along one axis and its coupling with another axis — for example, D one one, D one two, D one three, and so on. This means the diffusion process is anisotropic, or directionally dependent. In biological tissues such as white matter in the brain, water molecules tend to diffuse more easily along the fibers than across them, because the cell membranes and myelin sheaths restrict motion in the perpendicular directions. Similarly, in vasculature, blood and water tend to follow the principal axes of the vessels.

Now, if we extend this idea further, we can describe the random diffusion process mathematically. The probability that a molecule starting at position r-naught moves to position r after some diffusion time, tau, can be expressed using a multivariate Gaussian function. This is written as: p of r given r-naught and tau equals one over (4 pi tau) to the power three-halves times the square root of the determinant of D, multiplied by e to the power of negative (1 over 4 tau) times (r minus r-naught) transpose, times D inverse, times (r minus r-naught). In simpler terms, this function describes the probability distribution of molecular displacement — how far and in what direction the molecules are likely to move during diffusion.

So, diffusion tensor imaging, or DTI, is the MRI technique that measures this tensor for every voxel in the image. It gives us a complete three-dimensional map of how water diffuses in tissue — allowing us to infer the orientation of fibers and structural connectivity in the brain and other organs.

slide27:

Now that we understand diffusion tensors, we can visualize them in a colorful way by encoding the three main spatial directions—x, y, and z—into red, green, and blue channels respectively. So, at each location in the image, the color tells us the dominant direction of diffusion, and the brightness shows the degree of anisotropy, or how directional the diffusion is.

For example, a bright red region indicates that diffusion is strongest along the x-direction, green shows diffusion along y, and blue along z. When these combine, you get intermediate colors—just like mixing red, green, and blue light—forming a vivid, color-coded diffusion map.

This type of image is called a fractional anisotropy, or FA map.

It tells us how strongly the diffusion is confined to one particular direction. In regions such as the white matter of the brain, where axonal fibers are highly aligned, the diffusion is strongly anisotropic—so these areas appear bright and vividly colored.

What makes this especially powerful is that we are not only seeing the brain's anatomy, but also its function and connectivity.While CT imaging primarily gives us anatomical information—showing structures, shapes, and densities—MRI can provide functional information, such as how water and blood move inside tissues.

Using diffusion-weighted and diffusion-tensor imaging, we can track how water molecules travel along nerve fibers, revealing the brain's communication pathways. This is one of the reasons MRI is considered both a structural and a functional imaging modality.

And with that understanding of diffusion and perfusion imaging, we are ready to move to the final concept of this section—spectral imaging, which explores the frequency domain information contained in MRI signals.

slide28:

Spectral imaging is like adding another layer of information to MRI. You can think of it as extending the same Fourier-based idea we've been discussing — but now, instead of looking at a single signal per voxel, we analyze multiple frequencies coming from that same voxel. Let's pay attention to the sentence highlighted here: each voxel emits signals at multiple frequencies, rather than just a single frequency.

In conventional MRI, after phase and frequency encoding, every pixel or voxel produces one main signal — one frequency — corresponding to its local environment. But in spectral imaging, or what we often call MR spectroscopy, that same voxel emits several frequencies at once. Why does this happen? Because not all water molecules are in the same chemical environment. Some are bound within proteins, some interact with lipids, and others are part of metabolites. Each of these environments slightly shifts the resonance frequency of the hydrogen nuclei — what we call a chemical shift.

As a result, a single voxel gives rise to multiple peaks — each corresponding to different chemical compounds. For example, in the spectrum shown here, you can see peaks for creatine, choline, glutamate, glutamine, and N-acetylaspartate, or NAA. These molecules tell us about the brain's biochemical composition — for instance, NAA is a marker of healthy neurons, while choline indicates cell membrane turnover. So, spectral imaging gives us access to biochemical information, not just structural or functional detail. By measuring the exact frequencies and amplitudes of these peaks, we can infer what molecules are present and in what quantities.

The next natural question is — how do we obtain these signals from a specific pixel or voxel? There are two main approaches, just like the passive and active methods we discussed earlier in angiography. Both techniques let us excite or isolate specific regions in space to record their unique spectral signatures — and we'll explore those next.

slide29:

Now let's look at one of the two methods used in MRI spectroscopy — the PRESS method, which stands for Point-Resolved Spectroscopy. This method is commonly used to acquire the MR spectrum from a single voxel — that is, a small three-dimensional volume within the body. Here's the idea: you begin by applying a

90-degree RF pulse to select a specific slice of tissue, as shown in the figure. Then, you apply two 180-degree pulses, each combined with a gradient field — one along the x direction and another along the y direction. These three pulses — one 90-degree and two 180-degree — together define a small box-shaped region in space.

Only the nuclei located inside that box experience all three RF pulses, and therefore, only that region contributes to the detected MR signal. Everything outside that voxel is effectively suppressed. Once the voxel is defined, the signal collected from it contains contributions from all the chemical compounds present there — for example, metabolites like creatine, choline, and N-acetylaspartate. When you perform a Fourier transform on that signal, you get a spectrum, which shows the characteristic peaks of these molecules.

So essentially, PRESS isolates a specific voxel in three dimensions using three orthogonal gradients — Gx, Gy, and Gz — and extracts its spectrum. The next method, called chemical shift imaging, extends this idea further. It's a bit more complex but very powerful. Instead of acquiring data from just one voxel, it gathers spectral information from multiple voxels across a slice — creating a spatially resolved chemical map. This is a more advanced, and often more time-consuming, approach, but conceptually, it builds directly on what you see here.

slide30:

Now, let's move to the second spectroscopy method — CSI, which stands for Chemical Shift Imaging. This method extends the idea of PRESS. Instead of selecting a single voxel, CSI allows us to collect spectra from multiple voxels simultaneously, across an entire slice. Here's how it works. First, you select an axial slice in the z-direction using a 90-degree RF pulse combined with a slice-selection gradient, as shown in the figure. Then, phase encoding is applied in the x and y directions — that's what the Gx and Gy gradients represent here.

As a result, the MRI system collects signals from many small voxels arranged in a grid pattern, as you can see on the right. Each voxel now contains its own chemical spectrum, showing peaks from various metabolites, just like the one we saw in PRESS. To reconstruct this data, we perform two-dimensional inverse Fourier transformations with respect to the spatial dimensions x and y, and a forward Fourier transformation with respect to time. This process gives a full spectrum for each voxel, mapping both spatial and chemical information across the slice.

In simpler terms, PRESS gives you the spectrum from one small volume, while CSI gives you a whole map of spectra from many volumes at once. It's computationally heavier and takes longer to acquire, but it provides much richer biochemical information throughout the tissue. This technique is particularly valuable in brain spectroscopy, where you can visualize metabolic differences between normal and abnormal tissues. And just like other MRI techniques, the image quality and contrast can be further improved using specialized contrast agents — typically compounds based on magnetic elements with unpaired electrons, which help enhance local magnetic field effects and improve visibility.

slide31:

Now, let's look at an example of an MRI contrast agent — Magnevist. MRI contrast agents are usually paramagnetic materials, meaning they have unpaired electrons that create local magnetic field variations. These variations enhance the local signal intensity and improve the signal-to-noise ratio, allowing us to

better differentiate between tissues. When a contrast agent is introduced into the body, it tends to accumulate in regions with rich vasculature, such as tumors. This makes those areas appear brighter in the MRI image, helping radiologists identify abnormal growths or tissue changes.

The most common class of MRI contrast agents is based on metal ions such as gadolinium, manganese, or europium. Among these, gadolinium is particularly effective because the ion Gd^{3+} has seven unpaired electrons — the maximum possible — which gives it a very strong magnetic moment. However, gadolinium by itself is toxic, so it cannot be injected directly. Instead, it is safely enclosed in a chemical cage, called a chelate, that binds the gadolinium tightly. One of the most widely used formulations is Gadolinium DTPA, or Gadolinium Diethylenetriaminepentaacetic Acid, which is marketed under the trade name Magnevist.

This compound has been used clinically for many years to enhance MRI scans of the brain, spine, and other organs. It works by creating small magnetic inhomogeneities that shorten relaxation times, giving a strong contrast enhancement in vascular and pathological regions. However, although MRI itself is considered a green and non-ionizing imaging technique — unlike CT, which involves radiation — recent studies have shown that these contrast agents are not entirely harmless. Traces of gadolinium can sometimes remain in the body, leading to concerns about long-term safety, especially for patients who undergo repeated contrast-enhanced MRI scans.

So while MRI remains fundamentally safe, it's important to be aware of potential risks associated with certain contrast materials, and to continue developing safer, more biocompatible alternatives for the future.

slide32:

Here we see an example of how a contrast agent like Gadolinium-DTPA, or Magnevist, works in clinical imaging. On the left is the chemical structure of Gd-DTPA. In this molecule, the gadolinium ion — which has seven unpaired electrons — is enclosed within a stable chelating cage. This cage prevents the toxic metal from freely interacting with biological tissue, while still allowing it to influence nearby water molecules. When water molecules come close to the contrast agent, they experience local magnetic field changes that speed up their T_1 relaxation, producing a brighter signal on the MRI image. Water molecules that directly interact with the free binding site of the gadolinium ion experience what we call inner-sphere relaxation, resulting in a strong signal increase. Those that are nearby but not directly bound experience outer-sphere relaxation, which is weaker but still noticeable.

On the right, you can see a brain MRI taken after administration of Gd-DTPA. Notice the bright region within the brain — this enhancement clearly reveals a tumor. The contrast agent has accumulated in that region because tumors typically have leaky vasculature, allowing the agent to diffuse more freely into the tissue. However, it's important to mention that these contrast agents are not entirely safe. Although gadolinium is held within a chemical cage, studies have shown that residual traces of gadolinium can remain in the human body — particularly in the brain and kidneys — for a long time after injection. This has raised significant safety concerns.

As a result, the European Union has restricted or prohibited the clinical use of several popular gadolinium-based contrast agents. Researchers and manufacturers are now working to design safer alternatives — either agents that can be excreted more efficiently or new imaging methods that require little or no contrast agent at all. So while MRI contrast imaging is extremely powerful, it reminds us that even advanced, "non-ionizing" imaging technologies come with important safety considerations.

slide33:

Now let's talk about MRI safety, especially regarding the use of gadolinium-based contrast agents. Gadolinium, symbol Gd, is one of the 17 rare-earth elements and belongs to the lanthanide series of metals. It's a paramagnetic material, meaning it has unpaired electrons that interact with magnetic fields. This property makes gadolinium extremely useful for enhancing the contrast in MRI scans, helping us visualize blood flow, tissue abnormalities, and tumors more clearly.

However, gadolinium itself is highly toxic to humans. It has no natural biological role and is not found anywhere in the human body under normal conditions. The problem arises because gadolinium ions are similar in size and charge to calcium ions, which are essential for many cellular processes. When free gadolinium ions enter the body, they can interfere with calcium-dependent mechanisms, disrupting normal cell function. To minimize this risk, gadolinium is used in a chelated form, meaning it is locked inside a chemical cage — such as DTPA — to prevent the toxic metal from interacting directly with tissues.

However, studies have shown that not all of the gadolinium is completely excreted from the body after an MRI. Small amounts can remain in the brain, kidneys, and other organs, even months or years later. This residual gadolinium may lead to inflammation, oxidative stress, neurological effects, or even genetic damage at the cellular level. Over the past decade, these concerns have prompted major regulatory changes. The European Medicines Agency has restricted or banned the clinical use of several widely used gadolinium-based contrast agents, particularly those that release gadolinium more easily.

So, while MRI itself is a non-ionizing and generally safe imaging modality, the contrast agents we use must be handled with caution. This is why the field is actively researching safer alternatives — agents that are more stable, biodegradable, or even completely free of gadolinium — to make MRI an even safer tool for medical diagnosis.

slide34:

So, as we've just discussed, MRI contrast agents are not entirely harmless. Research in recent years has revealed that gadolinium-based agents — while extremely useful for improving image clarity — can sometimes remain in the body long after the scan, particularly in the brain and even within bone tissue. In response to this growing concern, the European Medicines Agency, or EMA, launched an extensive investigation. In 2017, their committee reviewed evidence showing that traces of gadolinium could persist in brain tissues for many months after a single injection. MRI scans even showed increased signal intensity in regions where gadolinium had accumulated.

Following this review, the EMA decided to take precautionary action. They suspended several widely used linear gadolinium-based contrast agents, including MultiHance, Omniscan, Magnevist, and OptiMARK. These were once very common in clinical MRI but are now restricted in Europe due to the risk of residual gadolinium retention. This is not something you'll find discussed in most textbooks yet — it's a relatively recent development, but an important one to know. It reminds us that while MRI is a non-ionizing and generally safe imaging technique, it still comes with chemical safety considerations when contrast agents are used.

Now, from here, I'll move quickly through a few additional slides that are not required for this course, but they're worth knowing to give you a broader perspective. They illustrate how MRI, beyond diagnostic

imaging, plays a major role in neuroscience and brain research, helping us study how the human brain functions and connects at a fundamental level.

slide35:

So far, we've covered the main aspects of MRI — the scanner hardware and the pulse sequences. These are the two core components that determine how MRI works and what kind of images we can obtain. Now, the remaining two topics — Brain Initiatives and Novel Ideas — are not required for this course, but I'll go through them briefly because they're very interesting and show how MRI has evolved beyond traditional medical imaging.

MRI is an exceptionally powerful tool, especially in brain research. It's not only used to visualize anatomical structures, like the brain's gray matter and white matter, but also to explore neurological function — how different regions of the brain communicate and interact. Researchers around the world, through major brain initiatives, are using advanced MRI techniques to map brain connectivity, monitor neural activity, and understand complex cognitive processes such as memory, decision-making, and emotion.

In this way, MRI has become one of the most important technologies for both clinical diagnosis and fundamental neuroscience research.

slide36:

This has become one of the hottest topics in science — not just in the United States, but also in Europe, China, and Japan. Governments around the world are investing heavily to unlock one of the greatest mysteries of all time: how the human brain works and how intelligence emerges. MRI plays a central role in this global effort because it allows researchers to look inside the living brain — not only to see its structure, but also to monitor its activity in real time.

Here you can see the timeline of the U.S. BRAIN Initiative, which was announced in 2013 by President Obama. The initiative brought together major agencies such as the NIH, NSF, DARPA, and later FDA and IARPA. Each agency contributes in its own way — from developing new tools and technologies, to studying cognition, rehabilitation, and neurological regulation. The research priorities include identifying brain cell types, mapping neural circuits, developing technologies to monitor brain activity, and creating precise intervention tools.

The ultimate goal is to integrate all of these approaches to achieve a comprehensive understanding of how the brain functions — from the molecular level to behavior. MRI is a key part of this vision, serving as one of the most powerful noninvasive tools for studying the human brain.

slide37:

Now let's look at this fascinating experiment. Here, subjects were shown two different types of images — human faces and buildings. The MRI scans you see here show how the brain reacts differently in each case. When a person looks at a face, certain brain regions light up — you can see those areas in red and orange. When the same person looks at a building, other regions become more active. This happens because the brain's oxygen-rich blood flow changes depending on what you are thinking or perceiving.

When neurons in a region start firing, they consume more oxygen. The body responds by sending in oxygenated blood, which creates a stronger MRI signal in that area. This mechanism is called the BOLD effect, short for Blood Oxygen Level Dependent contrast. By analyzing these signal patterns statistically, we can tell what kind of object the person is viewing — for example, a face versus a house. This was one of the early experiments demonstrating how brain activity reflects thought and perception.

In more recent studies, researchers have gone even further. Using advanced image analysis and machine learning, they can now reconstruct, with rough resolution, what a person is actually seeing — essentially allowing us to peek into human perception. It's an incredible step toward understanding how the brain encodes visual information. If you have time, I highly recommend watching a short video on this topic — it's truly amazing to see how far this technology has come.

slide38:

This short video, which I highly recommend you watch later, is truly fascinating. It explores how modern MRI technology can be used to study thought processes — in other words, how it can "read the mind."

In this demonstration, researchers show that by monitoring brain activity patterns, we can detect what a person is recognizing or recalling, even without them speaking. For instance, imagine a criminal suspect claiming, "I didn't commit the crime." If you place this person in an MRI scanner and show images of various locations — one being the actual crime scene — the brain's response will be noticeably different when that familiar location appears.

The brain's recognition centers light up, revealing subconscious memory traces, even when the person tries to hide it. This shows the incredible potential of functional MRI — not just for medical diagnosis, but also for applications in neuroscience, psychology, and even criminal investigation.

The video runs about ten minutes, and it's well worth watching. It gives you a glimpse into how MRI can connect biology, cognition, and technology — opening up entirely new frontiers for understanding the human mind.

slide39:

So far, we've learned about the major medical imaging modalities — CT, nuclear imaging, and MRI. Each of them has unique strengths and physical foundations. CT uses X-rays and attenuation; nuclear imaging uses radioactive tracers; and MRI uses magnetic resonance and the behavior of hydrogen protons.

Now, one very exciting direction in modern imaging research is modality fusion — combining these different imaging technologies into a single, integrated platform. For example, PET/CT and PET/MRI scanners combine anatomical and functional information together, so that we can see both the structure and the physiology in one unified image.

Beyond this, people are exploring hybrid imaging ideas, where optical imaging, ultrasound, or even photoacoustic methods are combined with MRI or CT to capture complementary information. These approaches can greatly enhance diagnostic accuracy and help us understand the body's processes at multiple scales — from cellular activity all the way to organ-level structure.

So, while MRI itself is a powerful and mature modality, the fusion of imaging technologies is opening up new horizons — what we might call the “next generation” of biomedical imaging systems. These are truly exciting times for both clinical medicine and research.

slide40:

So here we come to one of the most exciting directions — imaging modality fusion.

In this triangle, you see the three major imaging technologies: CT, MRI, and PET or SPECT. Each of them has unique strengths — CT provides excellent structural detail and shows calcification and bone; MRI gives superb soft-tissue contrast and functional information; and PET or SPECT provides metabolic and molecular insights.

In recent years, CT has already been combined with PET or SPECT, and MRI has also been paired with PET in what we call PET/MRI scanners — Siemens, for example, developed a system called the MRI Pattern Scanner to take advantage of this synergy.

But the combination of CT and MRI remains one of the most challenging yet most promising directions. In cardiac imaging, for example, CT clearly shows the coronary vasculature and calcifications, while MRI reveals the soft tissue and functional aspects of the heart. Together, they can provide a comprehensive picture of both structure and physiology.

Ultimately, our goal — and the vision our group has been promoting — is to bring all imaging modalities together in one integrated system, so that we can visualize anatomy, function, and molecular information simultaneously. That would be the true “all-in-one” scanner of the future — the complete fusion of CT, MRI, and nuclear imaging.

slide41:

And this idea leads us to the next concept — what we call Omni-tomography.

Omni-tomography means all-in-one, all-at-once, and all-of-couplings. In other words, we want to bring all imaging modalities — CT, nuclear imaging, and MRI — together into a single integrated system.

Instead of acquiring CT data first, then performing a nuclear scan, and finally running an MRI, the vision is to acquire all the information simultaneously, within one coordinated scan. Imagine a scanner capable of capturing structural, functional, and molecular information — all at the same time, perfectly aligned in space and synchronized in time.

This concept was proposed and discussed in a perspective paper titled “Towards Omni-Tomography — Grand Fusion of Multiple Modalities for Simultaneous Interior Tomography.” The paper outlines both the motivation and the technical strategies for making this vision possible.

The idea is ambitious — it represents the ultimate form of imaging fusion — but it’s also the natural next step in medical imaging, combining everything we’ve learned from CT, MRI, and nuclear imaging into one unified platform.

slide42:

To move toward this vision of omni-tomography, we began by building a toy system — a small prototype that integrates all the scanners together in one setup.

Here you can see a conceptual model where CT, PET, SPECT, and MRI components are arranged concentrically around a common patient table. The idea is to allow all the modalities to operate within the same geometry so that they can collect data simultaneously from the same region of interest.

In this design, the PET ring and SPECT detectors are positioned together with an X-ray tube and detector pair for CT imaging, while a specially designed open MRI magnet provides the magnetic field without blocking the other modalities. This required careful electromagnetic shielding and mechanical design to ensure that each subsystem could operate without interference.

After developing this integrated model, we focused on building a functional MRI subsystem, completing its design as a key step toward realizing a fully combined, all-in-one scanner. This system demonstrates the feasibility of truly simultaneous, multimodal imaging — a step closer to our long-term goal of achieving comprehensive, unified tomography for both research and clinical applications.

slide43:

This slide shows the next step — how we actually designed and integrated a CT–MRI scanner. The idea is to combine the strengths of both modalities within a single system.

Here you can see the components that make this possible. On the left, we have the CT subsystem, including the X-ray source and detector arranged in a compact, quasi-stationary ring. On the right are the MRI components, with magnetic coils and gradient coils designed to fit around the same imaging volume. The two systems operate together — CT providing high-resolution structural information, while MRI provides soft-tissue contrast and functional details.

All these elements come together in the integrated setup shown at the bottom. This is a simultaneous CT–MRI scanner, designed so that both modalities can collect data at the same time without interference.

For those of you who are interested in learning more, you can look up our paper published in Medical Physics titled “Vision 20/20: Simultaneous CT–MRI — The Next Chapter of Multimodality Imaging.”

In addition to this high-end system, we are also exploring a low-cost version, aiming to make hybrid imaging more accessible for both research and educational purposes.

slide44:

Here, we move toward a more practical version of hybrid imaging — a low-cost system that combines MRI and X-ray together. Instead of using a superconducting magnet as in a high-end MRI system, this design uses a permanent magnetic field, making it much simpler and more compact.

The concept is quite elegant. For each rotation, you obtain both an MRI projection and an X-ray projection at the same angle. By rotating the system around the patient, you can reconstruct the images together — effectively achieving CT–MRI fusion in a single setup.

This configuration opens many possibilities for future imaging systems. You can use MRI to guide X-ray imaging, combine anatomical and functional information in real time, and even explore new imaging modes for clinical or research applications.

So, while high-performance simultaneous CT–MRI scanners represent the top tier of multimodal imaging, this hybrid spiral MR–X-ray system offers a more accessible pathway to achieve similar goals — bringing advanced imaging capabilities within reach for smaller labs or specialized applications.

slide45:

When a patient undergoes radiation therapy for cancer, a highly energetic beam is directed toward the tumor to destroy cancerous cells. However, the accuracy of this beam is critical — we need to know precisely where the tumor is during treatment, and that's where image guidance becomes essential.

What you see here is a new concept that integrates interior MRI with a medical linear accelerator. The idea is to use MRI imaging to continuously visualize a small region of interest — the area where the tumor resides — while radiation therapy is being delivered. This allows us to target the tumor with submillimeter precision and adjust in real time if the tumor moves due to breathing or other physiological motion.

This work is being developed in collaboration with colleagues at UT Southwestern, where we are exploring how this interior MRI can guide radiation delivery and make cancer therapy safer and more effective.

slide46:

Building upon that idea, this slide shows our interior tomography approach for MRI-guided radiation therapy.

Here, the MRI system is designed not to image the entire body but to focus only on a specific region of interest, such as the tumor area. This design uses superconducting coils and gradient coils arranged in a compact geometry that allows radiation beams to pass through while maintaining the ability to generate precise MRI images of the target zone.

This compact "interior MRI" setup makes it possible to integrate the MRI system directly into a linear accelerator — the same device used for delivering radiation therapy. With this configuration, we can track the tumor in real time and adjust the beam accordingly, ensuring that radiation precisely targets the cancer while sparing surrounding healthy tissue.

Along this same direction, our latest research—published last year—introduces an even more advanced concept based on polarized radio tracers, which bridges the gap between traditional gamma-ray imaging and MRI-based methods.

slide47:

Previously, we learned about gamma-ray imaging and radio tracers, which allow us to visualize physiological functions inside the body by detecting radiation emitted from injected tracers. Now, in this work, we take the concept one step further by combining emission and transmission tomography into a single, simultaneous process—what we call Simultaneous Emission–Transmission Tomography, or SET.

This idea was presented at the Fully3D Conference in Xi'an, China, and it represents an important advance in multimodal imaging. Normally, in PET or SPECT, emission data and transmission data are collected separately—first one scan for anatomy and then another for function. But with SET, both types of data are acquired at the same time.

By combining emission and transmission information simultaneously, we can achieve more accurate reconstructions with improved spatial and contrast resolution. The images on the left show the structural information from the transmission data, while the ones on the right display emission activity from the radio tracer. When we fuse these two datasets, we obtain both anatomical and functional information in a single scan, reducing motion artifacts and improving quantitative accuracy.

This represents another step toward our broader goal—true multimodal imaging, where complementary information from different imaging mechanisms is acquired concurrently for the most comprehensive view of the human body.

slide48:

Now, let's look at the motivation for Simultaneous Emission–Transmission Tomography, or SET.This idea comes from a concept called Polarized Nuclear Imaging, abbreviated as PNI, which builds on the magnetic resonance (MR) framework.

In conventional nuclear imaging, gamma-ray photons are emitted randomly in all directions. However, with a polarized radio tracer, the emission becomes directionally dependent. That means the gamma rays are no longer isotropic — they are preferentially emitted along specific directions.

As shown in the figure, when nuclei are fully polarized along a magnetic field, the emitted gamma rays have a strong preference for the horizontal direction, and almost none are emitted vertically. By using RF pulses and magnetic field manipulations, similar to those in MRI, we can flip this polarization vector — sometimes called the M vector — to control the emission direction.

In other words, by flipping the spins with an RF pulse, we can switch from collecting gamma rays along the x–y plane to collecting them along the z direction. This enables us to encode spatial information into the atomic spin orientation, essentially merging nuclear imaging and magnetic resonance mechanisms.

The result is that polarized nuclear imaging provides much higher sensitivity than conventional MRI, since the signal is derived directly from gamma emission rather than from weak magnetic resonance. The only drawback is that it does not account for attenuation, which is usually handled by CT.

Our recent paper builds on this idea — showing that by using a polarized tracer, we can reconstruct both the radioactive tracer concentration and the attenuation background simultaneously, combining what MRI and CT each do best into a single integrated framework.

slide49:

With this polarized nuclear imaging framework, we can perform both emission and transmission tomography simultaneously.However, to make this practical, we need suitable polarized radioisotopes, or tracers.

This table lists a few existing candidates — isotopes of xenon and krypton — that can, in principle, be polarized and used for such imaging.For example, xenon-131m has a half-life of about twelve days and a gamma emission energy around one hundred sixty-four kilo–electron volts.Xenon-127m decays faster — with a half-life of sixty-nine seconds, emitting gamma rays between one hundred twenty-five and one hundred seventy kilo–electron volts.

And krypton-79m has a half-life of about fifty seconds, with gamma emission near one hundred thirty kilo–electron volts.

You can see the challenge here: ideally, a good clinical tracer should have a half-life around thirty minutes, long enough for preparation and imaging, but short enough to minimize radiation dose. Unfortunately, none of these isotopes fit that range — they are either too long or too short.

That's why our team has been collaborating with chemists and nuclear physicists to identify or synthesize new polarizable tracers with more suitable half-lives and emission properties.

If such tracers can be developed, they would make simultaneous emission–transmission tomography a realistic and powerful imaging tool for the future.

slide50:

(No notes for this slide)

slide51:

(No notes for this slide)

slide52:

Now, let's look at how all of this comes together in the system design.In this setup, you can selectively activate or turn on any voxel or pixel by controlling the magnetic gradients and RF pulses we just discussed.

Each polarized nucleus—such as xenon or krypton—can be individually manipulated using precise gradient fields in the x, y, and z directions. The coils you see here generate those magnetic field gradients, while the surrounding detectors collect the emitted gamma photons.

So, depending on which gradient and RF pulse combination you apply, you can excite a specific spatial region, measure its response, and reconstruct the image point by point.

This concept essentially merges nuclear emission detection with MRI-style spin control, enabling us to perform polarized nuclear imaging in a controlled and highly selective way. It's a very elegant system, because we can combine magnetic resonance encoding with nuclear emission detection—bringing together the best of both imaging worlds.

slide53:

Now let's talk about the data model behind this concept. As I mentioned earlier, each polarized tracer behaves like a tiny donut that can be flipped in any direction. If the polarization is not aligned with the detector, you won't detect any signal. But once you flip it to the right orientation, that voxel — or pixel — becomes active, and we can measure it. In this model, for each voxel we have two unknowns: the local attenuation coefficient, which we call "mu of r," and the local tracer concentration, which we call "lambda of r."

To estimate these two unknowns, we make two measurements — one from the left and one from the right — denoted as "m one of r" and "m two of r." These two measurements are expressed as: m one of r equals phi of r times the exponential of negative integral from minus infinity to r of mu of r d r, and m two of r equals phi of r times the exponential of negative integral from r to infinity of mu of r d r. Here, phi of r represents the gamma-ray flux, and mu of r is the attenuation coefficient along the ray path.

If we take the logarithm of the ratio "m one of r over m two of r," we obtain an expression involving the difference of two integrals of mu. By differentiating this relationship, we can recover mu of r directly as negative one-half times the derivative with respect to r of the logarithm of "m one over m two." Once mu of r is known, we can calculate phi of r, which equals the exponential of the integral of mu of r times m one of r, or equivalently, the exponential of the integral of mu of r times m two of r, depending on direction.

In simple terms, this means the gamma flux escaping from each side of an activated voxel is proportional to the tracer concentration at that point. So, with just two directional measurements per voxel, we can reconstruct both the attenuation coefficient and the tracer concentration — achieving simultaneous emission and transmission tomography in a single scan. That's the key idea behind SET, or Simultaneous Emission–Transmission Tomography.

slide54:

Now, let's look at how this approach can be extended to whole-body imaging. When we use a polarized radio tracer, we can perform both emission and transmission tomography within the same MRI framework. Because the tracer is polarized, it can emit gamma rays directionally — and since the polarization can be flipped, we can measure signals from multiple directions without rotating any hardware.

This setup naturally integrates with MRI. So not only can we collect MRI signals, but at the same time, we can measure attenuation, just like in CT, and tracer concentration, as in nuclear imaging. In other words, we can combine the strengths of all three modalities — CT gives us attenuation information, MRI gives us soft-tissue and molecular contrast, and nuclear imaging provides tracer distribution — all in one coordinated system.

Here, you see a concept for whole-body imaging. The patient lies inside a large cylindrical magnet, and the radio tracer inside the body is flipped periodically. Each time the tracer is flipped, gamma rays are emitted in opposite directions. These emissions are detected from both sides, allowing us to reconstruct a complete image of the entire body. A key advantage of this method is that the collimation — or the directional selection of gamma photons — is achieved magnetically. That means we no longer need a mechanical collimator, which normally blocks most photons and wastes energy.

Magnetic collimation captures more of the emitted photons, improving sensitivity and reducing noise. So this design minimizes photon waste and integrates the best parts of CT, MRI, and nuclear imaging — offering a very promising pathway for future multi-modality imaging systems.

slide55:

For your homework, I'd like you to draw diagrams for all the pulse sequences we discussed in this lecture. That includes spin echo, gradient echo, diffusion-weighted imaging, and spectroscopy sequences. Along with each diagram, please explain the working principle — how the sequence works and what kind of information it provides.

Also, please watch the YouTube video I mentioned earlier; it will help you better visualize how MRI can even be used for mind-reading research. Try to connect what you see in the video and what we discussed today about MRI signal generation and brain imaging. Then, take some time to read the textbook section on MRI so you're well prepared for our next lecture on ultrasound imaging.

Thank you, and I'll see you in the next class.