slide1:

In this lecture, we move into the second part of MRI imaging. The first part, as you saw in the video, focused on the underlying physical principles. Now, we will explain how MRI imaging is performed, based on those physical models and principles.

It's best to begin by reading the chapter first. You may naturally have questions, but that's exactly why you attend this lecture. Without that preparation, it's often difficult to build a solid understanding. MRI is quite tricky—you need to carefully think about the concepts, the relationships, and revisit them several times. Hence, an iterative approach should be used to preview, discuss, and review, again and again.

slide2:

Again, let us see where we are. We started with medical physics—specifically MRI physics—in the first lecture. Now, we move to MR imaging. In our next lecture, we'll talk about MRI pulse sequences, which will give us a more specific look at how MRI is applied in practice.

However, today, our goal is to first revisit what we have learned so far.

slide3:

The first part of today's lecture is a review, just in case you missed something earlier. Even if you understood it the first time, it's always good to refresh these key concepts.

The key point is to model and understand T1 and T2, which are critical relaxation parameters. Along with these, proton density is also an essential factor, and fortunately, that one is fairly easy to understand. Think of it this way: you have a collection of small magnets, whose density is proportional to proton density, and together they form a large magnetization vector, usually called the big M vector. Once you begin to manipulate this M vector, you immediately see the effects of T1 and T2.

T1 and T2 are intrinsic properties of tissues. They vary with tissue type, and they are very important for functional and even molecular imaging. You need to understand how the MRI signal changes under the influence of T1 and T2, and how these quantities can be measured. When you do this properly, you can create multi-contrast MRI images, such as proton-density-weighted, T1-weighted, and T2-weighted images.

This is much more complicated than what we encounter in X-ray imaging. In X-ray, the main parameter is the linear attenuation coefficient. Even though this coefficient is energy-dependent, it still gives us a straightforward measure: it tells us how many photons, or what fraction of photons, will be attenuated when an X-ray beam passes through a given thickness of a given tissue type.

By contrast, T1 and T2 are not so simple. They require more imagination—especially geometrical imagination—to grasp.

Now, the remaining two parts of today's lecture—slice selection and sampling in k-space—are where MRI becomes truly powerful. These steps allow us to extract MRI signals from the sample, or from the patient, in a spatially specific way. That's how we generate cross-sectional or even volumetric images of the human body.

In summary: today we will begin by revisiting T1 and T2 modeling and measurement, and then move on to slice selection and k-space sampling. These latter two are really the emphasis of this lecture. In the previous session, we talked about MRI imaging principles. Now we will extend that foundation and see how it translates into actual imaging techniques.

slide4:

So, T1 and T2 relaxation, this is kind of a review. T1 and T2 are quite different.

T1: We talk about the recovery of the M vector along the z-axis — that's the main field direction. T1 time is defined as sixty-three percent recovery of our original M vector, which we call M-naught. This is about spin–lattice interaction, something that involves energy.

When you turn the M vector into a certain direction, like a ninety-degree direction, the energy you used — injected into the tissue — will eventually be lost due to spin–lattice interaction. And it will gradually return to the steady state. The M vector wants to be aligned along the main field direction, which is the B-naught field, or the z direction.

On the other hand, the T2 parameter talks about the dephasing effect. The M vector is not just an individual thing — it's really a collective phenomenon. All of the small magnets, the spinning protons, are very small magnets. Altogether, they form the big M vector.

If you flip the M vector, say by ninety degrees, from the z direction into the y or x direction, this big M vector at first points all in one direction. But it is made of many small magnets, and some of those spins are faster, some are slower, because of inhomogeneity in the B-naught field or in the local field. The local field — ions, molecules, molecular motion — all these things contribute to that inhomogeneity.

Originally, all the small magnets line up along the x-axis. If you apply a ninety-degree pulse with respect to the y-axis, you flip the M vector along the x-axis. But with time, this will dephase. Some processes are a little faster, some slower. And after a long enough time, all these small magnets, or spinning protons, are pointing in different directions.

The overall field is no longer this original M-naught. Instead, it becomes effectively zero. The decay is modeled as an exponential curve — something like what you see on the slide. So, this is recovery, and this is decay. That is the T1, T2 story.

Let me give you a little more mathematical detail so you understand why both the recovery and the dephasing follow an exponential curve. One is growing, the other is decaying. It's something like this.

slide5:

So, the change rate, whether it is growth or decay, is described by an ordinary first-order differential equation. The first derivative — the change rate — is proportional to the original quantity.

When we talk about the change of y, the rate of change is proportional to y. You keep seeing this kind of first-order differential equation in multiple imaging modalities. For example, in X-ray imaging, you have the initial beam intensity. It will be attenuated. The amount of attenuation — the change of the X-ray flux — is proportional to the initial flux, which we call I-naught.

Similarly, here in MRI, with y, we are talking about the initial M vector, which is M-naught. So we have this relationship.

slide6:

Whenever you have a governing relationship like this, you can go through some mathematical steps.

With the first derivative like this, you just rearrange it a little bit. So you have:(one over y) d-y equals k d-t.

Then, integrate both sides of the equation. On the left, you get ln absolute y, and on the right, you get k t plus C.

Now, exponentiate both sides. That gives:y equals A e to the k t, where A is a constant.

This shows that whenever the process is described by a first-order differential equation, the solution must be of this form: y is exponential growth or exponential decay. In other words, a constant multiplied by an exponential function, with index k t, where t is time.

So, this explains why you see exponential decay, exponential growth, or attenuation in earlier imaging modalities like X-ray imaging, and why you see the same behavior here in MRI.

slide7:

Now, for T2, we need to look in more detail. T2 has two components, and it describes dephasing. The first dephasing is what we call the pure T2 effect, or spin–spin interaction. This happens because all the spinning protons are close to each other, and you do not have an absolutely uniform distribution. If you look microscopically, the spins interact with one another, and there are fluctuations in the local magnetic field. These effects are called pure T2, shown here in red, and they are due to spin–spin interaction. This behavior in soft tissue is physiological, and it cannot really be changed — it is a natural biological phenomenon.

On the other hand, when you place a patient, or even a small animal, in an external magnetic field — the B-naught field — we usually describe that field as uniform. But in reality, the B-naught field is never perfectly homogeneous. There will always be some inhomogeneity. This inhomogeneity is fixed: when you set up the B-naught field, some regions are slightly weaker, while other regions are slightly stronger. So, in one location, the spins process a little slower, and in another location, they process a little faster. That difference also contributes to dephasing.

So, both biological spin–spin interaction inside the tissue, and physical inhomogeneity of the external magnetic field, contribute to the same dephasing effect, but for different reasons. The external contribution is fixed and constant, while the biological part cannot really be corrected.

When we combine these two, we call the result T2 star. The external field inhomogeneity alone is called T2 plus. So, in summary: T2 star is the combined effect of external field inhomogeneity and biological spin–spin interaction. And importantly, there is a reciprocal relationship between them.

slide8:

Let’s explain the reciprocal relationship. We start with a constant, K. This K describes the strength of dephasing. If dephasing is due to spin–spin interaction, we can call it K1. If dephasing is due to external magnetic field inhomogeneity, we can call it K2. So, the total rate of dephasing is due to both K1 plus K2. Then, the time constant is simply the reciprocal of that total: one over K1 plus K2. That’s where the reciprocal relationship comes in.

If you still feel unclear, let’s connect it to something we already studied: nuclear imaging. When we studied nuclear decay, we said that a radioactive tracer inside the body disappears through two processes. One is physical radioactive decay, which happens naturally over time. The other is biological clearance, where the tracer leaves the body through circulation, mainly through urination. So, there are two mechanisms removing the tracer — one biological and the other physical — and together, they follow the same reciprocal rule.

In that case, we use the decay constant, lambda. The overall decay rate is lambda-one plus lambda-two, and the corresponding half-life is the reciprocal of lambda. So, just like in nuclear decay, where both physical and biological processes add together, in MRI, both spin–spin interaction and external inhomogeneity add together. Don’t be confused if you see these different equations — equation two-point-four for nuclear medicine and equation four-point-two-nine for MRI — because mathematically, they are describing the same idea.

To summarize: when we talk about T2, we really mean three different things. One is the biological T2, or pure spin–spin relaxation. The second is T2 plus, from external field inhomogeneity. And the third is T2 star, the combined effect of both. Now we have T1 and T2. The question is: how do we measure them?

slide9:

To measure T1, we often use a technique called inversion recovery. Here’s how it works. First, we apply a one-hundred-eighty-degree pulse. This flips the magnetization vector completely upside down. At this moment, there are no transverse components. Remember, only transverse components induce a signal in the nearby coil. Without a transverse component, if you place a coil close to the patient, there is no signal to measure.

Now, after a time delay, tau, we apply a ninety-degree pulse. This rotates some of the magnetization into the transverse plane. And once that happens, the coil can detect a signal. This is the signal you see here. Notice that the signal itself also decays. Why? Because of dephasing. And ask yourself — is this dephasing due to T2, or due to T2 star? The answer is T2 star, because it includes both spin–spin interaction and the effect of field inhomogeneity.

So, this is the inversion recovery method: a one-eighty-degree pulse, followed by a time delay, and then a ninety-degree pulse that produces a measurable signal. The key question we then ask is, what is the amplitude of that signal?

slide10:

This really depends on tau, the delay time, before you apply the ninety-degree pulse. And this is illustrated here.

Let's review carefully. First, you apply a one-hundred-eighty-degree pulse, and the magnetization vector M-naught becomes minus M-naught. If you don't do anything, this remains negative. Originally, you had a positive M-naught, but after the one-eighty pulse, it became negative M-naught. This negative magnetization will gradually recover back toward the original M-naught, returning to the normal, stable equilibrium. If you wait long enough, this negative value will rise, reach zero, and then continue climbing until it becomes fully positive again.

So it all depends on tau — the delay. At first, the amplitude is negative, then it grows smaller in magnitude, reaches zero, and then becomes positive, growing larger and larger toward M-naught. Now, if at some point during this recovery you apply a ninety-degree pulse, flipping it into the transverse plane — along the y-axis — you create a transverse component, and that can be measured. At this moment, you will have a signal.

But this signal is subject to T2 star decay, so what you measure is a free induction decay, or F-I-D signal. By repeating this process with different delay times, different values of tau, you can measure the signal amplitude as a function of tau-n. Then, you can plot this relationship. By arranging the equation properly, you get an exponential form with an index of minus tau-n divided by T1. If you take the logarithm, you can obtain a straight line. The slope of that line is minus one over T1, and from this slope, you can determine the T1 relaxation time.

Now, be careful: if you directly take the log, you will not immediately get T1. You need to rearrange the equation. But one important point is that M-naught is known, because you can measure it directly. That's why this works.

slide11:

Now, how do we measure T2? The most important and representative pulse sequence is the spin echo. Spin echo is a very clever design. It allows us to measure T2 directly. If you don't use spin echo, the signal decay you observe will be governed by T2 star, which includes both biological effects and external field inhomogeneities. But what we want is the pure T2, which reflects tissue physiology and pathology.

Pure T2 is a biological effect, while T2 star is dominated by external field imperfections. Since your patient's physiological status should not depend on imperfections in the external field, we really want to measure T2.

So how do we do it? Spin echo. First, you let the spins dephase for a short time. Then you apply a one-hundred-eighty-degree pulse. What happens? Some spins are lagging — they have a phase delay. Other spins are moving faster — they have an advanced phase. After the one-eighty-degree flip, the lagging spins continue to lag, and the faster spins continue to advance. But because their positions have been reversed by the flip, they will eventually meet again and refocus at a certain point.

At that refocusing point, the signal strength grows again, giving you the spin echo. This is a brilliant design. And in fact, you can repeat it multiple times for more echoes if you like.

slide12:

So here's another way to see the spin echo.

You start with the magnetization vector. Apply a ninety-degree pulse, and now you have transverse components. These transverse components begin to dephase — some spins are slower, shown in red, and some spins are faster, shown in black.

Next, you apply a one-hundred-eighty-degree pulse. This instantly flips the distribution. The red, slower spins are moved to the position where the black, faster spins were, and the black spins take the place of the red ones. But remember, the red spins are inherently slower. So even though they were moved ahead, they continue to process slowly. The black spins are inherently faster, so even though they were moved behind, they continue to process quickly.

Now, after the same amount of time passes, both groups realign again. The faster black spins catch up, and the slower red spins fall into place. At that moment, all the spins point together again, aligned in the same direction.

Because they are all aligned, their contributions add up constructively, and the transverse magnetization reaches its maximum. This produces a measurable signal in the nearby coils — the spin echo signal.

slide13:

Here's an even simpler way to explain spin echo. Imagine you have a group of students lining up at a race track. At the starting point, they are all together, right at the start line. After the signal, they all run toward the finish line. Some students are faster, and some are slower. So after a certain time, they will have covered different distances.

Now, at that moment, we ask them all to turn around and keep running, but in the opposite direction. The faster runners, who had gone farther ahead, now have a longer distance to come back. The slower runners, who didn't go as far, have a shorter distance to cover. But because they are slower, their speed matches the distance they need to make up.

After the same amount of time, both the fast and slow runners arrive back at the starting line together. They all realign again. This is the basic idea behind spin echo. By using this trick, we can measure T1, the pure T2, without the interference of the external field inhomogeneity.

slide14:

So this spin echo sequence is a very clever design. And from it, we get two very important parameters.

The first one is the echo time, which we write as T-E. This is defined as the time from the ninety-degree pulse to the echo peak. In practice, you apply a ninety-degree pulse, then a one-hundred-eighty-degree pulse, and the time from the ninety-degree to the one-eighty is T-E divided by two. Then, after the one-eighty pulse, you wait another T-E divided by two, and that's when the spin echo appears.

The spin echo amplitude reflects the decay due to pure T-two, independent of external field inhomogeneity. That's the significance of T-E.

The second parameter is the repetition time, or T-R. After the first sequence, you can wait for a time T-R, and then repeat the same trick, ninety degrees, one-eighty degrees, another spin echo. But in order for the system to be in a stable state, T-R has to be long enough so that the magnetization vector M-naught has fully recovered.

If, for example, T-R is chosen as two thousand milliseconds, that may not be long enough for full recovery. In that case, the longitudinal magnetization is still less than the original M-naught, and the signal will be reduced.

So this is why T-E and T-R are significant. By definition, T-E should always be less than T-R.

Now, this also connects to imaging. At the top, when we look at the spin echo for the whole sample, we get one signal that represents the entire patient or sample as a whole. But in imaging, we want more than that. We want to know the values of T-one and T-two for each pixel or voxel, not just the whole object. That requires us to encode spatial information. We'll get into that starting with the next slide.

slide15:

So now we move to the next part: being spatially specific.

In MRI, we don't just want a signal from the whole sample. We want to isolate a particular slice. Not every layer, but just one slice. And not only that — within that slice, we also want to know the signal from each location.

That means we need to encode each pixel in a way that makes its signal slightly different from every other pixel. Only then can we reconstruct an image that shows the T1 and T2 values for each voxel.

The first step in this process is slice selection. We begin by defining which slice we want to image.

slide16:

So this is about the processional frequency on B. In the textbook, we actually mentioned this in the previous lecture..

But let's focus on the conclusion. The derivation goes through all these steps, and if you take the time, you will understand the details.

The key point is this: the processional frequency is proportional to the local magnetic field. That means if the overall field is stronger, the M vector — or you can think of it as a spinning top — will process faster.

So just follow the picture here. The last line is very clear: the angular frequency of precession, omega, is proportional to the local magnetic field. If we call the main field B-naught, then omega is proportional to B-naught. Later on, if you change B-naught to B-prime, omega will be proportional to B-prime.

And then you multiply by a physical constant, called gamma.

So, in short, omega equals gamma times B.

slide17:

Now, if you want to explore a little more detail, you can check the derivation for this mechanical precession of motion.

In this case, the angular frequency is proportional to the gravitational force. I've marked it in green on the slide. This is basically a physics analogy using a spinning top and torque.

If you're interested, you can read it. Otherwise, don't worry too much — it's not essential for our discussion of MRI.

slide18:

Now here comes the great idea — introducing a gradient field on top of the background field B-naught. This was such an important idea that it led to a Nobel Prize.

So why is this so important? Because this is the key step that allows us to encode spatial information.

In the main field B-naught, the processional frequency of all spins would be the same. That would not allow us to distinguish one location from another.

But if we apply a gradient field, we change the situation. The magnetic field strength will now depend on position. For example, the larger the z-coordinate, the stronger the field. As a result, the processional frequency will also depend on position.

So now, instead of all spins having the same frequency, their frequency varies with location. At z equals zero, the frequency is the original B-naught frequency. But away from zero, it becomes stronger or weaker depending on the gradient.

This is the concept of a field gradient.

slide19:

Once we have a field gradient, we can apply an external rotating magnetic field to push the M vector at its resonant frequency — that is, its processional frequency.

Now think back to the analogy from the previous lecture: remember the photograph of a lady pushing a rotating door? That's the same idea. You have to push the door at just the right time to keep it moving smoothly. Similarly, you have to push the spins at just the right frequency to change their orientation.

And even earlier, we watched a short video where the instructor explained how an MRI works in less than ten minutes. He said you simply send in an RF frequency signal. That RF signal, or RF pulse, is a form of electromagnetic wave. And electromagnetic waves are oscillating magnetic fields.

So when you send in an RF pulse whose frequency is in resonance with the processional frequency of the spins, you are effectively pushing those tiny spinning tops — the protons — in a coordinated way. With each push, you tilt the magnetization vector a little more, and eventually you flip the M vector from its original orientation along the z-axis into one of the horizontal axes, like the x-axis or y-axis.

Now, here's the key: you don't want to flip every proton in the body. You only want to flip the spins in a specific slice of tissue. To do that, you must use both the gradient field and the bandwidth of your RF pulse together.

The frequency range of the RF pulse is centered at some frequency, call it omega-s, and it spans from omega-s minus delta-omega-s to omega-s plus delta-omega-s. So it has a certain bandwidth.

In the time domain, this looks like a rectangular pulse — a signal that is “on” for a finite duration. When you take the Fourier transform, that rectangular pulse in the time domain becomes a sinc-shaped function in the frequency domain. So, in practice, even though you are sending an RF signal in the time domain that looks like a sinc function, in the frequency domain it corresponds to a rectangular function.

This rectangular frequency band means that only those spins whose processional frequencies fall within that band will be excited. Now recall, because of the gradient field, the processional frequency depends on z-location. So different positions along the z-axis correspond to different resonance frequencies.

Therefore, when you send in your RF pulse with its defined bandwidth, only the spins whose z-positions correspond to that frequency band will resonate and be flipped. Spins outside that band will not be affected — they simply won’t respond, because the RF pulse is not in resonance with them.

So, the slice thickness directly corresponds to the frequency range of the RF pulse under the influence of the gradient field. If you make the bandwidth larger, the slice thickness increases. If you make the bandwidth smaller, the slice thickness decreases.

To summarize: you apply a gradient field to make the processional frequency vary with z-location. Then you send in an RF pulse with a defined frequency band. Only those spins in the matching z-range — that slice — are excited and flipped. Everything outside remains untouched. That is how slice selection in MRI works

slide20:

Graphically, we can show this as you see on the slide.

The gradient field creates a range of resonance frequencies along the z-axis. When you send an RF pulse with a defined frequency range, only the spins within that frequency range — and therefore within a specific z-range — are excited.

That z-range corresponds to the slice thickness. So this is how we select a slice in MRI.

What we’ve explained here is the principle for slice selection along the z-axis.

slide21:

Now, let’s look at slice selection in more detail. In your textbook, you’ll find this figure, labeled as 4.16. Here, you see that slice selection can be performed using gradients along the x-axis, the y-axis, or the z-axis. That means you can select a coronal slice, an axial slice, or a sagittal slice, depending on which gradient you apply.

And the beauty of MRI is that you are not limited to just these orthogonal planes. With the right gradient combinations, you can select slices at any orientation you like, even oblique slices at odd angles. This is one of the great flexibilities of MRI imaging.

Now let’s think about why this matters. Before we introduced gradient fields, all we could do was measure T1 and T2 signals from the whole sample — essentially, we had no spatial specificity. We were only getting the overall relaxation properties averaged over the entire tissue volume. But now, with slice selection, we can be much more specific.

We can tell the scanner: “Excite this slice only.” That means the signals we collect come from just that defined region. All the other spins, those outside of the slice — whether they’re above, below, or to the side — will not contribute. Because remember, resonance is only defined for those spins whose frequencies match the selected RF bandwidth, which corresponds to the gradient-defined slice.

So in practice, the signals from outside the slice, those spins located elsewhere, remain silent. Only the spins inside that slice range are excited and measured. This is the principle of slice selection in MRI.

slide22:

Now, there is an important technical subtlety about slice selection that we need to address.

When we apply a gradient field, by definition, it is not homogeneous. Instead, it varies linearly across space. But remember, any deviation from a homogeneous magnetic field introduces a problem: it causes dephasing.

Let’s break this down. On the right-hand side of the diagram, you see the gradient field. On one side, the magnetic field is a little stronger. On the other side, it is a little weaker. Stronger fields correspond to faster precession. Weaker fields correspond to slower precession.

So when you look across the entire slice, the spins are no longer processing together in sync. Some spins are ahead, some are behind. If you just add all these signals together, they partially cancel each other out. The result is a substantial reduction in signal strength. That’s the dephasing problem.

So how do we fix it? We use what’s called a rephasing gradient.

Here’s the idea: within the slice, after you’ve applied the slice-select gradient, the spins have accumulated different phases. That is described mathematically in your textbook by equation 4.41. The phase shift depends on the z-position and on the gradient strength.

But notice the formula includes tau over 2, not tau. Why is it tau divided by 2? Well, think about it this way: during the full pulse duration tau, the spins at the top of the slice accumulate the maximum phase, and the spins at the bottom accumulate the opposite phase. The middle spins are only halfway through. So if you take the average phase across the slice, it corresponds to tau divided by 2. That’s why this factor appears.

Now, to correct for this unwanted dephasing, we apply a second gradient — but this time with opposite polarity. We call it the rephasing gradient. It lasts for a short time, tau-ref, and its strength is chosen so that the overall effect cancels the earlier phase spread.

Mathematically, the product of gradient strength and time for this rephasing pulse must equal tau over 2 times the original gradient strength. That’s equation 4.42 in your textbook.

The result is that the spins, which had drifted apart in phase, are brought back into alignment. This cancels out the artificial dephasing caused by the slice-selection gradient itself.

So now, within the slice, you’re left with a clean signal. The signal strength, mathematically, is proportional to the number of spins inside the slice. That’s why the formula is written as a double integral over x and y within the slice. It represents the total proton density summed across the slice area.

This is a big improvement. Without slice selection, we had a triple integral over x, y, and z — the entire volume. But now, we've reduced it to just two dimensions, a double integral. This means our signal comes only from the selected slice, not the whole volume.

That's a huge step forward. But it's still not the final goal. Because in imaging, we don't just want the sum over the entire slice. We want to know the signal at each individual pixel or voxel. We want to resolve spatial detail — what is the proton density, the T1, or the T2 value at each location. That's what we'll cover in the next part.

slide23:

Now we move to the next major step: sampling in k-space.

This is probably the most important concept in MRI. And I'll be honest, it's also one of the trickiest. So I strongly encourage you to review this multiple times, both in the textbook and in these lectures.

The idea is this: when we add gradient encoding — both phase encoding and frequency encoding — the MRI signal we measure naturally corresponds to a Fourier transform of the object we're imaging.

So the scanner doesn't measure an image directly. Instead, it measures data in k-space, which is essentially the Fourier space of the image. Then, by applying the Fourier transform, we reconstruct the actual image in spatial coordinates.

This concept is formalized in what we call the k-space theorem. And as you'll see, the k-space theorem in MRI is directly analogous to the Fourier slice theorem we studied earlier for X-ray computed tomography.

So, in summary: slice selection gave us spatial specificity along one axis. Rephasing gradients fixed the signal within that slice. And now, sampling in k-space gives us the ability to encode and reconstruct every pixel inside the slice. This is the final piece that allows us to turn raw MR signals into detailed images.

slide24:

Up to this point, we've been talking about signals coming from a whole slice. But instead of just getting a collective signal from all the protons in the selected slice, our goal is to resolve individual locations inside that slice.

We want to know: what is the proton density, often called rho, at a given pixel location? What is the T1 relaxation time for that pixel? What is the T2 relaxation time for that pixel? Once you know all this information, pixel by pixel, you can construct a cross-sectional image. And that is the true purpose of tomographic imaging.

So here's the challenge. Let's pause for a moment and think carefully. We have an equation — you can see it here — which tells us that the signal is proportional to the sum over all spins inside the slice. That's already an improvement. It's much more specific than before, when we were averaging over the whole volume. Now, the signal comes only from one slice.

But here's the puzzle: how can we take this equation, which only tells us the total signal from the slice, and somehow extract the spatial information pixel by pixel? How can we single out, for example, the value of rho, or T1, or T2, at a particular coordinate — say, x naught, y naught?

I encourage you to really think about this problem. Imagine you are designing the experiment yourself. How would you go about resolving the spatial information?

How would you separate the contributions from different locations within the slice?

There are actually multiple ways to approach this. The method I'm going to explain to you is very elegant and very efficient, but it's not the only possibility. So think independently: if you were faced with this challenge, how would you do it?

The solution we use in MRI is to introduce phase encoding and frequency encoding. These tricks allow us to separate the signals spatially. When you see how it works, you'll realize it's both clever and practical. And that's what we'll now begin to explain.

slide25:

Let's start with phase encoding.

Suppose you do nothing special — what happens? All the signals from the slice come out at the same frequency, and they just add together. You can't tell which signal came from which location. It's like hearing a classroom of students speaking at once, all in the same pitch. You know there are many voices, but you can't separate them.

So how do we make progress? We apply a gradient encoding along one direction of the slice. This gradient is turned on for a certain amount of time. What happens then?

At the beginning, before the gradient is applied, all the spins are in phase. But once the gradient is turned on, the magnetic field is no longer the same everywhere. At one location, say at the top, the field is a little weaker. At another location, maybe lower down, the field is stronger.

Because precession frequency depends on field strength, the spins at different locations accumulate different amounts of phase shift over that same time interval.

Where the field is weak, the spins process more slowly.

Where the field is strong, the spins process more quickly.

And since the gradient is linear, the phase difference grows proportionally with position.

So, after the gradient pulse, spins at the top may have accumulated only a small phase change, while spins at the bottom may have accumulated a large phase change, maybe even 180 degrees.

Now, when you collect the signal, you know that spins with little or no phase shift must have come from one end of the slice, and spins with larger phase shifts must have come from the other end.

This is the essence of phase encoding. You haven't fully reconstructed the image yet, but you now have row-wise spatial information. It's like being able to say, "That sound came from the front row, or the middle row, or the back row," even if you don't yet know exactly which student was speaking.

Phase encoding is the first step toward resolving spatial information inside the slice.

slide26:

Now let’s connect this with the equations from your textbook.

As explained in the book, the phase factor is determined by the local processional angular frequency and the total time for which the phase-encoding gradient is applied. Mathematically, it’s the product of frequency and time. But remember, the local processional frequency is proportional to the local magnetic field. And the local magnetic field is not uniform — it’s linearly changing because of the applied gradient.

So we can write the local field as gamma times G y, times y. That is, the gradient strength G y, multiplied by the coordinate y, times the gyromagnetic ratio gamma.

This means the accumulated phase factor depends directly on position y. Spins at different y locations accumulate different phase shifts during the same encoding time. So after phase encoding, if you collect the signal, it still comes from the same slice, but now it’s modulated by this spatially dependent phase factor. That gives you extra information.

To put it simply, now you can resolve the signal line by line across the slice. You’re no longer dealing with one lumped signal. Instead, you can tell which “row” in the slice contributed to the measured data. Of course, this is not yet a full tomographic image. What we want is pixel-wise resolution — knowing the value at every point in the slice. Phase encoding alone only gives us line-by-line information. To complete the picture, we’ll need to add frequency encoding on top of it. That’s what we’ll discuss next.

slide27:

The next important idea in MRI is called frequency encoding. After we finish phase encoding, we do not immediately record any signal. At that stage, the information is already encoded into the spins as phase shifts, but the signal is not read out yet. The next step is to apply what we call the frequency-encoding gradient, often denoted as G-sub-s, or simply the readout gradient. Once this gradient is applied, we begin recording the MR signal.

What happens here is that the local magnetic field is no longer homogeneous. Instead, because of the gradient, the magnetic field strength changes linearly across the imaging plane. On the right-hand side, the field is stronger, and on the left-hand side, the field is weaker. This change in field strength directly affects the processional frequency of the spins. Spins that sit in regions of higher field will process faster, while those in lower field regions will process more slowly. This is a very clever way to extract spatial information from the signal.

Now, when we measure the signal, if we see a component oscillating at a very high frequency, we know that signal must have come from spins located on the right-hand side of the image, where the field is strongest. Conversely, if the signal oscillates at a lower frequency, we know it must have come from the left-hand side, where the field is weaker. In other words, the frequency of the signal directly encodes the spatial position of the spins along the direction of the gradient.

So, the field is no longer uniform across the entire plane. Instead, the gradient introduces a mapping between position and frequency. The faster oscillations in the measured signal correspond to spins located toward the right side, while the slower oscillations correspond to spins on the left side. All of the low-frequency signals come from the left-hand side of the slice, and all of the high-frequency signals come from the right-hand side. By combining this frequency information with the phase information we already stored in the earlier phase-encoding step, we can now determine not only which row the signal came from, but also which column. In this way, we can resolve the precise pixel location of the signal within the slice.

slide28:

Now let's put everything together. The phase angle that you see here carries important spatial meaning. Spins at the top of the slice accumulate a smaller phase angle, while spins at the bottom accumulate a larger phase angle. So when you combine phase encoding with frequency encoding, you can start to separate the signals point by point within the slice. This is the basic idea of how we move from bulk signals to localized, pixel-specific information.

You may wonder, why do we need two more gradient fields? At the very beginning, we already used a slice-selection gradient — let's say along the z-direction — to pick out one slice from the entire volume. But that only gives us localization in one dimension. To fully resolve the slice, we need to further encode along the other two directions: the y-direction for phase encoding, and the x-direction for frequency encoding. By applying gradients step by step, we gradually eliminate one dimension at a time. Slice selection removes the z-dimension, phase encoding handles the y-dimension, and frequency encoding takes care of the x-dimension. The end result is three-dimensional, point-specific localization of the MR signal. This is the foundation of tomographic imaging.

Mathematically, we can describe this process using the equations shown here. If we only apply frequency encoding without phase encoding, the acquired signal is expressed as an integral over the slice that includes a factor of "e to the minus j gamma G-x times x times t." This captures spatial frequency information along the x-direction. But when we combine both phase encoding and frequency encoding, the signal equation becomes the one you see at the bottom of the slide. Here, you have an additional factor of "e to the minus j gamma G-y times y times tau p-e." This accounts for the phase encoding gradient applied in the y-direction for a certain time tau p-e.

Now, looking at this combined equation, you can see it resembles a Fourier transformation. You have terms in both the x and y directions — frequency terms in x, and phase terms in y. Together, they allow us to map spatial coordinates into measurable signal space. This is the mathematical backbone of MRI imaging. We will go deeper into this Fourier interpretation later in the lecture. But for now, recognize that the combination of slice selection, phase encoding, and frequency encoding provides us with true point-wise information. Graphically, as shown here, this is exactly how the imaging sequence works.

slide29:

So here we come to the full sequence that combines slice selection, phase encoding, and frequency encoding. The process begins with a ninety-degree pulse. This radiofrequency pulse flips the magnetization vector, the big M vector, into the x-y plane. Once it is in the x-y plane, the transverse components begin to process, and these processing components induce an alternating magnetic field around the patient or the sample. That alternating field is what we measure as the MR signal.

Now, you want to make sure that you are not flipping every spin in the body. You only want to flip the spins within a single slice. That is where slice selection comes in. During the ninety-degree pulse, you apply a slice selection gradient. This gradient ensures that only the spins that are resonant within that slice frequency range are flipped into the x-y plane. Spins outside the slice are not flipped, and so they do not contribute to the signal. This is the first major step — slice selection removes one dimension, narrowing the signal to a single slice.

Next, you apply phase encoding. On the slide, you see these multiple horizontal lines under the label “G-phase.” This is because phase encoding must be repeated multiple times with different gradient strengths. You can start from a minimum value, then gradually increase step by step, covering a full range of gradient amplitudes. Each phase encoding step imprints a different spatially dependent phase shift onto the spins along the y-direction. Why do we need so many lines? The reason is the Fourier transformation. To fully reconstruct an image, we need information along two degrees of freedom — one for the phase dimension and one for the frequency dimension. Each phase encoding line is like collecting one row of data in Fourier space, which we call k-space.

After that, we apply frequency encoding, shown here as “G-freq.” Frequency encoding happens during the readout. With this gradient on, different spatial locations along the x-direction correspond to different precession frequencies. When you record the signal, you collect a series of oscillations that map into the x-axis of k-space. Each readout provides you with N data points along this horizontal axis.

So, putting everything together: for each phase encoding step, you collect one frequency-encoded readout with N data points. After repeating this process with many different phase encoding steps, you end up with a full grid of data — N by N points. Each row comes from one phase encoding, and each column comes from the frequency encoding readout. Altogether, this maps the entire cross-section of the slice into k-space. Then, by applying a two-dimensional Fourier transform, you reconstruct the final image.

This is the essence of Fourier imaging, or what we call k-space formulation in MRI. Slice selection removes one dimension, phase encoding handles the y-dimension, and frequency encoding handles the x-dimension. The result is a complete two-dimensional image of the selected slice, pixel by pixel.

slide30:

So, if we call the whole term for frequency encoding, G x times t, simply k x, and we call the phase-encoding term k y, then the previous formula we derived can be rewritten in this k-space format. The signal is proportional to the factors k x times x and k y times y, and this immediately shows that the signal we measure is essentially a Fourier transformation. For any fixed time point during readout, the time t sets a specific value of k x, while k y is determined by the strength of the gradient G y and the duration of the phase-encoding step, tau p e. In other words, the frequency-encoding factor is traced out by letting time run during readout, which sweeps continuously along one line in k-space, while the phase encoding only gives one value of k y per line, and this is why we need multiple phase-encoding steps. By repeating the experiment with many different gradient strengths for G y, we can step through multiple values of k y, gradually filling the two-dimensional k-space. This is the basis of Fourier imaging.

On the slide, you can see the formal definitions written explicitly: k x equals gamma over two pi times G x times t, and k y equals gamma over two pi times G y times tau p e. Substituting these into the signal equation, we get that S of k x, k y is proportional to the double integral over the slice of rho of x comma y times e to the minus j two pi times open parenthesis k x x plus k y y close parenthesis d x d y. This is shown as equation four point four nine, and it is the central mathematical statement of Fourier MRI. Notice that in your textbook, the two pi factor is missing inside the exponential, and this is actually a common oversight in many texts, including mine and in my lectures. In reality, the two pi must be included to match the conventional definition of the Fourier transform, so keep that in mind as a small but important detail.

The essential point here is that, through phase encoding and frequency encoding, the data we measure in MRI are already Fourier components. This is very different from X-ray imaging, where the Fourier slice

theorem has to be used: there, each projection profile must be one-dimensionally transformed to draw a radial line in Fourier space, and extra processing is needed. In MRI, the signals themselves directly populate k-space, which means the Fourier information is collected naturally as part of the measurement. Each phase-encoding step contributes to the vertical direction in k-space, and the continuous readout under the frequency-encoding gradient sweeps across the horizontal direction. By combining both, we directly fill k-space, and after applying the inverse Fourier transform, we reconstruct the final image.

slide31:

So now let's talk about how the data are actually collected in the k-space. When time t equals zero, the k value is also zero, which means you are sitting right at the center of k-space. As time increases during the frequency encoding readout, the trajectory moves gradually toward the right-hand side along the k x direction. In this way, you begin to cover half of the Fourier domain. Now remember, for real-valued functions, the Fourier transform has a very useful symmetry property: if you only collect data from one half of k-space, the other half can be inferred mathematically. This symmetry reduces the amount of data we need, but in practice, MRI systems still often collect both halves to account for noise and to achieve cleaner reconstructions.

This diagram shows how the data lines are filled. Each horizontal line corresponds to a readout under the frequency-encoding gradient, which traces across k x. After finishing one line, you step the phase-encoding gradient to a new value, moving up or down in the k y direction, and then record another frequency-encoded line. So, for example, line one at the bottom corresponds to the maximum negative value of the phase-encoding gradient. Then you step gradually upward, line by line, until you reach the maximum positive value of k y. The spacing between these lines is delta k y, and the spacing along the frequency-encoding direction is delta k x. By repeating this process, you fill up the entire k-space grid point by point, line by line.

This is why we call it k-space: each coordinate, k x and k y, corresponds to a Fourier component of the image. Once the k-space is fully sampled, an inverse Fourier transform is applied, and that directly reconstructs the image in the spatial domain. So the essential idea is very simple but very powerful: MRI does not record the image directly. Instead, it records data in k-space, which is really the Fourier space, and then uses the inverse transform to bring everything back into an image we can see.

slide32:

And here you see another important point: full coverage in k-space. If we look at this figure, you can see how we can slightly modify the pulse sequence to make sure the entire k-space is covered symmetrically. Normally, during frequency encoding, you simply start your data collection right away. That means your measurement begins at time t equal to zero, so the k x trajectory also starts at zero and moves only into the positive side. This means, by default, you are not covering the negative part of k-space.

So how do we fix this? Before the regular frequency encoding and data acquisition, we apply what is called a dephasing gradient. This is shown here as G dephase. By applying this dephasing pulse, we effectively shift the starting point of the k x trajectory into the negative region. In other words, instead of starting at zero, all the frequency components now begin with a negative k x value. Then, as time increases during the readout,

the trajectory sweeps across k-space and eventually moves into the positive side. By doing this, you cover the entire rectangular region of k-space, centered around the origin.

This is important because, according to the Nyquist sampling theorem, you need to cover the full bandwidth of frequencies to accurately reconstruct the image. If your sampling region is too narrow, you will miss some Fourier components, and that leads to aliasing or artifacts in the reconstructed image. By carefully choosing the dephasing and rephasing gradients, you make sure that this rectangular region in k-space is wide enough to include all of the significant Fourier components of the object.

So the essential idea here is that MRI imaging works in k-space, and by adjusting the gradient pulses — dephasing first, then rephasing with frequency encoding — you can guarantee full coverage. This is how the k-space theorem is applied in MRI. And remember, in our setup, phase encoding is along the vertical direction, while frequency encoding is along the horizontal direction. Together, they fill the k-space grid, ensuring that the reconstructed image contains all the spatial information of the object.

slide33:

Now let me explain this idea more generally. Here, tau, shown in green, is the time you spend applying the phase-encoding gradient. During this phase encoding, you introduce a phase difference across the spins, but you are not yet collecting any signal. After that, you let time t run, shown in red, and that is when you actually begin collecting the MR signal. This is the formulation with both phase encoding and frequency encoding included.

We first perform phase encoding longitudinally, then we do frequency encoding horizontally. But in fact, you can do this in different ways. You can apply gradients in both the vertical and horizontal directions simultaneously. When you do this, during the time tau, you move the starting phase point diagonally in k-space, depending on the relative strengths of the gradients. If you then start collecting signal immediately, the k-space trajectory begins from that shifted point. Alternatively, you might first move the starting point horizontally by applying one gradient, then switch to a vertical gradient and move to another position, and only then begin collecting the signal. Each different choice changes the trajectory you trace in k-space.

For example, if you use an x gradient for frequency encoding, then as time t increases from zero, the k x value changes continuously toward positive values. If instead you first apply a horizontal gradient as a phase-encoding step, the initial point shifts along the k x axis. Then, by applying a vertical gradient, you move the starting point vertically. Once you begin signal acquisition, you trace a horizontal or vertical line in k-space. Finally, if you apply both gradients at once and immediately collect the signal, you start at the origin and move along a straight line at a slope determined by the ratio of the two gradients.

The essential idea is that whenever you collect a signal with a gradient on, the frequency of precession is shifted, and that frequency determines the trajectory in k-space. If you apply a gradient without collecting a signal, the spins accumulate phase instead. By combining these options, you can navigate k-space in virtually any pattern. This flexibility is what allows MRI to encode spatial information and build an image slice by slice, pixel by pixel.

slide34:

At this point, you might be wondering — have you understood this, or are you confused? Honestly, if you have never studied this carefully before, it is very natural to be confused. Even for me, when I first studied this, I had to go through it several times, very carefully, before I felt I understood it clearly.

So let's check ourselves. If you basically follow the logical flow so far, you can feel confident. If you are quite confused, that's also expected — because these concepts are tricky and require repeated review. If you're somewhere in between, that's normal too. The important thing is not to get discouraged. What I will do now is go over the same ideas again, but from a slightly different angle, with more detail and more illustrations, so that the picture becomes clearer.

slide35:

Now let's revisit spin echo with another illustration. Spin echo is a very clever but also highly simplified signal model. In this figure, you can see the sequence step by step. You begin with the magnetization vector aligned along the z-axis. Then you apply a 90-degree pulse, which flips the vector into the transverse plane — here shown along the y-axis, though it could just as well be flipped along the x-axis depending on how the pulse is applied.

Once in the transverse plane, the spins begin to dephase because of local field inhomogeneities. Some spins process slightly faster, and others slightly slower. Over time, this produces a spread in phase, as shown here. Then comes the magic step: a 180-degree pulse is applied. This pulse flips the spins over, so that those that were lagging are now ahead, and those that were ahead are now behind. From that point forward, the phase differences begin to refocus. After the same amount of time has passed, the spins come back into alignment, forming an echo.

This is the echo signal that can be detected by the coil. Importantly, this echo arises from the selected slice, assuming you have already applied slice selection. The equation shown here, equation 4.43, expresses the signal mathematically as being proportional to the integral of the proton density over the slice. In other words, the echo represents the combined contribution of all the spins in that slice, after refocusing has occurred.

So spin echo illustrates in a very clear way how we can recover useful signal despite dephasing, and how gradients and RF pulses can be combined to control the evolution of magnetization in MRI.

slide36:

And now, let's talk about the actual signal we can measure in MRI, which is called the free induction decay, or FID. After applying a ninety-degree RF pulse, the net magnetization vector is flipped into the transverse plane. This transverse component of magnetization is constantly rotating around the main magnetic field, B zero. Because of this precession, the changing magnetic field induces an alternating electromagnetic field. And as you learned in high school physics, an alternating magnetic field induces an electric current in a nearby coil. So if we place a receiver coil close to the patient, that coil will detect this alternating signal.

However, the signal does not last forever. It decays over time, and this decay is due to the T2-star relaxation. The spins gradually lose phase coherence, and the signal amplitude decreases exponentially. What we record is not a constant sine wave of fixed amplitude, but a damped oscillation that fades away

with time. This decaying signal, directly measured from the whole sample, or from a whole slice if slice selection is applied, is the free induction decay, or FID. This is the fundamental NMR signal we detect in MRI.

slide37:

Now, let's look at how we time the acquisition of this signal. Two very important parameters in MRI are TE and TR. TE stands for echo time, which is the time from the application of the RF pulse to the moment we collect the peak of the echo signal. TR stands for repetition time, which is the time between successive RF excitations. You see in this diagram that we often use a ninety-degree pulse followed by a one-eighty-degree refocusing pulse to generate a spin echo. The echo forms at time TE. TR is measured from one excitation pulse to the next.

These two parameters, TE and TR, are under our control. They are not properties of the tissue itself but technical choices we make during the imaging sequence. By carefully selecting TE and TR, we can emphasize different tissue contrasts, highlight T1 weighting, T2 weighting, or proton density weighting. And at the same time, we can also balance image quality with scan time. If TR is long, the magnetization has more time to recover, giving us strong signals, but the scan will take longer. If TR is short, scanning is faster, but some magnetization may not have fully recovered, which reduces the signal. TE works similarly: shorter echo times reduce T2 contrast but keep signal strong, while longer echo times allow more T2 differences to emerge but at the cost of weaker signals. These parameters, TE and TR, are the main levers we adjust to create different kinds of MRI contrast.

slide38:

And I said that you can collect a signal with a nearby coil. Then, if you place another coil, you can also get a signal. Remember, the horizontal component of the magnetization is a rotational vector. So the signal actually has two parts. It depends both on amplitude and on phase. The signal detected along the x-axis and the signal detected along the y-axis naturally have a ninety-degree phase difference, because the magnetization keeps rotating. That makes the signal a complex-valued quantity.

So, let's describe it mathematically. You can record Sx, which is modulated as M zero. After a ninety-degree pulse, this transverse signal is subject to spin echo formation and also subject to T2 decay. The Sx term represents the amplitude and the frequency content along the x-axis. Likewise, you have a Sy component, which is just the complementary sine term. Together, these two components give you both cosine and sine contributions. If we put them into compact form, we can write the signal as a complex number: a real part and an imaginary part. This gives us a very compact, complex-valued signal representation. In this notation, I have chosen to place the imaginary part corresponding to Sy. This is just a convention. The coding system can be chosen differently. Here, I simply make Sy the imaginary component. But for physical measurement, both the real part and the imaginary part are equally meaningful. We use complex notation mainly for convenience.

This idea is similar to what we saw in Fourier series and Fourier transforms, where using a complex exponential makes the notation compact, but the actual measured signal is real. Any complex form can always be converted back into a purely real form if needed.

Now, the processional frequency is omega zero, determined by the main field B zero. The transverse magnetization keeps processing at this angular frequency omega zero. If we use a rotating frame—meaning

we mentally rotate our coordinate system at the same frequency as the spins—then the signal appears stationary. In the rotating frame, the oscillatory part is removed, and the signal just decays exponentially with T2. This makes the description much simpler.

So, this is the signal in the rotating frame: stationary, but still decaying, due to T2 relaxation. And remember, T two decay is a physiological property. We cannot control it; it reflects interactions within the tissue itself.

Now, what about repetition? The MRI sequence is repeated after a certain repetition time, TR. If TR is not long enough, the longitudinal magnetization, M zero, has not fully recovered. T1 relaxation determines how much of the original longitudinal magnetization is available at the start of the next cycle. So, after TR, the signal amplitude depends on both T1 recovery and T2 decay.

The total signal at time t is therefore a combination of these factors: proton density, hidden inside M zero; T1 recovery, which determines how much longitudinal magnetization has returned; and T2 decay, which determines how fast the transverse signal fades. Together, these biological properties determine tissue contrast.

On top of that, we have technical parameters: B zero, the main magnetic field; TR, the repetition time; and TE, the echo time. These are under our control. They don't reflect patient pathology, but our technical choices in the scan.

Now, why do we need repetition in the first place? Recall that in phase encoding, each time you collect data, you only get one line in k-space. Frequency encoding samples along that line, but since k-space is two-dimensional, you need multiple phase-encoding steps. That means repeating the sequence many times. This is very much like CT scanning, where helical scanning or fan-beam scanning requires repeated projections. In MRI, k-space must be filled line by line, and that requires repetition.

If TR is very long, M zero is fully recovered between repetitions. That gives you clean, strong signals, but scanning takes longer. If TR is too short, M zero hasn't fully recovered, so the signal is weaker. But shorter TR also means faster scanning. So there is a tradeoff: longer TR improves image quality, shorter TR reduces scan time. The goal is always to balance scan speed with image quality for the clinical application.

Finally, both TR and TE can be adjusted to emphasize different tissue properties. A short TR and short TE give you T1-weighted images. A long TR and long TE give you T2-weighted images. Proton density weighting lies somewhere in between. So, this signal model for a ninety-degree pulse captures both the biological properties of tissue—proton density, T1, and T2—and the technical factors—B zero, TR, TE, and time. Together, they shape the MRI signal and ultimately determine the image contrast.

slide39:

And you can use short TR or long TR. You can use short TE or long TE. And with different combinations, your image can reflect more about rho, about T1, or about T2. So with different combinations, you measure different parameters. This is not like CT, where you only measure the attenuation coefficient, mu. In MRI, you have T1, you have T2, and you have rho, the proton density. If you just do one measurement, one kind of image, then the image really reflects all of these parameters mixed together. But if you purposely want to emphasize one of the three parameters, you need to select the proper technical parameters.

So first, let me explain what I mean by long TR. TR is the time for repetition. If you wait long enough after each excitation, the m vector will return to the original position along the z-axis. So M0 aligns with the z-

axis. That is what we mean by long TR. Long TR means the magnetization has had enough time to fully recover. Short TR means you do not give the system enough time, so the full value of M0 has not yet been recovered along the z-axis. That's the meaning of long TR and short TR.

Now let's talk about TE, the echo time. What do we mean by short TE? Short TE means the echo time is very short. The spin echo is collected quickly, so the flipped M vector doesn't have enough time to dephase. Because the echo is measured soon after flipping, the dephasing is not serious. So M xy, the transverse component after the 90-degree flip, has had very little time to dephase. That means there is no significant T2 effect. If you measure the signal with a short TE, you cannot tell how much T2 contributes, because T2 hasn't had enough time to show its effect.

Now, if you use a long TR, then the magnetization has fully recovered. That means you have no information about T1, because T1 does not play a role when you wait long enough. If you use short TR, the recovery along the z-axis is incomplete. In that case, the effect you observe depends on the degree of recovery of the longitudinal component, so you see the T1 effect. That's why short TR with short TE is called T1-weighted imaging.

Next, if you use long TE, the echo time is long enough for T2 decay to take effect. And if you also use long TR, so the longitudinal magnetization has fully recovered, then there is no T1 effect. In this case, the signal is dominated by T2 decay, so it is T2-weighted. That is shown here with the green arrow.

Now, if you use long TR and short TE together, then there is no T1 effect, and no T2 effect. In that case, the signal is proportional to rho, the proton density, because you are basically measuring the full magnetization without either T1 or T2 altering it. That is called rho-weighted imaging.

There is also a combination shown here with short TR and long TE, but that is not possible in practice. Remember, TR is the repetition time, and TE is the echo time within that repetition. TE normally should not be longer than TR. So this case does not exist physically.

So, to summarize: short TR and short TE give you T1-weighted imaging, long TR and long TE give you T2-weighted imaging, and long TR with short TE gives you rho-weighted imaging. Again, this takes time for you to fully understand. And if this is your first time seeing it, it is natural to feel a little confused.

slide40:

Now, let's take a look at some real examples of brain images. On the same slide, you can see three different types of images: rho-weighted, T1-weighted, and T2-weighted. In MRI, rho here means proton density weighting, not P-weighted. In the brain, there are mainly five important tissue types that we care about: grey matter, white matter, cerebrospinal fluid or CSF, blood, and water. If you look carefully at the table, you'll notice that the T1 and T2 relaxation times are quite different for each tissue.

For example, grey matter has a T1 of about 1 second and a T2 of about 0.1 seconds, while CSF has a much longer T1 of around 2 seconds and a T2 of around 0.25 seconds. Water has even longer values, T1 about 4.7 seconds and T2 about 3.5 seconds. These differences are the key to MRI contrast. The very same slice of the brain can look dramatically different depending on whether you emphasize proton density, T1, or T2 weighting.

So, depending on your imaging parameters, you can choose to highlight different tissue characteristics. This is something CT cannot do, because CT only measures the attenuation coefficient mu. But MRI gives us

much richer information because of rho, T1, and T2. And for a special case, we can use Gx gradient encoding to separate frequencies along the x-axis. That allows us to further refine how we differentiate tissue types based on their signal behavior.

slide41:

Here we are looking at a special case: Gx gradient encoding. That means we apply a linear gradient along the x-axis, so that the precession frequency depends directly on the coordinate x. In this case, the signal we observe combines several effects: proton density, T1 recovery, and T2 decay.

You can see the equation here: the amplitude of the local signal depends on rho, multiplied by the term for partial recovery due to T1, and multiplied again by the decay due to T2. This gives the overall signal strength at each location. Then, with the Gx gradient, we impose a position-dependent frequency shift. That's why the exponential factor here involves Gx times x.

So, summarizing: the total signal is complex, but you can think of it as having two parts — the amplitude of the local signal, determined by rho, T1, and T2, and the frequency of the local signal, determined by the gradient. This separation is very useful because it tells us what portion of the signal comes from tissue properties, and what portion comes from our encoding scheme. This prepares us for data collection in k-space.

slide42:

Now let's move this into the k-space picture. We define a function f of x and y, which includes rho, the proton density, multiplied by the T1 recovery term and the T2 decay term. This function describes the true tissue properties in space. Then we define kx as gamma over 2 pi, times Gx, times t minus TE. So as time increases, kx traces out a trajectory in Fourier space. The signal of t, the thing we actually measure, is the Fourier transform of f of x and y, evaluated along the kx axis.

So in this case, we are scanning a single line in k-space, the horizontal line, because the gradient field is only along x.

This is a very special case, but it makes the principle clear: our signal directly corresponds to the Fourier components of the object we are imaging. That's why MRI data is often first collected in k-space and then reconstructed by inverse Fourier transform to get the final image.

slide43:

Now let's consider the general case. Instead of just Gx, we allow general gradients along x, y, and z. These gradients can vary with time, and their effects accumulate as the spins process. The local precession frequency is determined by the inner product of the gradient vector G of t and the position vector r. When we integrate this over time, we get a phase factor that builds up gradually. This leads us to define k of t as gamma over 2 pi times the time integral of G of tau d tau.

In other words, k of t is a trajectory through Fourier space, controlled entirely by our choice of gradient waveforms. The measured signal s of t is then the Fourier transform of the object's spin density, f of x, y, z, evaluated along this k-space trajectory. This is the k-space theorem. It tells us that by designing gradient

waveforms, we can sample any path we want in k-space — straight lines, zig-zags, spirals, anything. That's why MRI is so flexible. This general Fourier space formulation is extremely powerful because it connects what we measure directly to the spatial distribution of spins inside the patient. And again, by inverse Fourier transforming the collected data, we reconstruct the image.

slide44:

So at this point, I need to emphasize that this formulation is only approximate. Where exactly do the approximations come in? There are really three places.

First, if you look at this expression, you see that the function depends on time, t. But when we talk about the k-space theorem, we treat it as if the whole function is fixed in space, something like F of x, y, z, and not time-varying. In reality, there is a time dependence here, and when we ignore that, we are making an approximation.

Second, you see this vector r, the positional vector. We treat it as if it is constant, fixed in space. But in reality, spins may move. For example, spins in blood keep moving. Patients may also move. Even molecular motion introduces some effects. But in this model, we ignore those motions and assume r is constant. That is another approximation. And third, the factor related to signal decay.

In practice, the signal we measure is subject to T2 star decay. But in this approximation, we assume that the readout is fast enough, so that no significant T2 star decay occurs during acquisition. That is a third approximation. Under these assumptions, the formulation reduces neatly to a Fourier transformation, which makes image reconstruction straightforward. But if you want very accurate imaging, all these neglected factors must be included. And once you include them, the reconstruction is no longer a simple Fourier transform. It becomes much more complex.

In fact, this is an area where machine learning and advanced reconstruction methods can play a big role. In the last part of the course, I will mention how modern approaches are being used to solve these challenges.

slide45:

Now, let me give you a simple example to make this clearer. Under the three approximations we just discussed, we can measure the Fourier space information. The figure on the left shows what is recorded in k-space. Remember, the Fourier transform is a complex function. It has both amplitude and phase. But what we display here is only the amplitude. It looks something like this, a distribution of intensities in k-space.

Then, if you apply an inverse Fourier transformation, you obtain the image on the right, which is a cross-sectional image of the brain. So with MRI, by collecting data in k-space and applying an inverse Fourier transform, we can reconstruct detailed cross-sectional anatomy. I think this is really an amazing achievement — especially when you compare it with X-ray CT. MRI provides a completely different way of obtaining cross-sectional images, not from attenuation, but directly from signals in Fourier space.

slide46:

So now we come to the final part of this lecture. You see, we have obtained cross-sectional images using MRI. The last slides summarize the key points and also point to the MRI scanners themselves. I will not go

into the engineering details of scanners here, since these are mostly descriptive and you can easily read them from the textbook.

The textbook straightforwardly explains the scanner components. If you read carefully, you will find it both easy and interesting to understand. You will also find additional material about spin echo sequences, gradient echo, and slice selection, which can deepen your understanding. These are not required for the core lecture, but they are excellent supplemental resources if you want to learn more. And finally, you will have some homework related to this lecture, so that you can review the key ideas and practice the concepts we discussed.

slide47:

This diagram shows a block representation of how an MRI system works, starting from signal generation to image display. At the heart of the scanner, you have the patient placed inside the bore surrounded by gradient coils and the RF, or radiofrequency coil. The RF coil serves two purposes — it transmits the excitation pulses into the body, and it also receives the weak signals emitted back from the tissue. To generate those pulses, a frequency synthesizer and an RF amplifier create the correct radiofrequency signal. That signal is then sent through a transmit–receive switch, which ensures that the same RF coil can switch roles between sending and receiving.

Once the RF coil receives the returning signal, it is very weak. So the signal first goes through a preamplifier to boost it. Then, it passes through a receiver blanking unit to protect the electronics during transmission. After that, the signal is routed to an RF amplifier and a demodulator, followed by a quadrature mixer, which separates the signal into its real and imaginary components. These components represent the complex-valued MRI signal.

The analog signal is then converted into a digital form by an analog-to-digital, or A/D, converter. From there, the digital data is stored temporarily in RAM memory and controlled by the host computer. The computer performs the essential reconstruction step — the two-dimensional inverse Fourier transform, or 2-D IFFT. This is what converts raw k-space data into an actual image that can be displayed.

On the control side of the system, there are waveform generator timing boards and gradient amplifiers. These control the gradient coils inside the scanner, which are responsible for slice selection, phase encoding, and frequency encoding. The pulse programmer and software interface, also controlled by the computer, dictate exactly when and how these gradients and RF pulses are applied. Together, this timing sequence ensures precise spatial encoding of the MRI signal, allowing us to reconstruct high-resolution images.

So overall, this diagram gives us a complete picture: from the RF system and gradient amplifiers, to signal detection, digitization, Fourier reconstruction, and finally, the image display. It ties together all the engineering components that make MRI imaging possible.

slide48:

This 3D rendering gives us a cutaway view of how an MRI scanner is structured and how a patient fits inside the system. At the center, you see the patient lying on the patient table, which slides into the bore of the scanner. Surrounding the patient are several key components. Closest to the body are the radio frequency

coils. These coils are responsible for transmitting the RF pulses that flip the spins inside the tissue, and also for receiving the weak signals emitted back from the patient. Without these coils, no imaging would be possible.

Just outside the RF coils are the gradient coils. These coils are used to create small, linearly varying magnetic fields on top of the main magnetic field. By turning the gradient coils on and off in different directions, the scanner performs slice selection, phase encoding, and frequency encoding, which are the core steps that allow spatial localization of the MRI signal.

Encasing the gradient coils is the large superconducting magnet, shown here as the main cylindrical structure of the scanner. This magnet produces the strong, uniform field we call B-zero, typically one and a half Tesla, three Tesla, or even higher in advanced systems. The magnet is what aligns the spins in the body, creating the initial conditions necessary for MRI.

All of this hardware is integrated into the scanner body, which you see as the outer blue casing. The patient is positioned on the table and moved smoothly into the bore for imaging. This design allows precise alignment of the body part of interest with the center of the magnet, where the field is most uniform.

So in summary, this diagram shows the layered structure of an MRI scanner: the patient table for positioning, the RF coils for transmit and receive, the gradient coils for spatial encoding, and the large magnet that provides the powerful, stable field at the heart of the system.

slide49:

Here we see a modern clinical MRI scanner. This particular system is the Sparkler 1 point 5 Tesla MRI, which was made through a joint venture between Philips and Neusoft in 2008. What you notice first is the large cylindrical magnet housing, which contains the superconducting magnet that generates the powerful and stable magnetic field, known as B-zero. That main field is what aligns the spins inside the patient's body and makes MRI possible.

At the front of the scanner is the bore opening, with the patient table extending outward. The table is designed to move smoothly into the bore, positioning the patient precisely at the isocenter of the magnet, where the magnetic field is most uniform. Inside the bore, hidden from view in this picture, are the gradient coils and the radio frequency coils, which together perform the tasks of spatial encoding and signal detection. The control panels you see on either side allow the operator to manage some of the basic functions and safety interlocks.

This specific 1 point 5 Tesla strength is one of the most common clinical field strengths worldwide, balancing good image quality with manageable cost and patient comfort. It represents an industry standard for many diagnostic applications, including brain, spine, and body imaging. So this photo gives you a sense of what an actual MRI scanner looks like in the clinic, as compared to the system diagrams and cutaway renderings we discussed earlier.

slide50:

Here we see a PET-MRI scanner, which is an advanced hybrid imaging system. On the left side of the slide, there is a schematic diagram showing how the different components are arranged, and on the right side, you see an actual Siemens PET-MRI machine.

The MRI part of the scanner is built around the large cylindrical magnet, just like a conventional MRI. The magnet shielding coil, the primary magnet coil, and the gradient coils are all stacked concentrically, forming the core structure of the system. These generate the strong static field and the spatially varying gradient fields needed for magnetic resonance imaging. Inside this same system, PET detector modules are integrated. These include layers such as the PET camera, avalanche photodiodes, LSO crystals, and associated electronics like pre-amplifiers and driver boards. Together, these PET modules detect the gamma rays emitted by radioactive tracers inside the patient's body.

An RF head coil is also shown in the diagram, which is part of the MRI subsystem, used to transmit radiofrequency pulses and receive the MR signal. The magnet cryostat, shown as the outer casing, maintains the superconducting coils at very low temperatures so that the magnetic field stays stable.

On the right-hand side, the Siemens clinical PET-MRI scanner looks very much like a conventional MRI machine. But the key difference is that the PET detectors are integrated into the bore, allowing simultaneous acquisition of MRI data and PET data. This means that while MRI provides excellent soft-tissue contrast and anatomical detail, PET simultaneously provides functional and metabolic information. The result is a powerful hybrid modality that combines structural and functional imaging in a single scan.

This PET-MRI approach is especially important in neurology, oncology, and cardiology, where you want to see not just anatomy but also how tissues are functioning or metabolizing in real time.

slide51:

This photo shows a preclinical MRI scanner, specifically the Bruker BioSpec Avance III 94/20 system, which operates at 9.4 Tesla. Unlike clinical MRI scanners that are designed for human patients, preclinical MRI systems are mainly used for small animals such as mice, rats, or rabbits. These scanners are critical in biomedical research because they allow scientists to study disease models, monitor treatment effects, and investigate basic biological processes in a controlled setting before moving to human trials.

The scanner looks similar in shape to a clinical MRI, with a large cylindrical magnet and a bore in the center, but the bore is much smaller since it is designed for animal studies. The extremely high magnetic field strength of 9.4 Tesla gives very high resolution and sensitivity, allowing researchers to visualize fine structural details that would not be possible at lower field strengths. With this scanner, scientists can measure anatomical structures, functional activity, and even molecular processes, all non-invasively.

Preclinical MRI scanners like this one play a vital role in translational medicine, serving as a bridge between laboratory discoveries and clinical applications. By studying animal models of diseases such as cancer, stroke, or neurodegenerative disorders, researchers can optimize imaging protocols, test new therapies, and gather critical data before applying these methods in human patients.

slide52:

His photo shows a modern CT-MRI guided radiation therapy system. In radiation oncology, precise targeting of tumors is critical because you want to deliver a very high radiation dose to cancer cells while sparing as much of the surrounding healthy tissue as possible. Traditionally, CT has been used for radiation planning because CT provides good information on patient anatomy and electron density, which are needed to

calculate radiation dose. However, CT has limited soft-tissue contrast, which means tumors in organs like the brain, liver, or pelvis may be difficult to outline clearly.

By combining MRI with CT in radiation therapy, doctors can take advantage of the strengths of both modalities. MRI provides excellent soft tissue contrast, allowing the tumor to be visualized more clearly, while CT provides the necessary density information for radiation dose calculation. This hybrid approach makes the targeting more accurate and safer.

In this photo, you see the treatment couch where the patient lies, with immobilization devices to keep the patient still. The large circular gantry houses the imaging and radiation components. During treatment, the system acquires imaging data to confirm the tumor's position, and then the linear accelerator delivers a precisely shaped beam of radiation to the tumor. This technology represents a powerful integration of imaging and therapy — using CT and MRI together to guide radiation treatment in real time, which improves tumor control and reduces side effects for patients.

slide53:

This slide shows a potential MRI-guided radiation therapy design proposed by Xun Jia and colleagues. What you see here is a concept where the MRI magnet is split into two halves, with a separation of about 70 centimeters. That gap allows room for radiation beams to pass through while still maintaining the magnetic field needed for MRI imaging. The design also includes a couch that can rotate by plus or minus twenty-five degrees, giving flexibility in patient positioning and beam delivery angles.

The circular ring you see provides both mechanical stability and magnet shielding. This is very important because strong magnetic fields need to be contained and stabilized to prevent interference with surrounding equipment and to protect patient safety. In addition to the MRI magnet, the design also considers adding a cone-beam CT, or CBCT, mounted on the same ring. Although not shown in this illustration, the CBCT would provide complementary X-ray imaging, making this a hybrid system that combines MRI's excellent soft tissue contrast with CT's accuracy for dose calculation.

Clearly, this figure illustrates an innovative idea in radiation oncology — a split-magnet MRI system, integrated into a rotating gantry, designed to guide precise radiation treatment while maintaining real-time MRI imaging of the tumor.

slide54:

This slide is a list of extra resources for you to explore outside the lecture. These links are short videos that explain some of the key concepts we discussed in more detail.

The first group of links is about the spin echo. Spin echo is one of the most fundamental sequences in MRI, so watching those demonstrations will help reinforce how a ninety-degree pulse and a one-eighty-degree pulse work together to refocus the signal. The second group of links is about the gradient echo. Gradient echo is a variation where gradients are used to generate echoes instead of a one-eighty-degree pulse. These videos will show you the difference in how the signal is produced. Finally, the last link is about slice selection. As we discussed, slice selection is how we target a specific region of the body by combining a magnetic gradient with the RF pulse.

Note that although these videos are not required, they are helpful if you want to review the concepts visually. They can give you an extra layer of understanding, especially if you are new to MRI physics.

slide55:

This slide gives you the homework assignment for this section. You can see that it has three problems, numbered 4.7 through 4.9. In problem 4.7, you are asked to look at the five frequencies shown in the figure above and then state the order of the T1 and T2 relaxation times. For example, you might write something like T1 of brain is greater than T1 of CSF, which is greater than T1 of aqueous humor. The goal here is to compare how different tissues relax under different frequencies.

In problem 4.8, you are working with hydrogen nuclei in fat and water. You are told that the T2 value of fat is 100 milliseconds and the T2 value of water is 500 milliseconds. In a spin echo experiment, you need to calculate the delay time between the ninety-degree pulse and the one-eighty-degree pulse that maximizes the difference in signal between fat and water. This problem makes you think about how relaxation differences translate into image contrast.

Finally, in problem 4.9, you are asked to write an expression for the longitudinal magnetization, M z, as a function of time after a one-eighty-degree pulse. You should then determine the time at which M z becomes zero. After that, you need to plot what happens if, instead of a one-eighty-degree pulse, you apply a one-thirty-five-degree pulse. This problem connects the mathematical model to how pulse angle changes affect the signal.