slide1:

Last time, we explained to you the physical and mathematical principles for ultrasound imaging.

In this lecture, we will be more specific about how we can form different kinds of ultrasound images. We're moving from principles to practice—how we send sound, receive echoes, and turn those echoes into the different image types for various clinical tasks.

slide2:

I feel quite good today, moving into the second part of ultrasound imaging. After this lecture, we'll teach the last imaging modality —optical imaging, in one lecture.

Now, we'll follow the green textbook fairly closely. I strongly recommend that you read the chapter on ultrasound imaging — it's about twenty-three pages long. If you read through the chapter carefully, you'll have a very good understanding of everything we're covering. I encourage you again to review or re-review that chapter, because it will really help you master today's lecture.

slide3:

So here, I'm explaining the key points — the way I think helps you summarize what really matters. But remember, you still need to read the textbook. It's not that you can just sit back, listen to my talk, and skip the reading. No, you really have to read the textbook. That's how you build a good memory and a solid foundation. Two ultrasound imaging lectures in a classroom is not enough if you only listen passively. What I'm giving you here is a general yet quick descriptioin of ultrasound imaging methods — and specifically, how we actually form images.

Now, this topic can be divided into two main parts. The first part is about imaging modes. We'll talk about different types of ultrasound imaging modes — basically, A-mode, B-mode, C-mode, and M-mode, sometimes also called I-mode. What do these mean? Each mode represents a different way of collecting and displaying ultrasound data. There are multiple ideas behind them, and one very important concept that the textbook doesn't explain in detail — but I believe you really need to know — is called Huygens' principle. I checked the pronunciation; it's "Huygens' principle." This is a very important physical idea. Understanding this principle will give you a much deeper grasp of how array-based scanning and beam forming work. It shows how you can focus ultrasound beams at different depths and times, giving you tremendous flexibility in image formation.

In the second part, we'll move on to Doppler imaging. This part deals with dynamic image formation — it allows us to measure both the speed and the direction of motion, such as blood flow in vessels or motion of the heart chambers. Doppler imaging can be implemented in several ways: you can use continuous ultrasound waves, or you can send out sequences of pulses. It's a very important imaging mode. We'll also discuss echo correlation and color Doppler. All these topics are explained clearly in the textbook, so please read that section carefully to strengthen your understanding.

Finally, we'll look briefly at image artifacts and at some new things — this last part is more for fun and inspiration. We won't go into all the advanced new ideas in detail. The first three parts — the imaging modes, Huygens' principle, and Doppler imaging — are what you'll be tested on. So, let's begin by asking: what exactly is the A-mode?

slide4:

(No notes for this slide)

slide5:

Now let's move on to the second one — B-mode. B stands for brightness. That's the keyword to remember here. B-mode is a very powerful imaging technique. You can think of it as an extension of A-mode — instead of a single amplitude trace, you're now using multiple A-mode scans to create a two-dimensional image.

Here's how it works: in A-mode, you send one ultrasound pulse into the tissue and receive echoes from a single line of depth. In B-mode, you repeat that process at many different positions. You scan line by line — one A-mode here, another next to it, and so on. By combining all these one-dimensional signals, you form a two-dimensional image that shows amplitude, or intensity, as brightness. So, what you finally get is a picture — an intensity map of the internal structures.

You're not limited to scanning straight, parallel lines either. Just like in X-ray CT, you can use a fan-beam geometry — instead of sending one beam directly forward, the transducer sweeps across an angle, covering a larger area. This sweeping motion allows you to reconstruct a full cross-sectional view of the tissue. I'll show you another short movie clip to give you a better visual sense of this process.

Now, let's look at the movies. First, you see the A-mode. Here — look carefully. The pulse on the right is the ultrasound wave being transmitted into the biological tissue. The blue dots represent backscattered echoes — reflections from different structures inside the tissue. Notice that the same transducer both sends and receives the signal. So, when the pulse goes out, the transducer then listens for the returning echoes.

Let's play that again. Watch the blue dots — those are the echoes. You can see how they appear as small reflections. The tail of the pulse gets weaker as it travels deeper — that's where you receive the last few echoes. This gives you a one-dimensional signal profile, which is what we call A-mode.

Now, look at the next one — this is B-mode scanning. Each line in the image corresponds to one A-mode scan. When you stack many A-mode scans side by side, you form a two-dimensional brightness image. This example shows a phantom scan. You can see the B-mode image — it's not as crystal clear as a CT or MRI image. There are speckles, some noise, and scattering effects. But you can still recognize the major anatomical structures.

Ultrasound transducers are cost-effective, small, and compact — that's one of their biggest advantages. They can operate at high speed, producing real-time imaging. That's B-mode scanning — simple, efficient, and very practical for clinical use.

slide6:

Alright, let's move on. So, we've talked about A-mode and B-mode — now we come to the next one: M-mode. The letter M stands for motion.

M-mode is very similar to A-mode, but with one important difference. In A-mode, you send a pulse and get one amplitude profile at one point in time. In B-mode, you scan across space — one A-mode line after

another — so the difference between two A-mode profiles represents a change in spatial location. You're scanning across different positions or angles to build an image.

In M-mode, however, you stay at the same location — the same angle — and you repeatedly collect ultrasound signals over time. So instead of scanning across space, you're scanning across time. What changes here are the internal structures that move, like the walls of a vessel or a heart chamber.

For example, imagine selecting a region of interest, such as part of the heart wall. As the heart beats, that wall expands and contracts. Even though the probe stays fixed, the echoes you receive from that region change over time. Each A-mode measurement gives you a one-dimensional profile, but now, because of motion, those profiles vary with every pulse. Sometimes the wall moves closer to the transducer; other times, it moves away.

If you keep recording these signals, you can build a motion profile. You can still display it as a two-dimensional image: the vertical axis represents depth — just like a typical A-mode — and the horizontal axis represents time. The result is a real-time record of motion. This is what we call motion mode, or M-mode.

It's especially useful for studying moving organs, such as the beating heart, where you can visualize the rhythmic motion of cardiac walls and other dynamic structures over time.

slide7:

So, the last one—and I would say—is C-mode, C for coronal image. You know, three-dimensionally, if you cut a patient with a virtual plane this way, we usually call that the transverse plane. If you cut the patient this way, you call it sagittal. But if you cut the patient this way—think I am facing you and you cut like this—that is called the coronal image. So you have three views. I will show some images.

You do A-mode imaging along this direction. You have all the information, and the spatial dimension along this direction is your coordinate system—just one axis. But you can do some echo-gated imaging. You send the pulse in that direction, you wait a certain amount of time, then you see—within a small time window—if the echo comes back. Given the delay time, you really target a specific depth. So from this one line, you roughly calculate: the ultrasound goes a round trip from the transducer to the plane and back. Suppose this round-trip time is, for example, 100 microseconds—let's say 100 milliseconds in my example. You only see the amplitude at 100 microseconds. That is the round-trip echo time—the echo goes here, hits the plane, comes back—so, 100 microseconds.

This 100-microsecond sample is reported as your pixel value. Then, on this two-dimensional coronal plane, you keep scanning—line by line, pixel by pixel. You make sure the round-trip time is targeted to that depth position. You put all these amplitude values together, and you form a two-dimensional image. That is the so-called coronal image.

Here is an example: you have a hand here. You put the ultrasound into the hand, and you record the echoes. With all the timing coordinated right, you form an internal image that cuts through the hand. You can see certain pieces of bone. There is some terminology I just learned—metacarpals—these light-blue bones; you can see these light-blue bones. The transducer is operated in C-mode here. There are some good ideas; details are in this article.

You have the image—the image is really formed here. Then the image in a pixel here is reflected by a half-silvered mirror and viewed by the reader or doctor. Through this half-silvered mirror, you see the patient's

body, and the image you see is really superimposed onto the patient's internal structure. It gives you the illusion that you see into the patient directly—you see the beating heart—and you can maybe use this as guidance to put the catheter there, make a blood vessel open wider.

This is a good idea. It kind of works with augmented-reality things, and it's a good idea. But the C-mode has nothing to do with this viewer. C-mode can be understood only with this sub-figure—so that's the idea. And ultrasound imaging so far is straightforward. I have explained to you A, B, C, M—so you see later, a word you need to remember.

OK, this is a single-crystal–based transducer that can be used to do all these kinds of things. But when you want to form two-dimensional images, you need an easy scan, or you need to do multiple times—so you have a time dimension. With the time dimension, with a single transducer, you can still form two-dimensional images. Certainly, this is quite a lot like pencil-beam X-ray imaging. In the first generation, you have a pencil beam of X-rays and a single detector; you can do a scan. That is very inefficient.

slide8:

So that's why we say: we have to use ultrasound transducer arrays. This will allow you to do parallel data acquisition. Parallel data acquisition is always desirable. And in the case of clinical—I said chiral—scenarios, faster imaging speed will reduce most artifacts and improve the report. So, a lot of good points to make.

There are many types of ultrasound transducers. You have a linear array; you can use it in different ways. You have this linear phased array. And there is a one-point-five-D array—so this is a one-dimensional array, but you have several one-dimensional rows; the number is not too big. If the number is very big—so the x- and y-directions are kind of symmetric—you have a truly 2-D array.

Phased means each ultrasound pixel—each element—does not fire in the same way. Each pixel you can think of as a source, as a vibrator. You can make a single piezoelectric material vibrate in the way you want. You can control the vibration with electrical signals through the so-called piezoelectric effect—the reaction plane. And once we have a two-dimensional array or a 1.5-D or 1-D array, each pixel—each transducer—can be individually controlled.

With that individual control, we can freeze—computer. Let me see how to deal with it. Let me close the program. I will actually play a little bit with this. These are the reasons—I think a computer problem. This is the tool.

slide9:

OK, think about a linear array. A linear array is to put multiple piezoelectric transducers along a line. So you have multiple elements, and you can selectively make a subset of transducers on. This is the light—this subset is selected in this case. These elements are selected. You use these three as your ultrasound sources, so this beam will go down the road this way.

You do scanning as if this were just a single transducer aperture. Then you can define the next line—you just group elements differently. Then you have a neighboring line, line 2, sending ultrasound waves downward. OK. Next instant, you generate line 3. So you have a line of piezoelectric elements; the elements remain, but each time, the first second—you use this part; the second second, you use the second group. You keep doing this.

The physical effect is that you send one beam here; next, you send another beam. You keep doing A-mode scanning along many lines. But you do the beam scanning to form a B-mode image not by manually changing location or angle—you do it mechanically or electronically. This is very convenient. You have a reform, so you use the acoustic transducer array, and you can perform scanning easily. And not only scanning, you can also do beam forming—a lot of flexibility.

So, here, I would like to explain to you.

slide10:

So, the very famous Huygens' principle. This slide is to guide the idea. Once you have this idea clear, the different schemes for ultrasound scanning and beam forming with an array-type ultrasound transducer will become very clear to you. The so-called Huygens' principle basically says: if you have a source—just a simple example—if you have a point source, it will generate a spherical wave. The wave will keep moving out, and at a certain time, you have a wavefront here.

Think of the ultrasound wave: you have compressions and you have rarefactions. That boundary is the wavefront. The principle suggests that if you decompose the current wavefront into small elements, you get a small element here, a small element there—multiple small elements. Because the wavefield keeps vibrating—like an ultrasound wave—the small volume element keeps getting smaller and bigger, oscillating around a nominal volume; the pressure keeps changing, following a simple sinusoidal curve.

By this reasoning, if you view things individually, you isolate a small voxel—you see the voxel keep changing, vibrating, oscillating. You can think: OK, this voxel is a single, small, secondary ultrasound source—or for electromagnetic waves, a small secondary source. Treat that element as a source and ask: how do we predict or calculate the next wavefront?

One way is to solve the original wave equation and follow the whole propagation process. The other way—the Huygens way—is to forget the original source and say: at a given time, the current wavefront—shown here, A to A-prime—can be decomposed into many small segments. Each individual segment is regarded as a secondary small source. Each small source keeps changing, and each small source generates its own small spherical wave. The next wavefront is the superposition of all those secondary spherical waves added together, forming the new wavefront.

So this is Huygens' principle. The current wavefront is decomposed into numerous small sources, and the next wavefront is computed from all these secondary sources. In words, we say: any current point on a wavefront is a new source—a secondary source. And all such new secondary sources, collectively, determine the future wavefront. All the subsequently generated waves added together will form the next wavefront. That is the idea.

slide11:

Looking at this picture, you think: this is a current wavefront. You have multiple small secondary sources here. Each source forms a small secondary spherical wave. These waves together—the envelope formed by all these secondary spherical waves added together—form the next wavefront. In this first case, the next wavefront has the same shape as the previous or current wavefront, so a spherical wavefront propagates as a spherical wavefront. That looks right.

If you have a plane wave propagating forward, the plane wavefront can also be decomposed into sources—A, B, C, D, and so on. Each is a secondary spherical wave. Again, the envelope of all these small spherical waves forms a plane, which is a new wavefront parallel to your current wavefront, parallel to the previous wavefront. That illustrates how a plane wave propagates along one direction—the direction of beam propagation.

Now, if you have a small aperture, the plane wave goes down to the opening. Because you have multiple yellow sources across that aperture, and each yellow source is treated as an individual point source vibrating out sound or light, you observe diffraction. The light or sound will no longer go perfectly straight; it will bend around the edges—turn around, so to speak.

If you have two slits, the waves from the two openings will interfere due to phase differences. At some points the two spherical waves—one from here and one from the other opening—will add constructively; at other points they will cancel. If you place a screen there, you will observe bright and dark strips. This is interference.

The reflection, the refraction, and many old optical or ultrasonic phenomena can be explained similarly. You have a plane wave here; the wave is incident at a surface. At the surface, you again treat it with Huygens' sources. Depending on the location, this point is hit earlier and its secondary spherical wave is generated in this direction; that point is hit later and its spherical wave has a smaller extent. The envelope will be defined by straight lines—like plane fronts—depending on the speed of sound or light in the two media, the relative speed in the two media. You will see the transmitted beam change direction. The direction change is determined by the relative speeds of the wave in the first and second medium. This explains a lot of wave phenomena—very cool.

slide12:

Now, let's look at this example—an ultrasonic transducer array shown here inside the blue box. If you have a beam that is focused toward a certain direction, it shines in a specific mode, and you will have a wavefront forming within that blue region. So, what happens to the wave? Essentially, it behaves as if that entire front of energy is moving outward through the medium.

But here, instead of a continuous wavefront, we use a two-dimensional acoustic transducer to recreate that same wave behavior. In other words, we can replace the natural wavefront with this 2D transducer array. By doing so, we can electronically control each element in the array to operate exactly the way the original wavefront would. This means that the two-dimensional transducer can emulate the same effect as the real wave field.

With this approach, we can do the same kinds of operations—focusing, scanning, and steering the beam—with great flexibility. As long as each individual transducer element is driven properly, the array can reproduce the correct wave behavior. This is how we visualize and control wave propagation in a precise, programmable way.

slide13:

Let me make this clearer using a linear phased array example. If you fire all the pulses at the same time, all the elements are activated together, and the ultrasound beam travels straight ahead in one direction.

However, if you adjust the timing slightly, you can shape the beam. For instance, if you fire elements 1 and 5 a little earlier, as shown on the time axis here, where this point is time zero, and then fire the others with slight delays, you create different wavefronts that add up. At first, two spherical waves are generated from elements 1 and 5. Then, a bit later, elements 2 and 4 emit their waves, and finally, the central element fires last. When all these small wavelets combine, they form a curved wavefront, producing a focused ultrasound field downward, with the focal spot located here.

Now, the distance of this focus—how deep it forms from the transducer surface—depends on the relative time delays between the elements. If you apply the delay more gradually, the focus will form more deeply. By changing the delay pattern, you can also tilt the beam—to the left-hand side or the right-hand side—by linearly shifting the delay across elements. This gives you complete flexibility in how you steer and focus the beam.

The essential idea here follows Huygens' principle—each element acts as an individual source, generating its own waveform with a controlled relative phase. By adjusting the phase and timing, you can create different acoustic or optical effects. This is how beam steering and dynamic focusing are achieved in a phased array.

slide14:

Now, this is a phased array, and you can extend the same concept to the two-dimensional case. Instead of having one line of elements, you now have a grid—a 2D array of transducer elements.

By controlling the phase and timing in two dimensions, you can purposely focus along the principal axis, or you can steer the beam left, right, up, or down—even in a circular pattern if you wish. This allows for electronic focusing and beam steering in both dimensions, giving you precise control over the acoustic field.

So, in summary, the 2D phased array operates the same way as the 1D version but adds another level of control. It enables full 3D imaging capabilities, higher flexibility, and more advanced ultrasound applications.

slide15:

Now, let's look at the annular array. You can make the ultrasound beam radially symmetric, or what we sometimes call recirculately symmetric. Basically, you can design a circular symmetric arrangement of piezoelectric elements, like the one shown here. With this type of array, you cannot do circular scanning—because it's already circularly symmetric—but you can focus along the principal direction.

The advantage is that you can control the focal distance very easily. You can make the focus point closer or farther, depending on the time delay and how you drive those circular rings. So you can adjust the focal distance to be small or large. That's something you can easily do with this annular array structure.

slide16:

So, the beam-forming idea is that you introduce time delays across the array. Imagine you have scatterers located at different positions, labeled E1, E2, and E3. These are the echoes or signals you receive from different points inside the tissue.

On the receiver side, you receive all these signals and keep guiding or aligning them properly. If you add those signals together—but with individual delays as shown here—you can make them sum coherently. The tall bars in the diagram indicate a larger time delay. Different modulations of these delays determine which echo signal you want to detect.

At this location, the phase difference between the central element and the peripheral elements will be large, so the relative delay will be large too. You take the signal from the central channel and keep it. Then, you take the signal from the outer element—it takes a longer time for the echo to reach that peripheral element. That longer time becomes the time delay.

You add all the signals together, each with its own individualized time delay. When the time delays are properly distributed like this, and you add the five signals together, you form a strong echo—for example, for echo one. But if you use a uniform time delay for all, you are focusing to infinity—that's a remote point. So, by changing the distribution of time delays, you can guide coherent echoes at different depths.

This is the basic beam-forming idea. And of course, scanning and beam forming can be combined to achieve the best image quality and overall performance.

slide17:

Now, let me just copy this figure to help explain the idea. This shows a one-dimensional ultrasound array. If you use just this group of transducers, you can certainly do beam forming. You perform focusing in this way—based on the relative time delays across elements.

You can adjust the focus at a certain depth—from the transducer surface to the focal plane—and make the focus spot here. You receive signals mainly from this region of interest. Then, this group of signals can be moved laterally.

From left to right, you perform scanning—line by line—so you get information along each line. When you repeat this line-by-line process, you effectively build up a two-dimensional image.

So, in summary, with beam forming, you control how the echoes are added together in time, and with scanning, you move the beam position across the field. Together, they form the foundation of modern ultrasound imaging.

slide18:

Now, let's connect the ideas we discussed earlier. You can sweep the ultrasound beam from left to right, and when you perform ultrasound imaging, you can use this sweeping approach to create what we call compound imaging. Using a phased array transducer, you can steer the ultrasound beams toward the left side, collect multiple image lines—each one similar to an A-mode line—and combine them to form a B-mode image over the region of interest. Then, you slightly change the beam angle and sweep across again, this time pointing the beams toward the right side.

What you get are multiple images of the same region, each captured from a different angle. When you combine or "add" these images together, you average the results across those angles. The benefit is that the underlying structures—like the organs or tissues—stay consistent across frames, but the speckle noise and random scattering patterns vary from one angle to another. When you sum them together, the random

noise cancels out while the consistent features are reinforced, giving you a much smoother, clearer image with reduced artifacts.

So, this is what we call compound imaging. It's essentially a noise reduction technique achieved by imaging the same region from multiple directions. The figure here shows how the image looks smoother, with far fewer speckles compared to a single acquisition. This method is similar in principle to noise averaging or multi-angle integration.

In short, this is another powerful way to perform ultrasound imaging—whether A-mode, B-mode, or C-mode—based on Huygens' principle and the use of one- or two-dimensional transducer arrays. You can scan, focus, and form images in many flexible ways. So, read the green textbook carefully and make sure you understand these imaging principles. That completes the first part of this lecture.

slide19:

The second part of this lecture focuses on dynamic imaging, particularly on how we measure motion, such as blood flow velocity or heart wall movement. To do this, we rely on the Doppler effect.

So, what exactly is the Doppler effect? It's a physical principle that describes how the observed frequency of a wave changes when there is relative motion between the source and the observer. In ultrasound, this effect allows us to calculate how fast and in what direction blood is moving inside the body.

slide20:

Here's a simple cartoon that helps you visualize the concept. You've probably experienced the Doppler effect in everyday life. Imagine you're standing at a train station. When the train comes toward you, the pitch of its whistle sounds higher. When the train moves away, the pitch sounds lower.

The same thing happens with the siren of a police car. When the car is moving toward you, the sound waves are compressed, giving you a higher frequency or a sharper tone. When the car moves away, the sound waves are stretched, and you perceive a lower frequency.

In other words, the perceived frequency depends directly on the relative velocity between the sound source and the observer. This change in frequency with motion is the Doppler effect, and it forms the physical foundation for Doppler ultrasound imaging, where we measure blood flow speed and direction inside the body.

slide21:

Now, in your textbook, you have this diagram that shows how the Doppler shift occurs in flowing blood. Here, we are looking at a small vessel with red blood cells moving inside. The transducer sends ultrasound energy toward the blood vessel wall or directly into the blood stream. The wall or the moving red blood cells reflect the ultrasound energy back to the transducer.

If the wall or blood cells were perfectly stationary, the reflected frequency would be exactly the same as the incident frequency. However, since these scatterers are moving, the reflected frequency becomes slightly

different. The vertical component of the velocity—represented here as v cos θ—is what contributes to this difference. This component corresponds to the velocity of motion along the ultrasound beam direction.

The small difference between the transmitted and received frequencies is what we call the Doppler frequency shift. This effect is universal—it appears not only in ultrasound but also in optical and electromagnetic waves. The Doppler shift tells us how motion changes the observed frequency, and that is the basis for Doppler imaging in medical ultrasound.

slide22:

Here is the key paragraph from your textbook. The diagram is reduced in size, but the principle remains the same. You have motion, an incident frequency, and a reflected frequency. The derivation goes as follows: the frequency reflected by the moving blood cell or boundary is given by Equation 3.40. Here, c is the speed of sound in the tissue, v is the blood flow velocity, θ is the angle between the beam and the direction of motion, and the whole term is divided by λ, the wavelength.

Now, how do we derive this? The idea is straightforward. The wavelength times the frequency gives you the speed of the sound wave, meaning $c = f \times \lambda$. The term v cos θ represents the velocity component of motion in the direction of the beam. Together, these describe the relative velocity between the transducer and the moving target.

The incident frequency of the ultrasound is denoted f_i, which equals c divided by λ. The difference between the incident frequency and the reflected frequency is shown in the last equation. This frequency difference, called the Doppler frequency shift, allows you to calculate the velocity component along the beam direction. So, from the measured shift, you can infer how fast the blood or tissue is moving toward or away from the transducer.

This may not be immediately clear from the formulas, so let's use a simple cartoon to understand why this Doppler shift happens.

slide23:

So, you think—look at this. The transducer sends out ultrasound waves at the speed of c toward a stationary object. In this case, the speed c equals frequency times wavelength. The wavelength is basically how far the wave moves during one vibration. And how many times do you vibrate? That's the frequency. So, frequency times wavelength equals the speed c.

Now, in one unit of time, in this stationary case, you see four pulses. I should make that plural—four pulses are sent into the object. This is the reference situation. This simple example gives you the baseline idea, and I just drew more examples to make it clearer.

Now, think about the moving situation. I say, when the object is moving, you will see the apparent frequency change. How does that happen? Think about this: let's say this same stationary object now starts a steady motion toward the transducer. This is just for simplicity. In fact, the object's motion is also at a certain velocity, denoted v, and the sound itself travels with velocity c.

So, in one unit of time, the sound will travel a distance c. In that same unit time, you have four pulses sent into the object. But now, in the second case, the object is moving—also with velocity v—so within that same

unit time, the moving object will cover a certain distance. Because it covers that distance, if it were stationary, it would have received four pulses. But now, since it moves toward the transducer, those four pulses that were still traveling through space are “caught” or “eaten up” by the moving object.

So, this object will intercept not just the original four pulses, but also an additional four pulses that were on their way. Altogether, you now have eight pulses entering the object in one unit of time. Let’s assume the object is a perfect reflector.

In the first case—when the object is stationary—you have four pulses entering, and it will reflect four pulses within that unit time. So, you receive four reflected pulses, and that gives you the same frequency as the transmitted one—no frequency change. The object vibrates four times, and you see four reflected waves per unit time.

But in the second case—when the object is moving toward the transducer—it takes eight pulses within the same time period. If it is a perfect reflector, those eight pulses are now reflected back within one unit of time. So, the reflected signal contains eight pulses per unit time, meaning the reflected frequency is doubled.

This is how a moving object changes the reflected frequency. You get a signal reflected back with a higher frequency, and the difference between the stationary and moving cases is the Doppler effect.

So, this is the idea. Think about this simple diagram, and then review the equations again. If you think about the connection between this cartoon and those equations, you will gain a much better understanding of the formula—especially the one where we talk about the velocity component v times cosine θ.

Only the portion of the velocity component along the beam direction contributes to the observed frequency shift. That’s the part of motion you can measure as the carrying frequency. The incident frequency is the transmitted one, and the frequency shift—the Doppler shift—is this difference.

This normalization factor is simply the speed of sound, c. So, think about it this way: half of the transmitted power goes out, then it comes back after reflection. That’s the basic Doppler imaging process.

Doppler imaging is an important part of the ultrasound chapter. To really understand Doppler imaging, the key is to understand this formula. If you follow the derivation, it might not be totally clear at first, but if you understand this cartoon analogy, you can see how the Doppler frequency shift is generated.

The same idea applies across all situations. So, follow these formulas systematically. Then you’ll see what’s going on—why there’s a connection between the Doppler frequency shift and the velocity component that you’re measuring along the ultrasound beam direction.

If the motion is parallel to the beam, the Doppler shift is at its maximum. If the motion is perpendicular to the beam, you cannot measure it. You can only measure the Doppler frequency shift along the beam direction.

This is a simple idea, but it gives you a very good understanding of what’s happening physically.

slide24:

Now, let’s talk about how Doppler imaging is actually performed. You can use a continuous wave (CW) method, where you keep sending ultrasound waves continuously, or a pulsed wave (PW) method, where you send one pulse at a time—a train of pulses—to measure the Doppler information. Another important

technique is echo correlation, where you send one signal, receive it back, then send another signal, and again receive it back. By comparing the relative phase changes through correlation, you can detect motion.

Correlation tells you how much one signal matches another. In the foundation part, we mentioned the Cauchy–Schwarz inequality, which states that if two waveforms are perfectly correlated, they represent the same event at a specific location. When the reflector moves, the correlation pattern shifts, which reveals how much the object has displaced between the two measurements.

You can also combine Doppler imaging with B-mode imaging. The B-mode provides an anatomical background—essentially, a grayscale structural image. If you have multiple two-dimensional B-mode slices, you can reconstruct a three-dimensional ultrasound volume. On top of this anatomical structure, you can superimpose the velocity distribution obtained from Doppler measurements.

Velocity is a vector quantity, meaning it has both amplitude and direction. You can visualize these vectors using color coding. For example, motion in one direction may appear in red, motion in the opposite direction may appear in green, and another perpendicular direction may be shown in blue. The color intensity represents the velocity amplitude. With this color flow encoding, you can visualize the complete velocity field superimposed on the anatomical image. This combined visualization—structural plus velocity—is the essence of Doppler imaging. Now let's go step by step to see how each method works in practice.

slide25:

These concepts are quite straightforward, and it's good to know them as part of your general understanding. In continuous wave Doppler, the transducer is divided into two separate parts—one part for transmission and the other for reception. The transmitting part continuously sends ultrasound waves toward the region of interest, while the receiving part continuously collects the reflected signals.

So, you have a single transducer housing two elements. One is always sending, and the other is always listening. The signal that is sent out has a fixed transmission frequency, while the reflected signal may have a slightly different frequency because of the Doppler shift caused by motion. What you send and what you receive differ slightly in frequency, and that difference directly relates to the motion along the ultrasound beam direction.

The region of interest is where the transmitted and received beams overlap. This overlapping region is the sensitive zone from which the Doppler signals are detected. If the reflection occurs outside that overlap, the signal will not be received effectively.

In summary, continuous wave Doppler requires two separate elements—one for emission and one for detection—because a single element cannot transmit and receive continuously at the same time. So, CW Doppler uses continuous waves for both sending and receiving, providing constant velocity measurement but without specific depth localization.

slide26:

Now let's take a closer look at what actually happens to the signal. You have one signal going in and another signal coming out. The input signal is transmitted at a certain frequency, which we call f-i, while the received signal includes a small Doppler frequency shift, noted as delta-f. So, the frequency of the returning signal

becomes f-i plus delta-f. This delta-f can be either positive or negative, depending on whether the target is moving toward or away from the transducer.

Using digital signal processing, we can compare these two signals. The received signal is multiplied by the reference signal generated by the transmitter, and this multiplication is performed by a circuit called a mixer, which effectively combines the two signals. Mathematically, this creates two components—one at a very high frequency, which is the sum of the two, and one at a much lower frequency, which is their difference.

Since the transmitted frequency f-i is in the high ultrasound range and the Doppler shift delta-f is quite small, we use a low-pass filter to remove the high-frequency component and keep only the low-frequency term that carries the Doppler information. After filtering, the output signal is proportional to A over two times cosine of two-pi times delta-f times t, where A is the amplitude of the signal. This low-frequency signal represents the Doppler shift itself and allows us to estimate the velocity v using the formula you saw earlier: delta-f equals two times f-i times v times cosine theta divided by c.

So, by measuring delta-f, we can determine how fast the object is moving along the ultrasound beam direction. However, remember that continuous-wave Doppler does not provide depth information. It tells you the velocity but not the exact location of the moving target, because the transducer is transmitting and receiving continuously along the entire beam path. So, continuous-wave Doppler gives accurate velocity data, but without depth resolution.

slide27:

The second mode we use is called pulsed wave Doppler. In this mode, we do not send ultrasound continuously. Instead, we transmit a series of short pulses — almost like delta functions — separated in time. It's not exactly a delta function, but it's very narrow in duration.

So, at time one, you send out the first pulse into the tissue along the time axis. Then, after receiving the echo back, at time two, you send out the second pulse, and so on. Each pulse travels into the tissue, reflects from structures, and returns to the transducer. The process repeats again and again, one pulse at a time, to measure motion and flow at specific depths within the body.

slide28:

Now, I've included here two images from another textbook that show this more graphically. You can see how the pulses are sent one after another into the biological tissue. Between each transmitted pulse, there is a fixed time delay. That time interval is the pulse repetition period, and the inverse of that is called the pulse repetition frequency. This repetition rate is an important parameter because it determines how deep your ultrasound can penetrate and how fast you can sample motion.

Each pulse travels into the tissue and produces echoes that return at different times depending on the depth of the reflecting structure. These returning echoes — let's call them echo one, echo two, echo three — come from slightly different layers or scattering surfaces inside the tissue. If there is no motion, the interval between successive echoes remains the same. But if there is motion — for instance, blood flowing or tissue moving — that spacing changes slightly. By measuring the change in the timing or phase of these echoes, we can estimate velocity.

In pulsed wave Doppler, we can also target a specific depth, often referred to as the gate depth. This means we can focus the measurement on a particular region of interest within the body. The operator can adjust this gate — make it deeper, shallower, wider, or narrower — to control the resolution and sensitivity of the measurement. By analyzing how the returned echoes shift over time, we can extract information about both motion and flow at that specific location.

slide29:

Let's imagine we send a pulse to a stationary target — say, a student sitting still. The pulse travels to that target and comes back after a certain time. Suppose it takes exactly one hundred milliseconds for the sound to make the round trip. If we send a second pulse after one second, that second echo will return at the same delay — one hundred milliseconds — if the target hasn't moved. Both echoes will look identical, because the distance hasn't changed.

But now, if that student moves slightly toward us, the second echo will arrive a bit earlier — the return time becomes shorter. The peak of the echo signal shifts forward in time. This small timing change, or phase shift, tells us that motion has occurred. If the target moves rhythmically — for example, the surface is oscillating — the measured signal amplitude will fluctuate up and down around the average line.

We record this oscillating signal over time and then perform a Fourier transform on it. This mathematical operation converts the time-domain signal into a frequency-domain representation. The resulting frequency shift corresponds to the Doppler frequency, which reflects the motion or velocity of the surface or fluid. Usually, the motion is not perfectly uniform but has some periodic pattern. By analyzing these Doppler frequencies, we can quantify that motion precisely.

So, this is the basic idea behind pulsed wave Doppler. It's slightly more complex than continuous wave Doppler, but it has a major advantage — it allows us to measure motion at a specific depth. We'll look at some examples and waveforms shortly to illustrate how this works in practice.

slide30:

Now, let's talk about the pulsed wave Doppler mode. In this method, we don't use a continuous wave. Instead, we send out short, separated pulses—each one acting somewhat like a delta function. It's not exactly a delta function, but it behaves like one in time.

So, at time one, you send out a pulse into the tissue, and after a short delay, you receive the echo back. Then, at time two, you send the next pulse, and again you wait for the return signal. You keep sending these pulses one after another, in sequence. Each pulse travels into the tissue, interacts with the structures, and returns as a backscattered echo. The transducer then captures this echo signal, allowing you to measure reflections from specific depths—what we call the range gate.

slide31:

And I pause the wave, and I have two more ways to show you. So, let me just show you again. I hope the computer will not mess up. But I think good for you to see two more ways here. Just try. Okay. Okay. Try to pause the state, which means the reflector or surface does not have a

Let’s go a little deeper into how the pulsed wave behaves. Imagine you’re sending a sequence of pulses toward a target that is not moving. You send one pulse, two pulses, three pulses—each time the echo returns with the same amplitude and phase. If you line up all these received echoes at the same gated time instant, they all align perfectly. The phase remains constant, meaning the surface is stationary.

Now, let’s imagine a different case—one where the reflector is moving or oscillating back and forth. In that case, when you line up the received echoes, you’ll notice that their amplitudes and phases change from pulse to pulse. At a fixed gated time instant, the amplitude might be high for the first echo, lower for the second, and different again for the third. When you plot those amplitude values sequentially, they form a sinusoidal curve, representing a cyclic or oscillating motion.

If the target moves in a steady direction—say, a uniform motion instead of an oscillation—the resulting signal will not be sinusoidal but rather a linear change over time. So, each echo provides a single sampled data point from a given location, and by analyzing how those points evolve over a train of pulses, you can determine the relative motion of the scatterer or reflecting surface.

In your textbook, you’ll find several formulas that describe this mathematically, but the key idea is quite straightforward. The pulse wave mode works by collecting one data point from each returning echo at a fixed range. If there is no motion, the points stay constant; if there is motion, the points vary in amplitude or phase. By examining those variations, we can determine relative motion, which can then be related back to Doppler frequency and ultimately to velocity. However, if the object moves too fast, the returning echoes can overlap or become aliased, causing artifacts in the signal. This is one of the limits of the pulsed wave Doppler technique.

slide32:

Now, let’s look at another approach—the correlation method, which is quite intuitive. You start with a transducer that sends out a pulse of ultrasound and records the entire backscattered signal as a function of time. That gives you a complete waveform, not just a single data point. Then, you send a second pulse, and again you record the returned signal. By the time the second pulse arrives, the scatterers—such as red blood cells—will have moved slightly compared to their original positions.

To find out how much motion occurred, we compare the two signals using correlation. You look for the time delay, denoted as tau, where the two signals are most similar. This time delay tells you how far the scatterers have moved between the two measurements. When you multiply that tau by the sound speed along the beam direction, you get the relative displacement. From that, you can calculate the velocity or even derive the Doppler frequency if you prefer.

So, this is the basic idea behind the correlation method—it’s simple but very effective. By comparing the returning waveforms over time, we can accurately measure how fast and in what direction the blood or tissue is moving. This is often used alongside Doppler analysis to enhance motion estimation accuracy.

slide33:

Now, this slide shows how the Doppler effect can be used to measure and visualize velocity. The Doppler information can be superimposed directly on the B-mode image. The B-mode gives you the anatomical

structure — for example, you can clearly see a large blood vessel here. Once you have that anatomical background, you can add the velocity field on top of it using color coding.

In color Doppler imaging, the motion is encoded using colors, typically the RGB scheme. Motion toward the transducer is displayed in red, while motion away from the transducer is shown in blue. Sometimes a green component is used to indicate the variance of the signal — that is, how much fluctuation or turbulence there is in the flow. If the blood flow is smooth and uniform, you won't see much green at all. But if the flow becomes turbulent or irregular, you'll see more green appearing in the image.

This color map gives a one-dimensional view of motion — only along the direction of the ultrasound beam. It doesn't tell you about lateral movement, but it's very effective for examining flow in vessels. In more advanced cases, you can apply a Fourier transform to the Doppler signal to obtain the velocity distribution, because the motion is not always at a single frequency. Blood flow, for instance, has multiple velocity components due to pulsation and vessel geometry. So, color Doppler provides a dynamic, intuitive way to visualize those variations on top of anatomical images.

slide34:

Here we move to spectral Doppler imaging, as shown in Figure 3.23 from your textbook. This method provides more detailed information about the range of velocities within a vessel. You can see how the signal amplitude varies over the cardiac cycle — during systole, when the heart contracts, and diastole, when it relaxes. During systole, the flow velocity is higher, which shifts the Doppler frequency upward. During diastole, the flow slows down, so the Doppler shift becomes smaller.

The result is a spectrum showing how signal amplitude changes with Doppler frequency over time. Each vertical line in the display represents the distribution of frequencies — or velocities — at a given instant. When plotted together, this forms the spectral waveform, which you often see in clinical Doppler displays. By analyzing these patterns, clinicians can assess blood flow speed, direction, and uniformity through the cardiac cycle.

So, to summarize, Doppler imaging includes several complementary modes: continuous wave, pulsed wave, and correlation methods, along with color and spectral Doppler. Continuous and pulsed modes measure velocity; color Doppler visualizes it; and spectral Doppler reveals its frequency content. Together, they provide a comprehensive view of motion and flow in ultrasound imaging — anatomically and dynamically.

slide35:

Finally, let's talk briefly about ultrasound image artifacts. Like any imaging modality, ultrasound is not perfect — and various artifacts can appear due to limitations in the system or the physics of sound propagation. These artifacts may come from reflection, refraction, attenuation, or motion. Some are simply distortions, while others can mimic real anatomical structures.

Understanding and recognizing these artifacts is extremely important, because they can affect how you interpret the image. We'll look at examples of these artifacts and discuss what causes them, and how they can sometimes even be used advantageously in image interpretation.

slide36:

Now, let's talk about time-gain compensation. When we deal with very high-frequency ultrasound—ultrasound signals in the megahertz range—the signal amplitude can be extremely small. If we simply record the raw signal amplitude or intensity as it is, without doing any correction, you'll find that the deeper features inside the tissue appear much weaker than the ones near the surface.

This happens because of attenuation and wave scattering. As the ultrasound wave travels deeper, part of the energy is absorbed, and part of it is scattered in different directions. Both effects cause the returning signal—the echo—to decay exponentially with depth. Near the surface, the reflected wave is strong, but as you move deeper, the signal strength drops off very quickly. The decay roughly follows an exponential curve.

To correct for this loss, we apply what's called time-gain compensation, or TGC. The idea is simple: as the echo arrives later in time—which corresponds to greater depth—you gradually increase the amplification factor. In other words, you boost the gain as a function of depth to counteract the exponential decay. The amplification curve roughly follows the opposite shape of the attenuation curve. When you multiply the received signal by this depth-dependent gain, the resulting image becomes more uniform in brightness.

This compensation makes the ultrasound image look much better and more interpretable. Without it, the deeper regions would appear dark and faint, while the shallow regions would look overly bright. With TGC, you recover a more balanced image. However, remember that this correction is based on assumptions about tissue homogeneity—it works well when the tissue properties are roughly uniform. In real biological tissue, conditions are not perfectly homogeneous, so the quantitative accuracy is not perfect.

It's similar to what happens in CT imaging when we apply beam-hardening correction—we can reduce the artifacts, but it's never a perfect fix. So, keep in mind that in ultrasound, different echo times correspond to different depths, and the deeper the echo, the stronger the amplification you need to apply. That's the principle of time-gain compensation.

slide37:

Now, let's talk about the next type of artifact, which is caused by what we call side lobes. In ultrasound imaging, no matter how you modulate the transducer elements — whether you use a one-dimensional or a two-dimensional array, or even a single transducer — you can never create a perfectly straight, narrow beam. Each transducer element generates a spherical wave, and when these waves combine through interference, they form one strong main beam along the principal direction, plus several weaker beams on the sides — these are the side lobes.

Physically, this happens because we're dealing with waves, and waves naturally spread out. Even if we adjust the phase and amplitude across the array, we can't perfectly confine all the energy into a single pencil-thin line. The best we can achieve is a main lobe with high intensity, surrounded by smaller side lobes with lower intensity. If you've studied Fourier analysis, this is similar to what happens when you approximate a sharp function — you always get some rippling, known as ringing. Here, the same principle applies to the ultrasound beam.

Now, these side lobes can create misleading signals in the image. For example, if you're scanning a region with a single bright reflector, the main lobe gives a strong echo at the correct location. But the side lobes can also pick up reflections from off-axis points, which then appear as false echoes in the image — often as

smaller or dimmer copies of the real structure. So, you might see multiple bright spots, even though only one real reflector exists. In clinical practice, recognizing this artifact is essential — you may move the transducer slightly to confirm which echo is genuine and which is an artifact.

Another related effect is reverberation, which happens when echoes bounce back and forth between two reflective surfaces — such as tissue layers or an air-tissue interface. These multiple reflections cause several delayed echoes to appear on the image, even though only one actual surface exists. They appear as repeated, evenly spaced lines, decreasing in intensity with depth. You can see examples of these side lobe patterns and reverberation artifacts in the slide images here. Understanding them helps you identify what's real and what's simply an effect of wave interference and multiple reflection.

slide38:

Let's continue with more examples of ultrasound artifacts. One common type is caused by multiple reflections between layers. Imagine sending one ultrasound pulse into a structure like the lung. The wave first reflects from the top surface, giving you the first echo. Some energy penetrates deeper, hits the bottom surface, and reflects. But on its way, that echo may bounce again from the top surface and return once more to the transducer. As a result, even though you sent only one pulse, you detect several echoes arriving at different times. This produces a sequence of repeated bands on the image — very similar to hearing multiple "hellos" echoing in a mountain valley after shouting once.

Another important artifact is acoustic shadowing. This occurs when a dense object — for instance, a gallstone or kidney stone — strongly reflects or absorbs the ultrasound energy. Because the wave cannot penetrate beyond it, there's no signal coming from the tissue behind the object. On the image, this appears as a dark region, or "shadow," extending below the bright reflective surface. You can clearly see this in the figure on the right, where the shadow forms behind a gallstone.

In the images on this slide, you can see both types of artifacts — reverberation lines appearing as repeated echoes, and acoustic shadowing creating dark regions behind dense materials. Recognizing these patterns is essential for accurate diagnosis. They're not always errors; sometimes they actually help identify the nature of the structure — for example, confirming the presence of a stone or a metallic implant. But conceptually, you should always understand how these artifacts form and what they tell you about the ultrasound-tissue interaction.

slide39:

Now, in the last part of this lecture, let's take a look at a few interesting new ideas in ultrasound imaging. We won't go deep into equations or technical details here — I just want to give you a brief overview of some of the most exciting developments happening right now in this field.

slide40:

We have a faculty member, Professor Pingkun Yan, who has done excellent work in collaboration with Philips, a leading medical imaging company. Their research focuses on combining ultrasound and MRI imaging for prostate cancer surgery. The idea is to use MRI before the operation and ultrasound during the operation, so you can merge or "fuse" the information from both modalities.

MRI imaging provides excellent spatial resolution and sensitivity — it can clearly show anatomical and biological details, which makes it ideal for locating the tumor before surgery. However, MRI scanners are very expensive and not cost-effective for real-time surgical use. They also have slow temporal resolution, meaning it takes time to acquire the image. On the other hand, ultrasound imaging is fast, inexpensive, and easy to use in real time, but its spatial resolution and sensitivity are not as high.

By combining the two, surgeons can perform preoperative MRI imaging to identify the problem, and then, during the operation, use ultrasound as real-time guidance. Through image registration, or fusion, the live ultrasound image can be aligned with the preoperative MRI. This allows the surgeon to visualize the tumor location with both high spatial detail and real-time feedback. The fusion system developed by Dr. Yan's group has produced very promising results in clinical trials.

slide41:

Now, if you look at the images here, you'll see what happens if the image registration between ultrasound and MRI is not done properly — the structures won't line up, and the surgeon can't accurately target the tumor. Without good alignment, it's impossible to rely on the fused images for surgery.

To solve this, Dr. Yan's team developed machine learning algorithms to automatically register ultrasound and MRI images. Using these intelligent registration techniques, the alignment quality improves dramatically — for example, in their results, the similarity index increases from about 0.24 to nearly 0.99. That's a huge improvement.

With this level of precision, surgeons gain much higher confidence during the procedure, leading to better surgical outcomes. So this hybrid imaging approach — combining MRI's high spatial resolution with ultrasound's real-time imaging capability — represents a powerful direction for image-guided surgery. It's a very active research area with many exciting developments and a strong clinical impact.

Next, we'll look at another hybrid direction — combining ultrasound imaging with optical imaging.

slide42:

In the next lecture, we'll move into optical imaging, but before that, let's look at how ultrasound and optical imaging can be combined to create a powerful new imaging technique known as photoacoustic tomography. The idea is shown here. Imagine you inject an animal with contrast agents — in this case, graphene-based microbubbles. Once these microbubbles are introduced, they enhance the ultrasound signal because the presence of microbubbles makes acoustic reflections stronger.

Now, graphene has an interesting property: it can absorb light energy and convert it into heat. So, when you shine near-infrared light pulses onto the tissue, the graphene microbubbles absorb the light and heat up slightly. This heating causes a tiny thermal expansion — the microbubbles expand and contract in response to the pulsed laser. Each laser pulse causes a brief expansion followed by cooling, and this mechanical vibration generates an acoustic wave, much like a sound pulse. These acoustic waves are then detected by an ultrasound transducer.

So, you can think of the process as a chain: near-infrared excitation leads to graphene absorption, which causes heating and thermal expansion, generating acoustic waves that are picked up by an ultrasonic detector. From these signals, you can construct an image for diagnosis. The advantage is that the contrast is

introduced optically, while the detection is done acoustically, giving both sensitivity and depth penetration. In addition, some of these microbubbles can be used for therapeutic applications — for example, targeted drug delivery. There's a lot of exciting research happening in this area, exploring both diagnostic and therapeutic uses.

slide43:

Here you see some examples of photoacoustic tomography from Professor Lihong Wang, who made pioneering contributions to this field. This imaging method can be applied at multiple scales — from single cells and small biological structures to whole animals and even human patients. The remarkable thing is that it combines the strengths of both optical and ultrasound imaging.

Purely optical imaging provides excellent contrast but limited penetration depth — the light scatters too easily, so you can only see shallow structures. Ultrasound, on the other hand, penetrates much deeper and gives good spatial resolution, but the contrast is not as strong because it depends mainly on acoustic properties. When you combine the two in photoacoustic imaging, you get the best of both worlds: strong optical contrast and deeper acoustic penetration.

This hybrid approach allows researchers to visualize structures ranging from microvasculature to organ-level features, making it a powerful tool for both biological research and clinical applications. It's a perfect example of how multi-modality imaging can overcome the limitations of individual techniques.

slide44:

Finally, here's something still speculative – an iPhone-based ultrasound imager. With this device, we could turn an iPhone into a miniature ultrasound device. Imagine using your phone not just as a communication tool but as a health monitoring instrument. For example, we might use the phone's screen or back surface as a transducer array to measure things like blood velocity, blood pressure, or even blood sugar.

Let's consider blood pressure first. Traditionally, blood pressure is measured with a cuff that compresses the vessel, but perhaps we could use Doppler ultrasound principles instead. When blood flows through a vessel, the pressure from the heartbeat changes both the velocity and the cross-sectional area of the vessel. When the pressure is high, the vessel expands; when it's low, the vessel contracts. The velocity of blood flow also changes accordingly and can be measured. The shape of vessels can be measured as well.

If we could measure both the velocity (through the Doppler shift) and the vessel wall motion (by detecting the relative distance between the top and bottom surfaces), then, with proper calibration and machine learning, we might estimate blood pressure accurately — just by placing the iPhone over the skin. Over time, the system could learn your personal profile and give real-time blood pressure readings without the discomfort of a cuff. The same principle could even be extended to use an optical sensor for measuring blood sugar related signals, by analyzing light interactions with tissue. These ideas are still pre-mature, but with advances in materials, transducers and sensors, they may become realistic in the future.

slide45:

Alright, so this is the last slide for today. I'm not going to give you any regular homework, but it's a good time for you to start reviewing or re-reviewing the ultrasound chapter — what I usually call the green chapter. For my classroom teaching, I want to take a moment to share a few thoughts about how traditional classroom learning compares with what we're trying to do here.

Traditional teaching is often about knowledge transfer. In other words, you're given a bunch of information — formulas, numbers, and rules — and you're expected to become familiar with what other people have already done. You learn the steps, you can repeat them, you can reproduce the same results, and you can memorize the equations. There's nothing wrong with that — it's still a good foundation. But our goal goes a bit beyond that.

What I really want is to inspire you to have new ideas. To reach that point, you have to truly understand what's going on beneath the surface. For example, when we studied the Doppler system earlier, I didn't just want you to memorize the formula. I wanted you to see the reasoning behind it —to connect the physics to what you actually observe. That's the kind of understanding that helps you innovate. We encourage you not to simply copy what others are doing. Don't just see a bridge someone has built and make the same bridge again — instead, learn how to design your own.

You need to create, to innovate, to generate something new. That's far more important.

Now, when I talk about reviewing, we still use a closed-book or closed-form style of study. You can use the book, the slides, and the solutions. Even if you don't understand every detail right away, I want you to read through the green chapter. Learn the concepts, understand the examples, and become familiar with the existing answers.

At the same time, try to think hard where you can improve or add new functionality, such as iPhone-based possibilities.

You need to think beyond memorization — think about problem identification. When you review the material, try to notice how pioneers and earlier researchers identified their problems. How did they get their ideas? How did they use the tools available to them? Even if the textbook doesn't directly give you an answer, you can always look for more — use Google Scholar, use Web of Science, explore the literature, and ask ChatGPT and other AI models. Learn to find what you need and to solve problems independently.

And finally, motivation matters. The traditional way of teaching doesn't always emphasize new ideas or encourage creativity. But in real research, motivation drives discovery. That's what I want to underline. So, when you review and study, don't just focus on memorizing the content — look for the underlying meaning behind it. Try to understand the ideas, the motivation, and the reasoning that led to them.

So anyway — that's all for today. Take some time to reflect, review the green chapter, and start thinking creatively about how you can connect what you've learned to new ideas. That's how innovation really begins.