slide1:

Today, we're going to begin a new chapter on the physical principles of Magnetic Resonance Imaging, or MRI. This is Lecture 18 in our series.

MRI is one of the most powerful tools in modern biomedical imaging, and in this session, we'll look at the physics closely that makes MRI work. We'll build up from the physical foundations including Maxwell's equations, through the concepts of magnetization and precession, all the way to how signals are generated and detected. Along the way, we'll also discuss key parameters, such as proton density, T1 and T2, and how they enable image contrasts.

Okay, let's get started with MRI physics.

slide2:

So, we are still on schedule. If you look at the course outline, we have already covered topics such as Fourier series, signal processing, CT reconstruction, and nuclear imaging.

Now, we are moving into MRI materials. Today's focus is on MRI physics, which is the foundation for understanding how MRI systems generate and detect signals. This will prepare us for the next lectures, where we will go deeper into MRI techniques, systems, and applications.

slide3:

Let me start by giving you both a preview and a review. The preview is like a big picture of what MRI physics is about. I will provide you with some basic knowledge and introductory ideas that will help you understand the more detailed content later.

So first, let's take a look at this overall framework. We will begin with the physical foundation, then move on to signal generation, and finally discuss signal decay. This roadmap will guide us through today's lecture.

slide4:

First, let's step back and look at the big picture. Up to this point, we have studied CT and nuclear imaging. These are already used together in hybrid imaging. CT provides us with detailed structural and anatomical information, while nuclear imaging methods, such as PET and SPECT, provide functional information by tracking radioactive tracers that participate in biochemical reactions within the body.

Because these two types of information are highly complementary, combining them makes the results much more powerful. That is why PET/CT scanners have become standard in many hospitals, especially in radiation oncology. More recently, companies like Siemens and GE have developed PET/MRI systems. However, CT and MRI have not yet been commercially combined, although research is moving in that direction.

The long-term vision is to bring CT, MRI, and nuclear imaging together into one unified system, giving us structural, functional, and biochemical information at the same time. This is the major trend in imaging—moving from individual modalities to hybrid and, eventually, fully integrated scanners.

Now, with that context, let us move to our next imaging modality, MRI, which is unique in that it can provide both anatomical and functional information in a single technique.

slide5:

Now, let's go through some important milestones. MRI is unique because it can provide two types of information. On one hand, it gives us anatomical information, very much like CT, by showing the structure of soft tissues. On the other hand, it can measure blood oxygenation levels, which adds functional information. That combination is what makes MRI so powerful.

The story begins in 1946, when Bloch at Stanford and Purcell at Harvard discovered nuclear magnetic resonance. For this work, they received the Nobel Prize in Physics in 1952. At that time, NMR was mainly used to measure signals from a whole sample.

Then in 1973, Lauterbur at Stony Brook University introduced the idea of magnetic resonance imaging. The key innovation was to use gradient magnetic fields, which made it possible to turn NMR from a bulk measurement into a tomographic imaging technique. Now we could form images with pixels and voxels instead of just an overall signal. This breakthrough eventually earned Lauterbur and colleagues the Nobel Prize in Medicine in 2003.

By the late 1970s, the first human MRI images were produced. In the early 1980s, commercial MRI systems became available. Then, in 1993, functional MRI was developed, allowing us to look at brain activity through blood oxygenation changes.

And now, we are in the era of large-scale brain initiatives, where MRI continues to be at the center of neuroscience and clinical imaging.

slide6:

Here is the very general idea behind how MRI works. At the center of the system, we have a large superconducting magnet. This magnet creates a very strong, stable magnetic field. Around it, we place gradient coils, which allow us to vary the magnetic field in different directions and make tomographic, or slice-by-slice, imaging possible.

In addition, we use radiofrequency coils, or RF coils. These coils send radio waves to disturb the alignment of the protons in the body. When the protons are tipped away from alignment, they begin to precess, and in doing so, they emit their own radiofrequency signals. The RF coils then detect these signals.

With the help of signal processing and mathematical reconstruction, we turn these detected signals into detailed images. Right now, it may sound abstract, but as we go deeper into the lecture, we will carefully unfold each of these components. For now, just keep in mind the big picture: a strong magnet, gradient coils for spatial encoding, RF coils for excitation and detection, and signal reconstruction to form the final images.

slide7:

To get a rough idea of how MRI works, let's look at it step by step.

Magnetic resonance imaging—MRI for short—can be modeled in the following way. First, consider this conducting loop of wire. The magnetic dipole moment is defined as the current going through the wire, times the cross-sectional area enclosed by the wire, times a unit vector perpendicular to that surface area. So the magnetic dipole moment points in the perpendicular direction, like this.

Now, a proton is believed to spin on its axis at the subatomic level. Because it's spinning, the magnetic dipole moment can be thought of as the sum of many tiny circulating currents. Whenever you take a dipole moment and place it in an external magnetic field, the dipole experiences a torque. That torque is given by the equation: tau equals mu cross B. In words: the torque equals the magnetic dipole moment crossed with the magnetic field.

What this means is that the dipole moment will tend to align with the magnetic field. Imagine here is the B-field, and here is the dipole moment. When you take the cross product, you get a torque out of the board. That torque causes the dipole to rotate in this direction. The dipole will try to line up with the external B-field, but it won't line up completely. The angular momentum of the spinning top—so to speak—prevents complete alignment. Instead, it undergoes precession, just like a spinning top wobbling around a vertical axis.

So, think of the proton as a spinning top. It rotates on its axis, but also precesses about an axis that is parallel to the external magnetic field.

Now, let's introduce radiation. We emit an electromagnetic wave laterally, toward the precession axis. For simplicity, I've drawn only the magnetic component, oscillating in and out of the board as it travels laterally. If the angular frequency of this radiation equals the angular frequency of precession, resonance occurs.

We define the Larmor angular frequency:omega L equals gamma times B-naught.

Here gamma is the gyromagnetic ratio. For a proton, gamma is about two-point-sixty-seven times ten to the eighth radians per second per tesla. To convert from radians per second to hertz, we divide by two-pi. So,

two-pi f L equals gamma times B-naught.

Or equivalently:f L equals gamma over two-pi, times B-naught.

Numerically, this works out to about forty-two megahertz per tesla.

So, the frequency of the applied radiation matches the precession frequency of the proton. Whenever the proton comes around to a certain point, the applied B-field is right there to interact with it. The torque tends to flatten the precession, step by step, as the proton keeps rotating.

When the radiation is turned off, the proton flips back to its original motion, precessing around the external magnetic field axis. This is the fundamental mechanism of magnetic resonance. Historically, it was called nuclear magnetic resonance, but later changed to simply magnetic resonance—mostly for public perception, since "nuclear" tends to worry people.

Now let's take a group of hydrogen atoms. The body contains a large amount of hydrogen, because we are mostly water—H2O. Initially, the dipole moments of hydrogen protons are randomly oriented. When we apply an external magnetic field, the protons tend to align with the field, precessing around its axis.

When we emit radiation laterally, as we discussed before, the protons flatten out. Then, when the radiation is turned off, the protons flip back. Think of this like a spring being released. As they flip back, they emit radiation. That emitted radiation can be detected and processed to form an image.

In 1973, a breakthrough came from Paul Lauterbur—he showed that you could actually use this principle to create images. The idea was: take a subject, place them inside a tube with a strong superconducting magnet, on the order of five tesla. The magnetic field is aligned along the axis of the tube, so all the protons line up and precess.

Then, apply a secondary magnetic field that produces a gradient along the axis of the tube. That means the magnetic field is slightly stronger at one end, weaker at the other. Because the Larmor frequency depends on the field strength, different positions along the tube correspond to different precession frequencies.

For example, at a position one-point-three meters down the tube, the magnetic field might be B1, slightly larger than B0. Substituting B1 into the equation, you find a different resonance frequency. By tuning the applied radiation frequency, only protons at that slice will resonate and emit radiation. This gives spatial localization.

Engineers extended this idea. By using different gradient coils—like the saddle coil shown here—they can create gradients not only along the Z-axis, but also in the X and Y directions. That way, we can localize signals to small cubes of tissue, called voxels, inside the body.

Finally, only tissues containing hydrogen—meaning water and fat—produce signal. That's why MRI works so well for imaging soft tissues in the human body.

slide8:

So that gives us a general idea of how MRI works. Areas in the body that contain a lot of water, like soft tissues, produce strong signals, which show up bright in the image. Areas with little water, like bone, produce weak signals, so they appear dark. That is the basic contrast mechanism in MRI.

The short lecture we just saw is an excellent example of how to present these key principles clearly. I encourage you to revisit it after you study the details, because once you learn the full physics of MRI, you will be able to understand nearly all of it.

Now, let's review some of the mathematical foundations. One important concept is the dot product. If we have two vectors, v and w, their dot product is equal to the magnitude of v times the magnitude of w times the cosine of the angle between them. Geometrically, the dot product measures how much one vector projects onto another. You can think of it as relating to the area formed by the parallelogram between the two vectors.

This is a very elegant mathematical operation, and it also has a counterpart called the cross product, which we will look at next.

slide9:

Now let's look at the cross product. When you take two vectors, say vector a and vector b, their cross product is not a scalar like the dot product, but instead a new vector. The magnitude of this vector is equal to the length of a times the length of b times the sine of the angle between them. The direction of the result is perpendicular to the plane formed by a and b, determined by the right-hand rule.

So while the dot product gives us a projection, the cross product gives us an orthogonal vector that carries both magnitude and direction. This operation is not just a mathematical trick—it has deep physical meaning.

In MRI physics, cross products appear naturally when we describe torques, angular momentum, and precession, making them essential to understanding the behavior of magnetic moments in an external field.

slide10:

Here we define torque and see how it relates to angular momentum. Linear momentum is straightforward: it is mass times velocity. According to Newton's second law, the time derivative of linear momentum is force. That means if you want to change momentum, you need to apply a force, which changes velocity.

Now, when we move from linear motion to rotational motion, we deal with angular momentum. Angular momentum is defined as the cross product of the position vector r and the linear momentum p. So, L equals r cross p.

To change angular momentum, we need torque. Torque is also defined through a cross product: r cross F, the position vector crossed with the applied force. Just as force is the reason linear momentum changes, torque is the reason angular momentum changes.

This is really just an extension of Newton's second law into rotational dynamics. The derivative of angular momentum with respect to time equals torque. In other words, torque is the rotational analog of force.

This connection shows why the cross product is essential. It is not only a mathematical operation but also the key to describing how forces lead to changes in rotation. This principle is central in understanding precession and magnetization, which we will soon connect back to MRI physics.

slide11:

Let's quickly review some high school physics to set the stage. Electric charges interact through forces. If you place two charges in space, the nature of those charges determines how they behave. Like charges, meaning both positive or both negative, repel each other. Opposite charges, one positive and one negative, attract each other.

The strength of this interaction is given by Coulomb's law. The force is proportional to the product of the two charges, divided by the square of the distance between them. So the closer the charges, the stronger the force.

This simple idea about electric forces will be important when we connect electricity and magnetism, and eventually link these basic interactions to the physics of MRI.

slide12:

Here we see the behavior of magnetic poles. Just like electric charges, magnets also interact depending on their polarity. Every magnet has two poles: a north pole and a south pole. Like poles repel each other, while opposite poles attract.

This is a simple but powerful analogy. You can think of it like teamwork: if two people have the same role, they may push against each other, but if they have complementary skills, they come together to form a stronger team. In the same way, the north and south poles naturally attract and stabilize one another.

These basic principles of magnetism, combined with the behavior of electric charges, are the foundation for understanding the electromagnetic interactions that drive MRI physics.

slide13:

The Earth itself has a magnetic field, but it is actually quite weak. The Earth's field is about half a Gauss. To put this into perspective, one Tesla equals ten thousand Gauss. So the Earth's field is only about fifty microtesla, which is tiny compared to what we use in MRI.

MRI scanners rely on magnets that are tens of thousands of times stronger. Clinical MRI systems typically operate at one and a half Tesla, three Tesla, or even higher in research settings. These strong fields are essential for aligning enough protons in the body to generate measurable signals.

So while the Earth's magnetic field is important for navigation and protecting us from solar radiation, it is far too weak for imaging. For MRI, we need very powerful superconducting magnets. And now, let's connect electric and magnetic forces together, because in physics, they are two sides of the same coin.

slide14:

Electricity and magnetism are deeply connected. Whenever an electric current flows through a wire, it generates a magnetic field around that wire. If you coil the wire into loops, the magnetic fields from each loop add together, creating a much stronger overall magnetic field.

This is the principle behind an electromagnet. By sending current through a coil, we can generate powerful magnetic fields, and the more loops of wire we use, the stronger the field becomes.

This connection between electricity and magnetism is fundamental to MRI. The giant superconducting magnets used in MRI are essentially coils carrying current, producing extremely strong and stable magnetic fields needed for imaging.

slide15:

So far, we have seen that an electric current can generate a magnetic field. But the reverse is also true: a changing magnetic field can generate an electric current.

Here's how it works. Imagine you have a conducting loop. If nothing changes in the magnetic field passing through the loop, no current flows. But if you insert a magnet into the loop, or pull it out, the magnetic field inside the loop changes. This change induces a current in the loop.

This is Faraday's law of electromagnetic induction. The induced current always acts to oppose the change that caused it, a principle known as Lenz's law. So if you try to increase the magnetic field, the induced current will create a field that resists that increase. If you decrease the magnetic field, the induced current will create a field that resists the decrease.

This interplay between electricity and magnetism is fundamental for MRI. The scanner excites protons using magnetic and radiofrequency fields, and then relies on induced currents in coils to detect the returning signals.

slide16:

Now, let's bring everything together with Maxwell. Alongside Newton and Einstein, James Clerk Maxwell is regarded as one of the greatest physicists. His major contribution was formulating the four Maxwell equations, which elegantly describe how electric and magnetic fields behave and interact.

I don't expect you to master these equations right now, but it is important to recognize their role. They unite everything we have talked about: charges producing electric fields, magnetic fields produced by currents, and the deep connection between electricity and magnetism.

In these equations, the dot product is used to describe divergence—how field lines flow in or out of a point. The cross product is used to describe curl—how field lines twist or circulate. With just these two mathematical operations, Maxwell captured the full behavior of electromagnetic fields.

So at a basic level, remember this: Maxwell's equations are the foundation of all electromagnetic interactions. They are the framework that underlies MRI physics and many other technologies in modern science and engineering.

slide17:

Here we come to the Lorentz force, a very fundamental concept. The equation says that the total force on a charged particle is equal to the electric force plus the magnetic force. Mathematically, F equals q times E plus q times v cross B. Here, q is the charge, E is the electric field, v is the velocity of the particle, and B is the magnetic field.

This means that when an electric charge moves through a magnetic field, it experiences a force. The direction of that force is given by the cross product, which makes it perpendicular to both the velocity and the magnetic field.

At a deeper level, the Lorentz force can actually be derived from Maxwell's equations, showing that Maxwell's framework is completely self-contained. Together, these equations govern all of electromagnetism: electric and magnetic fields, waves, and interactions.

For us, the important point is to become familiar with these concepts—Maxwell's equations, electromagnetic interactions, and the Lorentz force—because they are the physical foundation for MRI signal generation and detection.

slide18:

That wraps up the first part of our discussion. Next, we will move more systematically, following closely along with the structure in your textbook.

In the first part, we covered some foundational physics: angular momentum, magnetic moments from a quantum mechanical perspective, magnetization in the classical sense, and the fascinating but at first mysterious idea of precession. These are the building blocks.

The magnetic moment represents the behavior of individual protons at the quantum level. Magnetization describes how many of these tiny moments add up into a collective vector that we can measure in bulk. And precession explains how that vector behaves in an external magnetic field.

Together, these ideas form the physical foundation that allows us to understand signal generation, relaxation, and ultimately the imaging sequences that make MRI possible.

slide19:

Let’s quickly revisit some high school chemistry to ground ourselves. Here is the basic atomic structure. In the outer shell, we have electrons moving around. In reality, their positions are described by probabilities, but we often picture them as orbiting the nucleus.

Inside the nucleus, we find two main types of particles: protons and neutrons. Protons carry positive charges, while neutrons are electrically neutral. Since all the protons are positively charged, they would naturally repel each other. But they are held together by the strong nuclear force, which balances out the electromagnetic repulsion and creates a stable nucleus.

This balance of forces is what allows atoms to exist in a stable form. And it is precisely these nuclear properties—especially the behavior of protons—that are central to MRI, because the hydrogen nucleus, which is just a single proton, is the main source of signal in magnetic resonance imaging.

slide20:

Now let’s focus on the proton itself. A proton is a positively charged particle, and one of its most important properties is spin. You can think of spin as being similar to a tiny top rotating on its axis. Because the proton is both spinning and charged, it behaves like a little current loop. And any current loop produces a magnetic field.

This is why we say a spinning proton has a magnetic moment. You can picture it as a miniature bar magnet, with a north pole and a south pole.

If we take just one proton—say, from a hydrogen atom in a water molecule—it acts like a tiny magnet. But in the body, we have countless protons, and in the absence of an external magnetic field, their magnetic moments point in random directions. As a result, they cancel out, and there is no overall magnetic effect.

However, when we apply a strong external magnetic field, something very interesting happens. These magnetic moments begin to align, and that alignment is what

MRI relies on generating signals. We will come to that point soon.

slide21:

Here we arrive at two key concepts you must know: angular momentum, usually written as P, and the magnetic moment, written as miu.

From mechanics, angular momentum describes the rotation of a particle—in this case, the proton. You can think of it as a spinning top. Just as torque can change the angular momentum of a top, external influences can change the angular momentum of a proton.

But unlike a simple spinning top, the proton also carries an electric charge. Because of this, it has not only mechanical properties but also electromagnetic ones. Its magnetic behavior is summarized by the magnetic moment miu.

The beautiful part is that these two quantities—angular momentum and magnetic moment—are directly linked. The relationship is miu equals gamma times P, where gamma is a constant called the gyromagnetic ratio.

Quantum mechanics adds one more layer: angular momentum cannot take on any arbitrary value but is quantized. It is expressed in terms of the spin quantum number, I. For a proton, I equals one-half, so the angular momentum follows a specific formula.

The main takeaway is this: the proton's mechanical rotation and its electromagnetic behavior are tied together in a very simple and elegant way. This link forms the basis of how MRI exploits proton spins to generate signals.

slide22:

When a proton with a magnetic moment is placed in a strong external magnetic field, its behavior becomes quantized. Classically, you might imagine that the magnetic moment could point in any direction. But quantum mechanics tells us that only certain orientations are allowed.

Specifically, we focus on the longitudinal component of the magnetic moment, meaning the part aligned with the external field. For the proton, this component can take on only two values: one parallel to the field and one antiparallel. These values are determined by the nuclear magnetic quantum number, m I, which for a proton can be plus one-half or minus one-half.

So, the z-component of the magnetic moment, miu z, equals plus or minus gamma times Planck's constant divided by four pi. This means the magnetic moment is never perfectly aligned—it is always tilted slightly. But it is restricted to these two possible states.

This quantum restriction, the fact that the magnetic moment can only take on discrete orientations, is what gives rise to the fundamental energy levels we use in MRI.

slide23:

When we place spinning protons, each with a magnetic moment, into a strong external magnetic field, their energy levels split. This phenomenon is called the Zeeman effect.

The magnetic moment of a proton can only align in one of two possible ways with respect to the external field: parallel or antiparallel. In the parallel case, the magnetic moment lines up with the field, and this state has lower energy. In the antiparallel case, the magnetic moment points against the field, and this state has higher energy.

So, instead of a continuous range of orientations, quantum mechanics restricts protons to just these two discrete states. The energy difference between them is proportional to the strength of the magnetic field.

This is similar to a mechanical analogy: imagine trying to rotate a bar magnet in a strong external field. If you align it with the field, that is the stable, low-energy position. If you force it to oppose the field, that requires work and corresponds to a higher-energy state.

This splitting of energy levels is fundamental in MRI. It is the small difference in populations between protons in the lower-energy parallel state and the higher-energy antiparallel state that generates the net magnetization we rely on to produce signals.

slide24:

When protons are placed in a strong magnetic field, they split into two energy states: the parallel state, which has lower energy, and the antiparallel state, which has higher energy. This is the Zeeman effect.

The energy of each state can be expressed as:

E equals negative mu z times B naught.

Here, mu z is the longitudinal component of the magnetic moment, and B naught is the strength of the external magnetic field.

Because mu z can take on two values, plus gamma h over four pi, or minus gamma h over four pi, the energies are:

E equals plus or minus gamma times h times B naught, divided by four pi.

The difference between these two levels is what we call delta E, and it is given by:

Delta E equals gamma times h times B naught, divided by two pi.

This difference is critical because it sets the resonance condition. Only when the energy of the applied radiofrequency matches this delta E can protons flip between the two states.

At body temperature, slightly more protons are in the lower energy parallel state compared to the higher energy antiparallel state. This small imbalance is what produces the net magnetization that we detect in MRI.

slide25:

Now, let's bring in some mathematics through the Boltzmann distribution, which describes how protons split between the two energy states.

We can write the ratio of the number of protons in the anti-parallel state to the number in the parallel state as:

"N anti-parallel divided by N parallel equals the exponential of negative delta E over k T."

Here, delta E is the energy difference between the two states, k is the Boltzmann constant, and T is the absolute temperature in kelvins.

Substituting the expression for delta E, we get:

“N anti-parallel divided by N parallel equals the exponential of negative gamma times h times B-zero, divided by two pi k T.”

With a first-order approximation, this becomes:

“N anti-parallel over N parallel is approximately equal to one minus gamma h B-zero over two pi k T.”

What does this mean in practice? It means the populations of the two states are nearly equal, with only a tiny excess in the parallel, or lower-energy state.

The net MRI signal comes from this small population difference. If we take the difference in numbers between the parallel and anti-parallel states, we get:

“Delta N equals N s times gamma h B-zero divided by four pi k T.”

Here, N s is the total number of protons in the body. The key point is that this difference is extremely small. For example, at a magnetic field strength of one and a half tesla, out of about one million protons, only a handful contribute to the net signal.

So, although the individual imbalance is tiny, the collective contribution of billions and billions of protons is what makes MRI signals measurable.

slide26:

Now, all of these models we’ve discussed, including the Boltzmann distribution, set the stage for the next key idea: precession.

When a magnetic moment sits inside a strong external magnetic field, it doesn’t simply align once and stay fixed. Instead, it undergoes a motion very much like a spinning top. The top leans and slowly wobbles around the vertical axis. In the same way, the magnetic moment traces out a precession around the magnetic field direction. We’ll go into more detail on this soon.

To explain why this happens, we need to go back to Newton’s second law—not just in the familiar form, force equals mass times acceleration, but in the rotational context. In rotations, the central quantity is angular momentum.

The law tells us that the time derivative of angular momentum, written as dL over dt, is equal to torque. So, torque is the rotational analog of force. Just as force changes linear momentum, torque changes angular momentum.

The derivation on this slide shows the details. If you’re interested in the full proof, you can follow it step by step. If not, you can simply trust the key result: the rate of change of angular momentum equals torque.

This is the foundation we need to understand precession in MRI physics.

slide27:

This is a classic example of precession, and it is the same principle that applies in MRI when the magnetic moment of protons precesses in a magnetic field.

Think of a spinning top. The top's axis of rotation points in one direction, representing the angular momentum. If no external force acts on it, the top will keep spinning steadily, just as a moving object keeps moving in a straight line.

Now, if a force is applied — here, gravity acting at the center of mass — it produces a torque. Torque is given by the cross product of the lever arm vector r and the force vector F. In this case, r cross F points into the page. That means the torque is perpendicular to the angular momentum.

When the torque is perpendicular, it doesn't stop the motion or flip the top over directly. Instead, it causes a small change, delta L, that is also perpendicular to the original angular momentum. This sideways change makes the tip of the angular momentum vector trace out a circle.

So instead of falling over, the spinning top precesses—it wobbles around in a circular path.

And that is the key insight: whenever a spinning object with angular momentum experiences a torque perpendicular to it, the result is precession. In MRI, the proton's magnetic moment behaves just like this spinning top, precessing around the direction of the external magnetic field.

slide28:

This idea is not too hard to grasp if we think about orbital motion.

Take the example of the Moon orbiting around the Earth, or a satellite orbiting a planet. You might wonder: why doesn't it just fall straight down? The answer is that there is always a centripetal force pointing inward, toward the center of the orbit.

That centripetal force doesn't change the magnitude of the velocity, but it does change its direction. Because the force is always perpendicular to the velocity vector, it continuously "bends" the path, keeping the satellite in circular motion rather than letting it fly off in a straight line.

Mathematically, the centripetal force is equal to mass times velocity squared, divided by the radius: F equals m v squared over r.

Now connect this to angular momentum. When the change in a vector is always perpendicular to the vector itself, the result is circular motion. In the orbital case, the velocity vector keeps changing direction, while in angular momentum, the same principle applies: torque causes a perpendicular change in angular momentum, and so the angular momentum vector itself traces out a circle.

This is the key point to understand precession: a perpendicular change leads to circular motion.

slide29:

When a magnetic moment is placed in an external magnetic field, it experiences a torque. This torque does not flip the magnetic moment directly into alignment with the field. Instead, it causes the moment to precess, meaning it rotates around the axis of the magnetic field, very much like a spinning top under the influence of gravity.

The rate of this precession is called the Larmor frequency. The Larmor frequency depends on two things: the strength of the magnetic field and a constant called the gyromagnetic ratio, which is unique for each type of nucleus.

Mathematically, we write the Larmor frequency as omega zero equals gamma times B zero. Here, omega zero represents the precession frequency, gamma is the gyromagnetic ratio, and B zero is the strength of the magnetic field.

This relationship tells us that the stronger the field, the faster the precession. And this precession frequency is the key to MRI, because it sets the frequency at which protons respond to radiofrequency pulses and generate the measurable MRI signal.

slide30:

Now let’s carefully look at how we can compute the precessional frequency.

When a magnetic moment, which we call mu, is placed at an angle to the external magnetic field B, it experiences a torque. This torque can be written as mu cross B zero. That torque is what drives the angular momentum to change over time.

Mathematically, we can say: torque equals the derivative of angular momentum with respect to time. In other words, dP by dt equals mu cross B zero.

Here, P is the angular momentum. Because protons both have mass and positive charge, their spinning motion not only produces angular momentum but also generates a tiny current loop. That current loop is associated with a magnetic moment. And so, whenever the magnetic moment is not aligned with the magnetic field, it feels a torque.

Now, to connect the pieces: magnetic moment and angular momentum are directly proportional. That means the change in angular momentum is tied to the change in magnetic moment.

To see how this becomes precision, consider a very small angular change, which we call d phi. For such small changes, the sine of d phi can be approximated as just d phi. With this, we can relate torque to the rate of change of the angular position.

So the angular frequency of precession, written as omega, is equal to the torque divided by the magnitude of angular momentum times sine theta. If you expand the cross product, you eventually find that omega equals gamma times B zero.

Here, gamma is the gyromagnetic ratio, a constant specific to the nucleus.

This final result is extremely important: the precessional frequency is directly proportional to the strength of the magnetic field. In a stronger field, the spins precess faster. In a weaker field, they precess more slowly. This proportionality is the foundation of magnetic resonance.

slide31:

It’s often much easier to demonstrate precision than to explain it with equations.

Here you see a gyroscope. Its circular motion around the stand is precession. This motion results from the combination of two effects: the downward force of gravity and the angular momentum around the spinning axis.

Notice how the wheel does not simply fall. Instead, because of torque, it precesses around the point where it is supported. If I increase the downward force—for example, by adding the weight of a wrench to the axle—you can see that the precession frequency increases.

This is the same principle that applies in MRI. The frequency of precession is directly linked to the strength of the magnetic field. Stronger forces or fields lead to faster precession. We call this frequency encoding, and it is the foundation for how MRI gathers spatial information.

slide32:

Now we move from individual protons to the idea of magnetization.

Each proton carries a magnetic moment, and as we have discussed, these moments precess around the external magnetic field. In the population of protons, slightly more occupy the lower-energy state, aligned parallel to the magnetic field, while slightly fewer are in the higher-energy, anti-parallel state. This imbalance, even though very small, is crucial.

Because the magnetic moments are vectors, those that are opposite to each other will cancel out in the transverse plane. What remains is a small excess of protons in the parallel state, and when we add up all of their contributions, we obtain a single net vector pointing along the direction of the magnetic field. This is what we call the net magnetization vector, often denoted as M.

From this point on, instead of worrying about the behavior of individual protons, we can focus on this collective magnetization vector. We can treat it as a classical vector that can tilt, rotate, or precess, just like a spinning top. This makes it much easier to understand the physics and to design MRI pulse sequences.

So, magnetization is the bridge from the microscopic quantum world of spins to the macroscopic signals we use in MRI.

slide33:

Now we bring everything together and look at the steady state of magnetization.

Earlier, we talked about the tiny surplus of protons in the lower-energy parallel state compared with those in the higher-energy anti-parallel state. That difference is very small, but when multiplied by the enormous total number of protons in the body, it adds up to a measurable effect.

For a single proton, the magnetic moment can be expressed as gamma times h-bar over two. But now, instead of just one proton, we have on the order of 10

to the power of 23 protons, and only a small fraction contributes to the net imbalance. Multiplying that tiny difference by the large population gives us the bulk magnetization.

We call this net magnetization M naught. It represents the collective contribution of all the protons and points along the direction of the external field B naught.

In this steady state, the magnetization only has a longitudinal component. That means M z equals M naught, while the transverse components M x and M y are both zero.

Even though individual protons are still precessing and flipping between parallel and anti-parallel orientations, when you take the average, the only surviving component is along the z-axis. This is what we refer to as the steady-state magnetization.

slide34:

So far, we've spent quite some time on the physical foundation. That's good, because now you have a solid grounding in the physics behind MRI — angular momentum, magnetic moment, precession, and magnetization. All of this is essential because without that base, the next steps would be very difficult to follow.

Now we're ready to move forward and talk about how the MRI signal is actually generated. This is the transition point — from theory into application. The key question becomes: once the protons are aligned and precessing in the magnetic field, how do we disturb them in a controlled way and measure the response?

That's where the radiofrequency pulse comes in. By applying a carefully tuned RF pulse, we can perturb the magnetization vector and create a measurable signal. In the upcoming slides, I'll explain this process step by step — beginning with RF perturbation, then the Bloch equations that describe the dynamics, and finally, how we detect the free induction decay signal.

So, let's pay attention to the next slide — this is where MRI really comes alive.

slide35:

So here is the key motivation for MR signal generation. We have this net magnetization vector, which is called M-naught. It reflects the water and lipid content in the body, because that's where most of the hydrogen protons are found. That sounds great, but the important question is: how do we actually measure M-naught?

If you just leave M-naught alone in its steady state, it aligns with the main magnetic field and stays constant in time. A static field like this will not generate any measurable electrical signal. Remember, electromagnetic induction only works when the magnetic field is changing. A constant magnetization produces nothing in our detection coil.

The key idea is this: if we can flip M-naught away from its alignment with the main field, then it will gain a transverse component. That transverse part does not just sit there — it precesses, meaning it rotates around the axis of the main magnetic field. This rotating transverse magnetization produces an alternating magnetic field. An alternating field is exactly what we need, because it induces an electrical signal in the detection coil.

This was the breakthrough that launched nuclear magnetic resonance. It was discovered in nineteen forty-six by Felix Bloch at Stanford and Edward Purcell at Harvard, and their work was recognized with the Nobel Prize in Physics in nineteen fifty-two.

So the essential trick is: flip the net magnetization, create a transverse component, let it precess, and detect the alternating signal that comes out. This is the foundation of MRI signal generation.

slide36:

When the magnetization vector is in steady state, it is balanced. No magnetic field is changing, so no signal can be detected. But once we flip it, we break that balance. Now there is a transverse component of magnetization. This transverse part does not stay still — it undergoes circular motion, what we call precession. And this precessing transverse component produces an alternating magnetic field, which can then be measured through electromagnetic induction, as described by Maxwell's equations.

So how do we flip the magnetization to break the balance? To answer that, we look at the energy difference between the parallel state and the anti-parallel state of the protons. Earlier, we derived this difference and called it delta E.

This energy difference must be supplied by an incoming electromagnetic wave. And the energy carried by a photon is equal to Planck's constant times the frequency. That means the frequency of the applied radiofrequency wave has to exactly match the energy gap between the two spin states.

When you rearrange the formulas, you find that the resonant frequency, written as f, is equal to gamma times B-zero divided by two pi. In terms of angular frequency, written as omega, the result is omega equals gamma times B-zero.

This is a fundamental result. It tells us that the Larmor precession frequency of the proton is identical to the frequency of the radiofrequency field we must apply to drive transitions between the parallel and the anti-parallel states.

So, from the quantum mechanical point of view, resonance occurs when the radiofrequency wave carries exactly the right energy to match that gap — and that is the basis of nuclear magnetic resonance.

slide37:

We can now look at RF excitation from the quantum mechanical point of view.Imagine we have two energy levels: a lower energy state and a higher energy state. When we send in a radiofrequency wave at the resonance frequency, that wave carries just the right amount of energy to promote a proton from the lower state up to the higher state.

But once the RF excitation stops, those protons in the higher energy state cannot stay there forever. They naturally relax back down to the lower energy state, which is the more stable configuration. As they do so, the excess energy is released. Importantly, the released energy is in the form of the same radiofrequency signal that was used to excite them.

This emitted signal is what we detect externally. It is the key step that allows us to convert the microscopic spin transitions of hydrogen protons into a measurable macroscopic MR signal.

slide38:

Now, let’s look at the classical view of RF excitation.We have our net magnetization vector, which represents the collective magnetic moment of all protons. In the steady state, this vector points along the z-axis, aligned with the static magnetic field B-zero.

When we apply an additional magnetic field, called B-one, along the x-axis, something interesting happens. The cross product between the magnetization vector and B-one produces a torque. This torque acts in the direction perpendicular to both vectors, and it causes the magnetization to tip away from the z-axis.

But remember, the system is still precessing around the main magnetic field. So the magnetization doesn’t simply tilt once and stop—it spirals away from the z-axis, tracing out a cone. The longer we apply the B-one field, the further the vector tips, increasing the flip angle.

This picture is very convenient to describe using the concept of a rotating frame. In that frame, the B-one field looks stationary, and the magnetization vector simply rotates around it. That’s why we often describe RF excitation as “flipping” the magnetization into the transverse plane.

slide39:

In the laboratory frame, the magnetization vector doesn’t just tip smoothly into the transverse plane. Instead, it traces out a spiral path, because while it is flipping down, it is also precessing around the z-axis at the Larmor frequency. So what we observe looks like a helical trajectory.

But there is a much simpler way to describe this motion. Instead of staying in the laboratory coordinate system, we can shift into a rotating frame of reference. This is a coordinate system that rotates at exactly the Larmor frequency.

In this rotating frame, the spiral disappears. The magnetization no longer seems to precess—it simply tilts, or “flips,” into the transverse plane. This makes the description of RF excitation much more intuitive.

That’s why, throughout MRI physics, we often switch between the laboratory frame and the rotating frame. The physics is the same, but in the rotating frame, the mathematics and the visualization become much simpler.

slide40:

Think about this in terms of a revolving door. If you just walk up and push the door randomly, sometimes your push helps it move, and sometimes it cancels out—because your force is not always aligned in the right way.

But if you keep pushing at exactly the right rhythm, always perpendicular to the door, you will steadily drive it around in a smooth motion. That is resonance: applying a force at the right frequency so that every push adds up.

In MRI, the same idea holds. The RF signal has to be applied at the exact precessional frequency—the Larmor frequency. If it matches, then every cycle of the RF field keeps nudging the spins in the same direction. In the rotating frame, this no longer looks like an oscillating wave. Instead, it looks like a steady force, a constant torque, applied to the magnetization vector.

That constant torque is what tips the magnetization away from the z-axis. It is what allows us to rotate the vector from its equilibrium position into the transverse plane, or even further.

So the revolving door is a nice analogy: push in rhythm with the rotation, and the system responds strongly. Push out of rhythm, and the effect cancels out.

slide41:

We describe the change of the magnetization vector M of t with time.The Bloch equation is:

dM of t over d t equals M of t cross gamma times B of t, minus R times the difference of M of t and M zero.

Here:

gamma is the gyromagnetic ratio,

B of t is the magnetic field,

M zero is the equilibrium magnetization,

and R is the relaxation operator that encodes T1 and T2 processes.

Now, if we expand this into components, we get three equations:

For the z component:dM z over d t equals gamma times the quantity M x times B y, minus M y times B x, minus the difference M z minus M zero divided by T1.

For the x component:dM x over d t equals gamma times the quantity M y times B z, minus M z times B y, minus M x divided by T2.

For the y component:dM y over d t equals gamma times the quantity M z times B x, minus M x times B z, minus M y divided by T2.

slide42:

Now let’s see how the Bloch equation works in practice.

First, consider precession. If the magnetization vector, capital M, is not parallel to the magnetic field B, then M will precess around B. On the other hand, if M is perfectly parallel to B, then the precession radius is zero. In that case, you do not see any visible precession, even though the vector is still technically circling like a single dot.

Next, think about what happens when we flip M. Suppose M is rotated away from the z-axis into the transverse plane. Now the magnetization has a horizontal component, which we call M x y. This component will oscillate and produce an alternating magnetic field. That is the field that induces a measurable signal in the nearby coil.

Physically, you can imagine the magnetization vector as a magnetic bar. If you keep flipping or rotating this bar, the field it produces also keeps changing. That changing field drives current in the coil, which is exactly what we detect as the MR signal. The amplitude of this signal depends on proton density: the more protons available, the stronger the signal.

Now, after the excitation pulse ends, the system will gradually return to equilibrium. The z-component of magnetization, M z, will recover to its steady-state value M zero. This recovery is governed by the T1 relaxation process, sometimes called spin-lattice relaxation. At the same time, the transverse component M x y will decay toward zero. This loss of coherence is governed by T2 relaxation, also called spin-spin relaxation.

So, to summarize:

T1 describes how the longitudinal magnetization M z returns to equilibrium.

T2 describes how the transverse magnetization M x y decays to zero.

Both of these effects are essential in shaping the MR signal we actually record.

slide43:

Now let's talk about how the MR signal is actually detected.

'

When the magnetization vector, capital M, is flipped away from the z-axis, it develops a transverse component, which we call M x y. This transverse magnetization does not remain static — it oscillates at the resonance frequency. That oscillation induces a sinusoidal current in the radiofrequency coil, exactly as predicted by Faraday's law of electromagnetic induction.

The strength of this induced signal depends directly on the size of M x y. The larger the transverse component, the stronger the oscillating magnetic field, and therefore the stronger the detected voltage in the coil.

The maximum possible signal is reached when the flip angle is ninety degrees. At this angle, the entire magnetization vector lies in the x-y plane, meaning that M x y is equal to M naught, the full equilibrium magnetization. This is why a ninety-degree pulse is often used in MRI — it gives the largest possible signal for detection.

slide44:

Once the radiofrequency pulse is turned off, the magnetization vector begins to relax back toward the main magnetic field, capital B naught. In this process, the excess energy that was absorbed is released.

As the system relaxes, the transverse magnetization — that is, M x y — begins to decay. This decay produces a signal that gradually decreases to zero. We call this the free induction decay, or FID.

The FID is the fundamental nuclear magnetic resonance signal. It oscillates at the resonance frequency, and its strength is proportional to the local proton density. In other words, the more protons you have in a region, the stronger the signal you record.

This decay curve, shaped by relaxation mechanisms T1 and T2, contains the raw information that MRI uses. By capturing and analyzing this signal, we can reconstruct images that reveal tissue structure and composition.

slide45:

Now that we have covered how MRI signals are generated and detected, the next step is to understand how those signals decay. This part is crucial, because the decay mechanisms — specifically T1 and T2 relaxation — are what ultimately give us image contrast.

So, in the final section of this lecture, we will look more closely at the physics behind T1 and T2. We will also see how these relaxation times can be measured, and how pulse sequences such as inversion recovery and spin echo are designed to take advantage of them.

Up to this point, our discussion has focused on the overall process of signal generation and detection. But to move from signals to images, we need to understand relaxation. T1 and T2 are the keys that translate raw MRI physics into meaningful image contrast.

slide46:

So now let's ask: why does the MR signal decay?

The answer is very much like the rolling ball analogy. A ball on the ground slows down because of resistance and friction. In MRI, once we flip the magnetization into precession, it cannot stay there forever. Over time, it will naturally return to its steady state, aligned with the main magnetic field, B naught.

This return happens through different relaxation mechanisms:

T1 relaxation, also called longitudinal or spin-lattice relaxation, is when flipped nuclei realign with the main magnetic field. The energy they lose is given back to the surrounding tissue as thermal energy.

T2 relaxation, also called transverse or spin-spin relaxation, is when the nuclei gradually lose phase coherence with one another. Even though they were flipped together, local differences in their environment cause them to fall out of step.

T2-star relaxation is an additional effect that comes from imperfections in the magnetic field itself. The field is never perfectly uniform, and small local variations cause the spins to dephase even faster. This effect is especially important in functional MRI, where susceptibility differences are actually used to generate contrast.

The key point is that the MR signal we detect — the NMR signal — is proportional to proton density, but reduced by these T1, T2, and T2-star factors.

And this is what gives MRI its powerful contrasts. For example:

T1 weighting helps us distinguish gray matter from white matter.

T2 weighting highlights tissues and fluid, such as cerebrospinal fluid.

T2-star weighting is especially useful for functional MRI, because it is sensitive to magnetic susceptibility changes, such as those caused by blood oxygenation.

So, signal decay is not a nuisance. In fact, it's what makes MRI such a versatile imaging tool.

slide47:

Now let’s take a closer look at the T1 effect.

T1 describes how the longitudinal magnetization, that is, the component along the z-axis, gradually returns to its equilibrium value after being disturbed by a radiofrequency pulse.

On the left side of the figure, we see what happens after a ninety-degree pulse. In this case, the longitudinal magnetization, which we call M z, has been tipped entirely into the transverse plane, so it starts at zero. Over time, it recovers back toward its equilibrium value, which we call M naught, and it does so in an exponential fashion.

On the right side, we see the case of a one-hundred-eighty-degree pulse. Here, the longitudinal magnetization is flipped completely upside down, starting at negative M naught. From there, it also recovers exponentially back to M naught.

This exponential recovery is the hallmark of T1 relaxation. The recovery speed depends on the tissue type, because different tissues exchange energy with their environment—the so-called spin-lattice interaction—at different rates.

So, to summarize: T1 tells us how quickly a tissue’s magnetization realigns with the main magnetic field, B zero. And these differences in T1 times are one of the key sources of image contrast in MRI.

slide48:

Now let’s talk about the T2 effect, which is also called transverse relaxation or dephasing.

After we apply a ninety-degree pulse, the net magnetization lies entirely in the transverse plane, along the x-y plane. At that initial moment, all of the tiny magnetic moments of individual protons are lined up together, pointing in the same direction.

But over time, each of those protons experiences slightly different local magnetic fields. Some of them precess a little faster, and others precess a little slower. As a result, they gradually drift out of phase with one another.

When that happens, their contributions begin to cancel out. The net transverse magnetization shrinks, even though each proton is still spinning.

On the graph at the bottom, you can see this process clearly. The transverse magnetization starts at its maximum value and then decays exponentially toward zero as the moments become more and more dephased.

This is the essence of the T2 relaxation process: it measures how quickly the spins lose phase coherence in the transverse plane, leading to signal decay.

slide49:

At this stage, we can summarize how T1 and T2 relaxation describe the time evolution of magnetization in MRI.

First, consider the longitudinal component, which we usually call M z. After a radiofrequency pulse, it does not instantly return to equilibrium. Instead, it recovers gradually, following an exponential curve. Mathematically, we say that M z of time equals M naught multiplied by one minus e to the power of negative t over T1. In plain words, this means the recovery depends on a time constant, T1, which tells us how quickly the spins realign with the main magnetic field.

Next, for the transverse components, which are M x and M y, these decay over time. Their rate of decay is described by the time constant T2. The equations are written as d M x over d t equals minus M x divided by T2, and d M y over d t equals minus M y divided by T2. In other words, both x and y components shrink exponentially toward zero, at a rate determined by T2.

So, together, T1 and T2 give us a semi-quantitative model for MRI. T1 governs how fast the longitudinal magnetization recovers, and T2 governs how fast the transverse magnetization disappears. Both processes follow exponential curves, but with different time constants.

slide50:

Here we have a summary table of T1 and T2 relaxation times for different tissues, measured at a field strength of one point five Tesla.

For fat, the T1 relaxation time is about two hundred sixty milliseconds, and the T2 is about eighty milliseconds.

For muscle, T1 is around eight hundred seventy milliseconds, while T2 is much shorter, about forty-five milliseconds.

In the brain, gray matter has a T1 of about nine hundred milliseconds and a T2 of about one hundred milliseconds. White matter is slightly shorter, with a T1 of about seven hundred eighty milliseconds and a T2 of about ninety milliseconds.

The liver shows a T1 of about five hundred milliseconds and a T2 of about forty milliseconds.

Finally, cerebrospinal fluid has the longest times by far: a T1 of about two thousand four hundred milliseconds and a T2 of about one hundred sixty milliseconds.

These values illustrate a key point: different tissues have different relaxation times, and that difference is what gives MRI its powerful ability to generate contrast between tissues.

slide51:

Now let’s distinguish between T2, T2 prime, and T2 star.

The loss of phase coherence in the transverse magnetization can arise from two different mechanisms. The first is the so-called pure T2 decay, which comes from intrinsic spin–spin interactions inside the tissue. This is a physiological and biological effect, and it represents the fundamental limit of how long spins can stay in phase with one another.

The second contribution comes from spatial variations in the magnetic field strength inside the body. In practice, no magnet is perfectly uniform. There are always small imperfections in the main field, B zero. In addition, different tissues have slightly different magnetic susceptibilities, especially at boundaries between air and tissue, or bone and tissue. These local variations cause spins to accumulate extra phase shifts, further accelerating the loss of coherence. This additional contribution is represented by T2 prime.

When we put these together, the overall transverse relaxation time we observe is called T2 star. Mathematically, the rate of decay of T2 star is the sum of the rates of T2 and T2 prime. That is, one over T2 star equals one over T2 plus one over T2 prime.

So in summary:

T2 reflects fundamental spin–spin interactions.

T2 prime reflects field inhomogeneity and susceptibility effects.

T2 star combines both, and it is the effective decay time we actually measure in many MRI sequences.

slide52:

Now, let’s look at how T1 can be measured using the inversion recovery method.

The idea is to start with a one-hundred-eighty-degree pulse, which flips the net magnetization completely upside down along the negative Z-axis. Once it is inverted, we allow it to recover toward equilibrium during a delay time, which we call tau. During this period, the longitudinal magnetization gradually regrows toward its steady state, following the exponential recovery curve we discussed earlier.

At the end of the delay, we apply a ninety-degree pulse. This flips whatever longitudinal magnetization has recovered into the transverse plane, where it produces a measurable signal. By repeating this process with different values of tau, we can track how the magnetization regrows over time.

Mathematically, the detected signal follows the expression:S of tau equals M zero times the quantity one minus two times e to the power of minus tau over T1.

By fitting this recovery curve, we obtain the T1 relaxation time. In practice, this method gives a very robust way to measure T1 across different tissues, and it is the basis for many T1-weighted imaging protocols in MRI.

slide53:

Now let’s visualize what’s happening in an inversion recovery sequence.

We start with the net magnetization pointing up along the Z-axis at its equilibrium value, M_0. Then, we apply a one-hundred-eighty-degree pulse, which flips it completely upside down, to minus M_0. From this inverted state, the magnetization begins to relax back toward equilibrium. Over time, the longitudinal component regrows along the Z-axis following the T1 recovery curve.

At some chosen inversion time, which we call TI, we apply a ninety-degree pulse. This rotates whatever longitudinal magnetization has recovered into the transverse plane. Once in the transverse plane, it generates a measurable free induction decay, or FID signal, that is subject to T2 star decay.

By repeating this process with different inversion times and recording the signals, we can map out the full recovery curve. Then, using a logarithmic fit, we extract the T1 relaxation time.

So, in short, this diagram shows how the one-hundred-eighty-degree inversion, the recovery period, and the ninety-degree readout pulse together make it possible to measure T1 directly.

slide54:

Now let's look at how we measure T2 relaxation using the spin-echo method.

We begin with a ninety-degree pulse that tips the net magnetization into the transverse plane. Right away, the spins begin to dephase because of small differences in their local magnetic environments. Faster spins move ahead, slower spins lag, and the overall signal decays quickly with a time constant that looks like T2 star.

But here is the clever trick: at a certain time, we apply a one-hundred-eighty-degree pulse. This flips all the spins over, so that the ones that were leading are now behind, and the ones that were lagging are now in front. As time continues, these spins re-converge, and at a later time, they come back into phase. When that happens, we see the spins add constructively and produce a strong signal, known as the spin echo.

The height of this echo decays with the true T2 relaxation time, not with T2 star. By repeating the sequence with different time delays, we can track how quickly the echo signal diminishes, and from that, extract the T2 constant.

This technique is both simple and powerful: it cancels out static inhomogeneities in the magnetic field and isolates the true spin-spin relaxation, which is exactly what T2 represents.

slide55:

To wrap up today's lecture, here is your homework assignment.

For problem 4.3, you will calculate the effects of different pulse sequences on thermal equilibrium magnetization. Your final answers should include the x, y, and z components of magnetization. The cases include a ninety-degree pulse about the x-axis, an eighty-degree pulse about the x-axis, two consecutive ninety-degree pulses about x and y, and finally, two consecutive eighty-degree pulses about x and y.

For problem 4.4, you will decide whether each statement is true or false and provide a brief explanation. These questions cover recovery of magnetization, the behavior of the static field B0, relaxation from the transverse to the longitudinal axis, and the interpretation of a short T1 relaxation time.

The due date is one working week from today.

Thank you for following along — we've covered a lot of ground, from the Bloch equations to T1 and T2 relaxation and even spin echo sequences. This sets the stage for more advanced imaging concepts in our next lecture.