slide1:
Welcome back, everyone.
Today, we move one step further in our journey – from the Fourier series to the Fourier transform. This is a big transition, and it’s a powerful one. So, what exactly changes when we move from a series to a transform? That’s what we’ll uncover together in this lecture.
Before we begin, I want to emphasize something important. You need to have a solid understanding of the Fourier series first. The transform builds directly on those ideas. If you’re still feeling confused about Fourier series – don’t worry – but do take action. You can interact with ChatGPT, discuss with your classmates, review the lecture materials, and review the chapter I wrote. You can also search online for different explanations or examples that make more sense to you. Group discussions can be constructive.
Fourier analysis is fundamental – not just for this course, but for understanding medical imaging technologies down the line. I want to make sure no one gets left behind here. Because if you miss this foundation, the rest of the material will feel considerably harder.
Alright – let’s dive in.

slide2:
Again, so this is our schedule. We are on schedule, so no problem.

slide3:
Now, let’s break things down to a very basic and powerful idea.
Suppose we have a function – let’s call it f of t – a one-dimensional function that varies over time. You can think of it as a smooth, continuous curve. That’s the usual way we look at functions.
But here’s a different perspective – instead of seeing f of t as one continuous piece, imagine it as a sum of impulses. Yes – a collection of sharp, narrow spikes, each carrying a little bit of information about the function at a specific time.
This is where the Dirac delta function, or delta for short, comes in. From linear systems theory, we know that any continuous function can be viewed as the convolution of that function with a train of delta functions.
What does this mean, intuitively? It means you can imagine slicing the original function into many tiny segments. Each segment becomes a small impulse – and when you add up all those impulses, you reconstruct the full function.
So, this is one way to represent a function – not as a smooth line, but as a weighted sum of sharp impulses. And this idea is going to be very helpful as we move toward the Fourier transform.

slide4:
Now let’s look at the same function from a completely different angle – instead of seeing it as a sum of impulses, imagine it as a sum of waves.
Here, we’re working with a periodic function – meaning it repeats itself over time. Let’s just focus on one complete cycle of that function. The rest are just copies.
The key idea is this: a periodic function like this can be broken down into many sine and cosine waves – we call these sinusoidal components. These waves can vary in three ways: their frequency, which tells us how fast they oscillate; their amplitude, which tells us how tall they are; and their phase, which tells us where each wave starts.
So what’s the trick? We want to find the right combination of amplitudes, frequencies, and phases – so that when we add up all these sine and cosine waves, we recover the original function.
At first, this might sound abstract, but conceptually it’s not that difficult. You’re just looking at the same function in a different way.
Earlier, we saw a kind of particle view – slicing the function into sharp impulses. Now we’re seeing a wave view – smoothing it into oscillating signals. And this wave-based view is the foundation for Fourier analysis.

slide5:
So now that we've talked about breaking a function into a sum of waves, let's formalize that idea using the real form of the Fourier series.

What we’re looking at here is just a mathematical way of saying: any periodic function f of t can be written as a combination of three types of components:
A constant term, called the DC component,
A sum of cosine terms,
And a sum of sine terms.
Each term has a different frequency and amplitude, depending on the value of n.
The higher the value of n, the higher the frequency of that term.
So the general formula goes like this:
f of t equals a-zero over 2, plus the sum from n equals 1 to N of a-n times cosine of 2 pi n t, plus b-n times sine of 2 pi n t.
Now, where do these coefficients come from?
Well, we have formulas to compute them. That’s the trick!
To find a zero, you integrate f of t from 0 to 1.
To find a-n, you multiply f of t with cosine of 2 pi n t, and integrate.
And for b-n, you do the same, but with sine of 2 pi n t.
So if I give you any function f of t, you just plug it into these formulas, do the math, and you get a set of numbers — the a’s and b’s. With those coefficients, you can then reconstruct the original function by summing all the sine and cosine waves.
As N gets larger — meaning you include more wave components — the reconstructed function becomes more and more accurate. In fact, when N approaches infinity, this series can represent the function exactly, under reasonable conditions.
Now here’s something deeper: These sine and cosine functions form an orthonormal basis — that’s just a fancy way of saying they’re like the X, Y, and Z axes in 3D space, but in an infinite-dimensional space of functions.
So geometrically, what you're doing is projecting the function f of t onto each basis function — kind of like breaking a 3D vector into its X, Y, and Z components. Except here, we’re breaking a function into sine, cosine, and constant components.
If this all feels a bit abstract, don’t worry. The key takeaway is:#We’re representing a function as a sum of waves.#And we have formulas that let us find the right weights — or coefficients — for those waves.
That’s what the Fourier series, in its real form, is all about.

slide6:
So far, we've talked about the real form of the Fourier series, which uses sine and cosine functions. But there's another way to write the same idea — in what we call the complex form.
Now don’t be alarmed by the word “complex.” This version is mathematically equivalent to the real form — it's just more compact and elegant.
Instead of separating things into sines and cosines, we use complex exponentials — specifically, e to the power of 2 pi i n t. That’s Euler’s formula at work.
In this form, the function f of t is expressed as a sum — from n equals minus N to N — of coefficients c n multiplied by e to the 2 pi i n t.
And just like before, we need to figure out those coefficients. So how do we compute c n?
We use this formula:#c n equals the integral from 0 to 1 of e to the power minus 2 pi i n t times f of t, with respect to t.
This is essentially an inner product — projecting your function onto the exponential basis function e to the power minus 2 pi i n t.
In my book, I like to write the imaginary unit i in front — as in minus i 2 pi n t — but it’s just a matter of notation. The meaning stays the same.
At the bottom here, you see the same formula written slightly differently — as f-hat of n — which is just another way of naming the coefficient.
Now, throughout this slide, we’re assuming that the function is periodic with unit period — meaning it repeats every interval from 0 to 1. If the function instead repeats from, say, 100 to 101, it’s still the same picture — just shifted.
So the complex form doesn't change the logic — it simply gives us a cleaner, more powerful way to write and manipulate Fourier series, especially when we move into Fourier transforms.

slide7:
So far, we’ve been working under a simple assumption — that the function we’re

analyzing has a unit period. That means the function repeats every one unit of time, say, from 0 to 1.
But what if the period is something else — like 5, or 100? That's what we're covering here.
This slide shows how to handle a function with an arbitrary period, which we'll call capital T. The good news is: the concept stays exactly the same — we're still breaking the function into complex exponential components — but we just adjust the formulas a bit.
Here's how the function f of t is now written:
f of t equals the sum from n equals minus infinity to infinity of c n times e to the power i 2 pi n t over T.
So you can think of this as a generalized version of the complex Fourier series.
And to compute the coefficient c n, we use the formula:
c n equals 1 over capital T, times the integral from minus T over 2 to T over 2 of e to the power minus i 2 pi n t over T, which is multiplied by f of t, with respect to t.
Now let me break that down for you:
The exponential term — that's your harmonic basis function, adjusted for the new period.
t over T is just a way to normalize time, so that we're still operating over a standard interval.
The 1 over T in front is a scaling factor that makes everything work out correctly.
And the range of integration — from minus T over 2 to plus T over 2 — gives us a symmetric interval, which is often more convenient for analysis.
So, whether your period is 1, 10, or any other positive number, you simply use T in place of 1 and apply this formula. You're still summing wave-like components — just stretched or compressed in time depending on the period.
Now here's something helpful:#If you set T equals 1, this formula reduces back to the unit-period version we saw before. So you only need to remember this general form — and adjust T as needed.
In practice, the real form of the Fourier series is nice because it gives you a clear geometric picture — sines, cosines, and their amplitudes.#The complex form, like the one we see here, is more compact and elegant, especially when we move into transforms. It just takes a little abstract thinking in complex space.
So this wraps up our review of Fourier series — both real and complex — and now we're ready to move on to the Fourier transform itself.

slide8:
By now, you might be wondering — why go through all this? Why even bother with the Fourier transform?
Well, let's take a step back and think with common sense.
We started with a function — f of t — and we've looked at several ways to represent it. First, we saw it as a sum of impulses. Then, as a sum of waves — sines and cosines. Each of these views gives us a different lens to understand the same signal.
And that's the point. Multiple perspectives give us more flexibility and smarter strategies for solving problems.
In real life, when you're facing a tough task, what do you usually do? You don't try to tackle it head-on in the hardest way possible. You try to simplify it. You look for shortcuts. You use tools that make the problem easier.
That's exactly what we're doing with Fourier analysis.
Sometimes a problem looks really messy in the time domain — but if we switch to the Fourier domain, that same problem might become simple and elegant.
Here's an analogy: imagine trying to do multiplication using Roman numerals. It's a nightmare! But switch to Arabic numerals — or even binary — and suddenly, multiplication becomes easy. Especially for computers.
So choosing the right representation can make all the difference.
That's what the Fourier transform is all about. It gives us a new way to look at a function — in terms of its frequency content — and often, that view makes analysis or computation much simpler.
There's also a deeper strategy here:#Sometimes a function is complicated as a whole, but if we break it down into smaller, simpler parts, we can understand it better.

This is the classic divide and conquer approach.
We’ve already seen this with impulse decomposition. If you know how a system responds to a single impulse, then you can figure out how it will respond to a complicated signal — just by adding up the responses to each impulse.
The same logic applies to sine waves — if you understand how a system reacts to one sine wave, and you can express a signal as a sum of sine waves, then you’ve got the whole picture.
So, in short, Fourier analysis follows two powerful principles:
Use simple tools to solve complex problems.
Divide big problems into small, manageable pieces.
And that’s what we’re building toward with the Fourier transform.

slide9:
So now that we’ve covered the Fourier series — and we’ve seen why changing representations can make problems easier — it’s time to dive into the main topic of today’s lecture: the Fourier transform.
This corresponds to Chapter 4 in the book draft I shared with you. I do plan to revise it soon — polish up some explanations and fix a few typos — and I’ll send you the updated version once it’s ready.
Here’s how we’ll structure today’s lecture:
First, in Section 1, we’ll talk about how the Fourier transform is derived from the Fourier series. This transition is logical and elegant, and it gives us a solid foundation for everything that follows.
Then, in Section 2, we’ll explore some of the most important properties of the Fourier transform. These properties aren’t just mathematical curiosities — they’re practical tools used in signal processing, imaging, and many other fields.
After that, in Section 3, we’ll extend our thinking to higher dimensions. I’ll show you how the Fourier transform applies not just to one-dimensional signals, but also to images — and beyond. I’ll even walk you through an example where converting an image to the Fourier domain makes it easy to remove noise.
Finally, in Section 4, we’ll wrap things up with a few remarks and reflections.
So by the end of this lecture, you should not only understand the mathematical structure of the Fourier transform, but also see its real-world value, especially in areas like medical imaging.
Let’s get started with the first section: the derivation of the Fourier transform.

slide10:
You’ve actually seen this slide before — it’s a summary of what we covered in the last lecture, and now it serves as our starting point for deriving the Fourier transform.
Let’s quickly revisit the main idea:
We said that any arbitrary function can be treated as if it were periodic, defined over a fixed interval. In this case, we’re looking at a symmetric interval — from minus capital T over 2 to plus capital T over 2.
Within this interval, we can represent the function f of t as a sum of sinusoidal components — specifically, complex exponentials of the form:
e to the power i times 2 pi n, times t over T.
This is a compact way to write sine and cosine waves using Euler’s formula. So even though the expression looks complex, it represents a mix of sine and cosine waves at various frequencies, where n runs from minus infinity to plus infinity.
And just like before, each component has a corresponding coefficient, which tells us how much of that wave appears in the overall function.
To compute the coefficient c n, we use the formula at the bottom:
c n equals 1 over capital T, times the integral from minus T over 2 to T over 2, of e to the minus i 2 pi n t over T, times f of t, with respect to t.
This is just an inner product — projecting f of t onto a complex exponential basis function. That projection gives us the weight, or the contribution, of that particular frequency.
So, this slide wraps up all the key ideas from the Fourier series — including:
Treating functions as periodic,
Expressing them with complex exponentials,
And computing coefficients through integration.

With this solid foundation in place, we’re now ready to take the next step — and derive the Fourier transform, which generalizes everything we’ve done so far.

slide11:
Alright — now let’s take the next step together.
Earlier, we learned how to calculate the Fourier coefficient c n for a function with arbitrary period T. Now, we’re going to insert that expression for c n directly into the Fourier series formula.
What happens when we do that?
Well, here’s what we get: we take the formula for c n — which is an integral — and plug it into the sum. So now, each term in the series contains an inner product multiplied by a complex exponential.
Notice this part out front: 1 over capital T. This shows up because we’re working over a general period T, not the unit interval. So we need to normalize — or average — over the full interval from minus T over 2 to plus T over 2.
This 1 over T acts like a scaling factor — a way of balancing the sum so that everything works out correctly.
Now look inside the brackets. What we’re really doing here is taking the inner product of two functions:
One is f of t,
And the other is the complex exponential e to the power minus i 2 pi n t over T.
You multiply them together point by point, and then integrate over the interval.
That’s what we mean by an inner product in this context.
And we’re doing this for many values of n, both negative and positive. So you can imagine we’re computing inner products at a whole series of frequency points — where each frequency is n divided by T.
As T becomes large, the spacing between these frequencies — that is, the difference between one n over T and the next — becomes smaller and smaller.
Eventually, this sum will start to resemble an integral over a continuous range of frequencies, and that’s the key idea behind the Fourier transform.
So, to summarize:#We’re inserting the formula for c n, and that gives us a series of inner products across many frequencies, with a spacing of delta u equals 1 over T. This step sets us up perfectly to move from a discrete sum to a continuous integral — which is what the Fourier transform is all about.

slide12:
Now to understand that last idea more clearly, let’s take a look at this picture.
What we’re doing here is sampling inner products at many frequency points. Each of these points corresponds to a value of n divided by capital T — that is, u equals n over T. As n varies from negative to positive integers, you get a collection of frequency points spread along the u-axis, which represents frequency.
When n equals 0, you’re at the DC component — the zero-frequency term. When n equals 1, you get one frequency point. When n equals minus 1, you get another. And so on.
So you can imagine each vertical line here as one of those frequency points — u equals n over T — spread across the axis.
Now here’s the key idea:#If we let the period T become very large, then the range from minus T over 2 to plus T over 2 will grow wider and wider. In fact, in the limit as T goes to infinity, that interval covers the entire frequency axis — from minus infinity to plus infinity.
That’s what we mean by dense sampling of frequency.
As T increases, the spacing between neighboring frequencies gets smaller and smaller. Mathematically, the spacing — which we call delta u — becomes 1 over capital T. So the gap between u equals to n over T and u equals to (n plus 1) over T shrinks.
Eventually, these discrete frequency samples get so close together that they begin to form a continuous spectrum.
And that’s the magic.
We started with a periodic function, which gives us discrete frequencies. But as the period stretches toward infinity, the function becomes non-periodic, and the discrete frequency points turn into a continuous frequency range.
So, what is a non-periodic function in this view?

It’s just a periodic function with an infinitely long period. Anything that happens beyond that infinite window doesn’t affect what we see — and that’s how we bridge from the Fourier series to the Fourier transform.
This slide captures the big picture:
Discrete frequencies spaced by 1 over T,
Becoming densely packed as T increases,
And ultimately forming a continuous frequency axis.
That’s our gateway into the Fourier transform.

slide13:
Alright — this slide is the heart of the lecture. Take a moment to follow this carefully, because here is where we formally derive the Fourier transform — and also see how we can recover the original function using the inverse transform.
Let’s begin at the top.
We’re starting with our expression for f of t, written as a Fourier series. We substitute the expression for c over n — the Fourier coefficient — right into this formula.
Now remember, c over n contains an integral — and it includes a 1 over T term out front. Since T is constant, we pull it out of the sum.
At this point, we apply a key idea we introduced earlier:#Let T become very large, so large that we’re effectively dealing with a non-periodic function.
When T approaches infinity, the spacing between the frequency samples, delta u equals to 1 over T, becomes tiny — almost zero.
So what happens?
Instead of summing over discrete frequencies, we move to a continuous frequency variable, which we call u — where u equals n divided by T.
This transition lets us replace the sum with an integral.
So now we’ve gone from a sum of discrete terms to an integral over a continuous range of frequencies. The inner product becomes f-hat of u, the Fourier transform of f of t. And the exponential term becomes e to the power i 2 pi u t — our new kernel in the transform.
Let’s pause here and look at the final expression.
You’ll notice two things:
First, to compute f-hat of u, the Fourier transform, we integrate f of t multiplied by e to the minus i 2 pi u t. That’s the forward transform — moving from the time domain to the frequency domain.
Second, to reconstruct f of t, we take f-hat of u, multiply it by e to the power i 2 pi u t, and integrate over all u. That’s the inverse transform — bringing us back from frequency to time.
So these two formulas — the forward and inverse transforms — are the foundation of the Fourier transform.
They let us take a non-periodic function and decompose it into a continuous spectrum of frequencies. And then, using that spectrum, we can reconstruct the original function with complete accuracy.
This is what makes the Fourier transform so powerful — not just mathematically, but also practically — in fields like signal processing, image analysis, and, of course, medical imaging.
From this point forward, we’re now working with non-periodic functions, and we’ve fully transitioned from Fourier series to the Fourier transform.

slide14:
Now let’s bring things down to earth a bit and look at a concrete example — something visual, something simple.
This function right here is called the rectangular function, or sometimes the gate function. You’ll also see it written as pi of t.
And as you can see from the graph — it really does look like a gate.
Let’s break it down:
The value of the function is 1 when the absolute value of t is less than one-half — that is, between minus one-half and plus one-half.
Outside that interval, the function drops to zero.
So, the entire “on” region is one unit wide, and the height is 1, which means the area under the curve is also 1.
This function is simple, but it’s also very important — and we’re going to use it often in later examples.

Now here’s the key point:#This is a non-periodic function.
And back when we only had Fourier series — which apply to periodic functions — we couldn’t directly express this kind of shape using sines and cosines.
But now, thanks to the Fourier transform, we can.
We’re no longer limited to repeating signals. We can now handle functions like this — compact, finite, non-repeating — and still express them in terms of sinusoidal components, just over a continuous range of frequencies.
So this rectangular function — although it’s simple — gives us a great opportunity to visualize the power of the Fourier transform.
Let’s keep going and see what it looks like in the frequency domain.

slide15:
Now let’s do something clever.
We just looked at the gate function, or rectangular function, which is not periodic. And because it’s non-periodic, we can't represent it using a Fourier series — only the Fourier transform works in that case.
But what if we wanted to use Fourier series?#Well — we can! We just need to make the function periodic.
And here’s how we do it: we take that single gate function — and simply repeat it over and over again, at regular intervals. This is called periodization.
In this slide, the original gate function is being repeated with a period of 16.
So now you’ve got a train of rectangular pulses, equally spaced along the time axis.
Visually, you can think of it like flipping a light switch on for a second — and then leaving it off for a long time — then flipping it on again — and repeating that process over and over.
The key is: now we’ve created a periodic signal.
And because it’s periodic, we can now apply Fourier series to analyze it — just like we did before.
All you need to do is plug this periodic function into the Fourier series formulas, and you’ll get a representation as a sum of sine and cosine waves — or complex exponentials.
This kind of construction is very useful. In fact, it shows up in engineering and physics quite a bit — especially in systems that switch on and off repeatedly. Sometimes you’ll hear people talk about the duty cycle, which refers to how much of each period the function is “on” versus “off.”
So again — by making the function periodic, we unlock the power of the Fourier series.
And soon, we’ll compare this with what happens when we go back to the non-periodic version, and apply the Fourier transform instead.

slide16:
Now, let’s visualize the powerful idea we’ve been building toward — how we move from Fourier series to the Fourier transform.
What you see here is the frequency-domain representation of the rectangular function as we increase the period of its periodic extension.
Let’s start at the top left.
These vertical lines represent the Fourier coefficients — discrete spikes at different frequency points. These are the outputs of the Fourier series.
You’ve got:
A DC component at zero frequency,
The first harmonic,
The second harmonic, and so on.
Each coefficient is calculated using the formulas we saw earlier. And because the original function is real-valued, these coefficients appear symmetrically around zero when using the complex form.
Now — here’s the trick:#We begin to increase the period T of the function.
That means we’re spacing the rectangular pulses farther apart in the time domain. And in the frequency domain, this has a very important effect:#The spacing between the Fourier components, which is 1 over T, becomes smaller and smaller.
You can see this happening in the middle image — the spikes begin to cluster more tightly. The frequency axis, labeled u equals n divided by T, gets more densely sampled.

And as we continue increasing T — moving to the bottom image — the spacing becomes infinitesimally small. At this point, we're no longer looking at a discrete spectrum. We now have a continuous curve.
That's the moment when we've transitioned from the Fourier series to the Fourier transform.
So here's the big picture:
When a function is periodic, we get a discrete Fourier spectrum.
But when the function becomes non-periodic, by letting the period go to infinity, we end up with a continuous Fourier spectrum.
This spectrum is smooth and continuous — and it tells us how much of each frequency exists in the original signal.
So this slide gives us the geometrical insight — the visual transition — from the world of periodic signals and discrete spectra, to the world of non-periodic signals and continuous spectra.
And this is what the Fourier transform captures so beautifully.

slide17:
Now that we've built up all the ideas step by step, let's bring it all together into one elegant concept — the Fourier transform pair.
What you see on this slide captures the very heart of Fourier analysis.
Let's begin with the first equation, which is the forward Fourier transform.#Suppose you have a signal in the time domain, which we'll call f of t. The goal is to understand how much of each frequency is present in that signal.
So how do we do that?
We project f of t onto a family of complex exponential functions. These functions look like e to the power minus i, 2 pi, u, t, where u is the frequency. This projection acts like an inner product, telling us how strongly f of t resonates with each frequency component.
When we integrate this product across all time, we obtain f hat of u, also called the frequency spectrum or the Fourier transform of f of t.
So this first equation maps the signal from the time domain to the frequency domain.
Now look at the second equation. This is the inverse Fourier transform.
Here, we take all those frequency components — that is, each f hat of u — and multiply them by their corresponding complex exponential, this time e to the power i, two pi, u, t. Then, we integrate over all frequencies.
And what do we get?#We reconstruct the original signal — f of t — exactly.
So in summary, the forward transform takes you from time to frequency, and the inverse transform brings you back from frequency to time.#That's the two-way mapping shown symbolically at the bottom:#f of t goes to f hat of u, and back again.
Now, a few mathematical notes to make this complete:
First — this transform assumes that f of t is square-integrable, meaning if you square the function and integrate it over all time, the total is finite. This ensures convergence and makes the math work properly.
Second — even though we're dealing with a non-periodic function here, the idea still grows out of the Fourier series. Remember what happens when the period of a function becomes very large — the discrete set of frequency components becomes a continuous spectrum. The Fourier transform is essentially the limit of the Fourier series as the period tends to infinity.
And finally — here's a nice geometric interpretation.
Think of projecting a 3D vector onto the x, y, and z axes.
In the same way, when we apply the Fourier transform, we're projecting our function onto an infinite set of sine and cosine waves — or more precisely, complex exponentials. Each frequency contributes a tiny slice of the original signal.#And by adding all those slices back together — using integration — we reconstruct the full signal.
This isn't just a clever trick.#It's one of the most powerful tools in all of engineering, physics, and mathematics.#It lets us analyze and manipulate signals in both time and frequency — without losing any information.
With this foundational concept in place, we're ready to explore some practical examples and dive deeper into the properties of the Fourier transform.

slide18:
Let’s go through a concrete example – the Fourier transform of the gate function, also called the rectangular function.
We’ve already seen this function earlier. It’s a flat function that equals 1 when the time variable t is between minus one-half and plus one-half, and it’s zero everywhere else. So it looks like a rectangle centered at zero.
Now, we want to find its Fourier transform. To do that, we apply the definition. We take the integral of the original function multiplied by a complex exponential – that’s e to the power negative i 2 pi u t – and we integrate over all time.
But since the gate function is zero outside the interval from minus one-half to plus one-half, we only need to integrate over that range.
So the integral becomes: from minus one-half to plus one-half of e to the negative i 2 pi u t, with respect to t.
This is a standard exponential integral. When you solve it, you get the function sine of pi u, divided by pi u.
This result is known as the sinc function and it shows up everywhere in signal processing and imaging.
So here’s what we’ve found: a rectangle in time turns into a sinc wave in frequency. That’s a beautiful and very useful result.
Now, what if we stretch the gate to make it wider?
Let’s say the gate equals 1 from minus capital T over 2 to plus T over 2, and zero elsewhere. Then the Fourier transform becomes a scaled version of the same sinc function – it just shrinks or stretches based on the value of T.
So you could try computing it directly, or you could use a shortcut – the scaling property of the Fourier transform, which we’ll talk about soon.
But the key point is this: the rectangular function in time gives us a sinc function in frequency.
And here’s one final note. You might wonder what happens when u equals zero, because both the top and bottom become zero. But don’t worry – if you take the limit using calculus, it turns out the value is exactly 1. So the function is smooth at the center.
We’ll use this gate-to-sinc pair often, so keep it in mind as we move forward.

slide19:
Now let’s take a look at the actual shape of the sinc function – the one we just derived from the rectangular gate.
What you see here is the graph of that function – sinc of u – which is defined as sine of pi times u, divided by pi times u.
It has a smooth, wave-like shape. The peak is right at the center – at u equals zero – and the height there is exactly 1.
As you move away from the center in either direction, the function oscillates – it goes up and down – but the peaks get smaller and smaller.
That’s because the sine wave in the numerator keeps swinging, but the denominator grows, so the overall value shrinks.
This is the signature look of the sinc function – one strong central lobe and then smaller and smaller ripples on the sides.
So this is the frequency domain representation of a simple gate in time.
And just to recap – this sinc function came from a rectangular gate function of width 1.
If we use a wider gate – let’s say the width is capital T instead of 1 – then the sinc function would stretch out horizontally. That is, the ripples would become narrower in frequency, because time and frequency are inversely related.
We’ll explore that in more detail when we talk about the scaling property. But for now, just remember: a narrow gate in time gives a wide sinc in frequency. A wide gate in time gives a narrow sinc.
This is a key idea we’ll return to again and again.

slide20:
Now let’s move on to our second example – the triangle function.
This function, often called the triangular function, is shaped just like its name suggests – a triangle. It peaks at 1 when x equals 0, and then it decreases linearly to zero as x moves out to minus 1 and plus 1. Outside of that range, the value is zero.

Mathematically, it’s defined like this: one minus the absolute value of x when x is between negative 1 and 1, and zero otherwise.
So what we want to do now is find its Fourier transform.
In other words, we want to express this smooth, non-periodic triangular function as a combination of infinitely many sinusoidal waves — just like we did with the gate function.
We’ll add up all these waves — each with a different frequency and amplitude — and by doing so, we’ll be able to reconstruct this triangle shape exactly.
That’s the idea behind Fourier transform: take a function, no matter how it looks, and rewrite it as a sum of wave components. Each component carries part of the shape — and when you combine them all, you get the original back.
Let’s see what that looks like in this case.

slide21:
Let’s now take a look at the result of the Fourier transform of the triangle function.
If we go through the steps — plugging the triangle function into the Fourier transform formula and carrying out the integration — we end up with a very elegant result. The Fourier transform turns out to be the square of the sinc function.
That is, we get sinc squared.
You might remember that the Fourier transform of the gate function gave us a sinc function. So this result — sinc squared — is not just a coincidence. There’s actually a deeper reason behind it.
It’s related to an operation called convolution. We saw this in the context of the Fourier series, and we’ll explore it again in more depth when we talk about properties of the Fourier transform.
But the key takeaway here is this: the triangle function is, in a way, the convolution of two gate functions. And in the frequency domain, convolution in time corresponds to multiplication in frequency. That’s why the sinc function — from the gate — becomes sinc squared for the triangle.
And down at the bottom, we can see the graph of sinc squared. It has the same overall shape as sinc — a central peak and decaying ripples — but now the oscillations are all positive and fall off more smoothly.
So, this example shows another beautiful connection between shapes in the time domain and their patterns in the frequency domain.

slide22:
Let’s now take a look at a few more examples of Fourier transform pairs — some of them elegant, some of them a bit more abstract.
First, we have the Gaussian function. In the time domain, it’s written as e to the power negative pi t squared. This is the familiar bell-shaped curve — smooth, centered, and decaying quickly. What’s remarkable is that its Fourier transform is also a Gaussian — e to the power negative pi u squared. The shape stays the same, just transformed into the frequency domain. The Gaussian is one of the few functions that’s unchanged, except for scaling — we say it’s self-Fourier.
Next, let’s consider a constant function — just a flat value c across time. Now, technically this function is not square-integrable, meaning we can’t just apply the Fourier transform in the usual way. But in a generalized sense, we can still assign it a meaning. And it turns out the Fourier transform of a constant is a scaled delta function — specifically, c times delta of u. This makes intuitive sense: a constant signal has no frequency variation, so all its energy is concentrated at zero frequency.
And finally, look at this interesting example: the shifted delta function delta of t minus a. If you perform the Fourier transform of this, you get a complex exponential — e to the power minus i 2 pi a u. This tells us that shifting a delta function in time introduces a phase shift in frequency. And that phase shift depends on the amount of translation — the value a — and also on the frequency u.
Now, I should point out — when you’re working with generalized functions like constants and delta functions, the usual rules don’t always apply directly. These are not square-integrable functions. So we use the tools of distribution theory to handle them more rigorously. If you’re curious, I explain this in more

detail in the book chapter.
But for now, just keep in mind — even with these edge cases, the idea of expressing a signal in terms of wave components still holds.

slide23:
Now that we’ve gone through several examples, let’s take a moment to highlight some core properties of the Fourier transform — the features that make it such a powerful and elegant analytical tool.
We begin with linearity. If you combine two functions, say f and g, using constants a and b — meaning you form "a times f plus b times g" — then the Fourier transform of that sum is simply "a times the transform of f plus b times the transform of g." This reflects the principle of superposition, which is central to all linear systems.
Next is the time-shifting property. If you shift a function in time — let’s say you move it by x naught — then in the frequency domain, the transform stays the same in shape but gets multiplied by a complex exponential factor. Specifically, it’s multiplied by e to the power minus i 2 pi x naught u. So shifting in time introduces a phase shift in frequency.
Now let’s flip that idea — this is called modulation. If you multiply a function in the time domain by a complex exponential, then you shift its frequency content. So time-domain modulation causes a translation in the frequency domain.
Another important concept is scaling. If you compress or stretch a function in time — say, you use f of a t — then the frequency representation stretches or compresses in the opposite way. And the result is scaled by one over the absolute value of a. So if the time signal gets narrower, the spectrum spreads out.
We also have conjugation. If you take the complex conjugate of a function, then its Fourier transform reflects across the frequency axis — that is, the transform at minus u becomes the conjugate of the transform at u.
And finally, if the original function is real-valued, its Fourier transform has Hermitian symmetry. That means the spectrum is symmetric in a complex-conjugate sense.
Each of these properties gives us deeper insight into how signals behave across time and frequency. We’ll explore them in more detail as we continue — but for now, keep these tools in mind. They’ll help you decode and design systems with much greater clarity.

slide24:
Let's look into more detail. Okay, the first linearity. So this is just, again, this part copied from the Stanford textbook. So you have a function f and g. There are Fourier transformations. So you have Fourier transformation. You have this summation, the summation of the original function.
Then you perform Fourier transformation. The result is the same as the summation of the Fourier transformation of f and the Fourier transformation of g. So this is additivity. So in the system, linear system lecture, we explain that.
How about the scaling or homogeneity? So if f is scaled by alpha, then we perform Fourier transform. That's the same as performing the Fourier transform of f. Then you scale the result by the same scaling factor. So you can verify these two properties according to the formula by defining the Fourier transform. So you can just see here, just as an example, you show additivity. And you can similarly show the scaling property.

slide25:
Now let’s look at another important idea in Fourier analysis — the shift property, also called the translation property.
This property explains what happens when we shift a function in time. Suppose we have a function, f of t, and we shift it by some constant b. That means our new function is f of t minus b. So what happens to its Fourier transform?
The answer is: the shape of the spectrum stays the same, but we multiply it by a phase factor. Specifically, we multiply it by e to the power minus 2 pi i s b. In other words, shifting the function in time causes a rotation in the phase of the frequency spectrum.
Let’s go through why that’s true.
We start with the definition of the Fourier transform of f of t minus b. Inside

the integral, we make a substitution: we let u equal t minus b. That’s just a change of variable — and it doesn’t affect the limits of the integral, which still go from negative infinity to positive infinity.
Now, when we rewrite the integral in terms of u, something interesting happens. The exponential term splits into two parts. One part depends on u, and the other depends on b. The part with b is just a constant, so we can pull it outside the integral.
What remains inside is the same integral we started with — the Fourier transform of the original function, f of t.
So in the end, we have the original spectrum, multiplied by this exponential phase term.
That’s the core idea of the shift property:#When you shift a function in time, it introduces a phase shift in the frequency domain.
This is an extremely useful result when working with signals that are delayed or moved in time — it tells you exactly how the spectrum changes, and helps preserve the full mathematical relationship between time and frequency.

slide26:
Let’s take a closer look at the scaling property of the Fourier transform.
Here’s the idea. Suppose we have a function, f of t. And now, instead of f of t, we look at f of a times t, where a is just some constant — not zero.
We’re changing the time scale. The question is: what happens to its Fourier transform?
Well, this depends on whether a is positive or negative. But in both cases, something interesting happens.
First, let’s say a is greater than zero.
We plug f of a t into the Fourier transform formula. If you do the math — which involves changing variables and rearranging terms — you end up with this:
The Fourier transform of f of a t becomes#one divided by a, times F of s divided by a.
So what does that mean?
If we stretch the function in time — that is, make it slower — its frequency content gets compressed.#And if we squeeze it in time — make it faster — the frequency content stretches out.
Now, what if a is less than zero?
In that case, the same formula still works, but with a minus sign that comes from flipping the limits of integration.
So we just take the absolute value of a, and the final result becomes:
The Fourier transform of f of a t equals#one over the absolute value of a, times F of s over a.
So in short:#Scaling in time causes the opposite effect in frequency.#That’s the scaling property.
And this is one more way to see how time and frequency are tightly connected in the Fourier world.

slide27:
So now, let’s bring this scaling idea to life with a visual example.
Take a look at the top plot. This shows different versions of a triangular function — blue, orange, and green. They all have the same shape, but different widths.
The blue triangle is the narrowest.#The orange one is a little wider.#And the green one is the widest of all.
Now, let’s look at what happens in the frequency domain — shown in the bottom plot.
Here’s the key idea:#When the function gets wider in time, its frequency content gets narrower.
You can see that clearly:
The narrow blue triangle has the widest Fourier transform — that’s the blue curve in the bottom plot.
The medium-width orange triangle gives a more focused frequency response — that’s the orange curve.
And the wide green triangle gives a sharp, tall peak in the frequency domain — the green curve below.
This is the heart of the scaling property.#Wider in time means tighter in

frequency, and vice versa.
It’s a kind of duality – as if time and frequency are pulling on opposite ends of a rope.
And as a special case, if you use a delta function – which is infinitely narrow in time – its Fourier transform becomes perfectly flat, spread across all frequencies.
On the other hand, a Gaussian function is balanced – it’s the only shape where both the time and frequency profiles remain Gaussian.
So visually, scaling is not just a math trick – it’s a powerful way to understand how functions behave in both domains.

slide28:
Now let’s talk about a beautiful and surprising property of the Fourier transform – its connection to derivatives.
Suppose we start with a function, f of t, and its Fourier transform is capital F of s.
Now you might ask:#What happens if we take the derivative of f – that is, f prime of t – and then perform the Fourier transform?
The answer is elegant:#Taking a derivative in time becomes multiplication in the frequency domain.
More precisely,#The Fourier transform of f ' of t is just 2 pi i s times the Fourier transform of f.#So again, taking a derivative becomes a simple multiplication – and the factor depends on frequency.
There’s also a second version of this rule:#If you take the derivative of the Fourier transform itself with respect to s,#That’s the same as applying -2 pi i t to f of t – and then transforming it.
So in summary:
Derivatives in one domain correspond to multiplications in the other domain.
And this duality works both ways – from time to frequency, or frequency to time.
This relationship is incredibly useful in practice, especially in physics and engineering, where differentiation often pops up in systems and signals.
And best of all, it’s not magic – you can prove this directly from the Fourier transform formula, just by carefully taking the derivative under the integral.
So this derivative property really adds to our collection of powerful tools in the Fourier toolbox.

slide29:
Now here's one of the most important Fourier transform pairs – the impulse train, also called the comb function.
Think of it this way:
We place delta functions at evenly spaced intervals – for example, every delta t along the time axis.#This gives us a repeating sequence of sharp spikes – like the teeth of a comb – so we call it a comb function in time.
Mathematically, we write this as the sum of delta of t minus n delta t, summed over all integers n.#That just means we have delta functions at 0, at delta t, at 2 delta t, and so on, going both directions.
Now here’s the key:#When we take the Fourier transform of this comb in time, what do we get?
Surprisingly, we get another comb – this time in the frequency domain.#But the spacing changes – instead of delta t, the spikes are now spaced by 1 over delta t.
So if we sample more tightly in time, the frequency comb spreads out.#And if we stretch out the spacing in time, the frequency spikes get closer together.
This is exactly the scaling property we saw earlier – a tight structure in one domain leads to a broad structure in the other.
Now you might be wondering – how do we handle all these delta functions?#They’re not ordinary functions – they’re distributions or generalized functions.#And we’re dealing with an infinite sum of them.
This touches on deeper mathematics – involving convergence and rigor – but for our purposes, this formula holds and is extremely useful.#We’ll come back to this paired comb concept again in the next lecture, especially when we talk about sampling theory and the Fourier series.
So just remember:#A comb in time transforms to a comb in frequency.#And their spacing is inversely related – one over the other.

slide30:
Let’s now talk about one of the most powerful tools in the Fourier world — the convolution theorem.
In plain terms, the convolution theorem says this:
Convolution in the time domain becomes multiplication in the frequency domain.
Let me say that again, but clearly:#If you have two functions — let’s call them f of t and g of t — and you convolve them together, then their Fourier transform is simply the product of their individual transforms.
So:#If the Fourier transform of f of t is capital F of s,#and the transform of g of t is capital G of s,#then the transform of f convolved with g is F of s times G of s.
That’s the core idea.
Now why does this matter?
Because convolution — though useful — is often hard to compute directly.#It involves flipping one function, shifting it, multiplying, and integrating.
That’s a lot of work!
But thanks to this theorem, we can skip all that.#Instead, we go to the frequency domain, multiply two functions — much simpler — and then just come back to the time domain using an inverse transform.
Let’s walk through a concrete example.
Think about a gate function — also called a rectangular pulse. Its Fourier transform is a sinc function — which looks like a smooth wave with side ripples.
Now, if we take two gate functions and convolve them in time, we get a triangle function.
So what happens in the frequency domain?
Each gate becomes a sinc function. When we multiply the two sinc functions together, we get sinc squared.
That is — sinc of s times sinc of s gives you sinc squared of s.
So the triangle function — which was a convolution of two gates — has a Fourier transform equal to sinc squared.
That’s a beautiful result, and it all comes from the convolution theorem.
It also shows how Fourier analysis turns complicated operations into simpler ones, especially in engineering, signal processing, and physics.

slide31:
Now let’s take a moment to understand why the convolution theorem actually works. What’s the reasoning behind it?
So far, we've seen that the Fourier transform of the convolution of two functions — say, f of t and g of t — turns into a simple multiplication of their individual transforms.
But why is that true?
Let’s walk through the logic behind it, one step at a time.
We start with this:#Multiply the Fourier transform of f of t by the Fourier transform of g of t. That means you're multiplying two integrals. One for g of t times a complex exponential, and the other for f of x times the same kind of exponential — but with a different variable.
To combine them, we use different dummy variables inside the integrals — like t in one and x in the other — so they don't interfere. Then we combine them into a double integral.
Now we group the exponential terms. The two exponentials become a single exponential with t plus x in the exponent.
This sets up our key trick: we change variables. In the inner integral, let’s say#u equals t plus x. That means t equals u minus x, and since we're just changing variables inside an integral, the limits stay the same.
So now, what used to be g of t becomes g of u minus x, and the exponential becomes a function of u. This lets us express the entire inside of the integral in terms of u, and f of x stays outside.
At this point, we switch the order of integration — that’s allowed because the functions we’re working with are well-behaved. So now we’re integrating over u, with x nested inside.
What you get in the end is this:#You have e to the power minus 2 pi i s u — that’s the same complex exponential from the Fourier formula — and it’s multiplied by the convolution of g and f at position u.

That’s the key. The inner integral has now become the convolution of g and f, evaluated at u.
So finally, this whole thing is just the Fourier transform of that convolution. That’s the magic of it.
Multiplication in the Fourier domain really does correspond to convolution in the time domain — and now you’ve seen where it comes from.
Of course, these steps are dense, and I encourage you to go over them again slowly. Grab a coffee, sit back, and follow the substitutions. Each step follows cleanly from the last — and the full picture is quite elegant.

slide32:
So let’s take a step back and ask—why is convolution in the time domain equivalent to multiplication in the frequency domain?
To understand this, let’s think about a shift-invariant linear system. If you feed a sinusoidal signal into such a system, the output will also be sinusoidal. And—here’s the key—it will be at the same frequency. The system might change the amplitude or the phase, but it doesn’t generate new frequencies. That’s the hallmark of a shift-invariant, or time-invariant, linear system.
This leads to a deep and beautiful conclusion. If sinusoidal signals are preserved in shape and frequency, then the effect of the system on each frequency component can be described by a single number—just a scaling factor that depends on frequency.
So, imagine you take any general signal, and you decompose it into a sum of sinusoidal components. The system acts on each component individually, scaling each one by a different amount. That’s exactly what multiplication looks like in the Fourier domain. Each frequency is being multiplied by its own gain factor. That’s why convolution in the time domain turns into multiplication in the frequency domain.
Now, here's something important: this only works because of that special property of sinusoids. Among all functions, only sinusoids maintain their shape when passing through a shift-invariant linear system. Delta functions, for instance, don’t behave that way—they get transformed into something entirely different. But a cosine remains a cosine. A sine remains a sine. Just possibly scaled or phase-shifted.
And this is what makes the Fourier transform so unique. It’s the only transform for which the convolution theorem holds—because only sinusoids have this invariance.
You might wonder—could we define a similar convolution theorem using other transforms, like wavelet transforms or Hadamard transforms? The answer is no. Those basis functions don’t have this shape-preserving property under linear shift-invariant systems. So the convolution theorem doesn’t hold for them.
Now if you’re curious, you can actually find some advanced papers on this. Try searching for “characterization of the convolution theorem.” You’ll find rigorous mathematical proofs that ultimately say the same thing: the convolution theorem is special to the Fourier transform.
But you don’t need pages of math to get the intuition. If a function keeps its shape through a system, and we can describe the system’s effect with a frequency-dependent multiplier, then that’s multiplication in the Fourier domain.
And if you're passionate about the math behind this, you could even write a short paper exploring this idea. I'd be happy to discuss it with you.

slide33:
Now let’s talk about another important result in Fourier analysis — Parseval’s Identity.
This identity says that the total energy of a signal, when measured in the time domain, is exactly equal to the total energy in the frequency domain.
Mathematically, it's written as:#the integral of the square of the function f of t over all time, equals the integral of the square of its Fourier transform — f hat of s — over all frequencies.
Let me give you some intuition here.
Suppose you have a function, maybe it’s an electric signal or a sound wave. In the time domain, the quantity f of t squared represents the power of that signal at each instant. When you integrate this over time, you’re summing up all that

power — you get the total energy.
Now, that same signal can be broken down into its frequency components using the Fourier transform. In that domain, each component has an amplitude, and squaring that amplitude gives you the energy associated with that frequency. Integrating over all frequencies gives you the total energy — again.
So Parseval’s Identity tells us something very beautiful: energy is conserved between the time domain and the frequency domain.
This has a strong connection to what you learned in physics. Think about alternating current. If you have a sinusoidal current flowing through a resistor, the instantaneous power is proportional to the current squared. And when you integrate that over time, you get the total energy consumed by the resistor.
Now, from a geometric perspective, you can think of f of t as a vector in an infinite-dimensional space. That sounds abstract, but it’s just like a regular vector — except with infinitely many components. And the length of this vector is found by summing up the square of each component — just like in ordinary geometry.
The Fourier transform is a kind of coordinate transformation — like a rotation of the space. So when we go from the time domain to the frequency domain, we’re rotating the vector. But that doesn’t change its length. So geometrically, Parseval’s Identity simply says: the length of the signal vector is preserved under the Fourier transform.
You can also prove this identity using standard calculus and the definition of the Fourier transform. And yes, that’s a good exercise. But I want you to also grasp the meaning behind it.
So next time you see this identity, remember — it’s not just a formula. It tells you that energy, or vector length, is preserved when moving between the time and frequency domains. And that is both physically and mathematically very powerful.

slide34:
So far, we’ve been talking about transformations in one dimension — mainly between time and frequency using the Fourier transform. But to truly understand what's going on, it helps to think geometrically.
Let’s start with something more familiar — a simple two-dimensional rotation. Imagine we have a point in the 2D plane, described by coordinates x and y in a standard X Y coordinate system. Now, suppose we want to express this same point in a new coordinate system, X-prime and Y-prime, which is rotated by some angle phi from the original axes.
As the diagram shows, the new coordinates — x-prime and y-prime — can be computed from the original ones using the rotation formulas.
So x-prime equals x times cosine phi plus y times sine phi,#and y-prime equals negative x times sine phi plus y times cosine phi.
You can also write this transformation in matrix form. It’s a simple 2-by-2 rotation matrix.
Now why is this relevant?
Because this kind of rotation is a basic example of an orthonormal transformation, just like the Fourier transform. When we rotate the coordinate system, we’re not changing the length of the vector — we’re just viewing it from a different angle. The structure stays the same, just expressed differently.
This gives us a visual and intuitive way to think about more abstract transformations. In higher dimensions — or even infinite dimensions — the concept is similar. We’re taking a function, representing it in a new coordinate system, and the transformation preserves important properties like length and energy.
That’s why understanding this simple 2D rotation helps us better appreciate what the Fourier transform is doing. It’s not magic — it’s geometry, just in a much bigger space.

slide35:
Now that we understand the Fourier transform in one dimension, let’s take it a step further — to two dimensions.
Just like in the 1D case, the idea is to take a function — in this case, a function of two variables, x and y — and express it in terms of its frequency content. But instead of waves traveling along a line, now we’re dealing with

wave patterns that extend in all directions across a plane.
The 2D Fourier transform lets us analyze how these wave components — traveling in the x direction, the y direction, or even diagonally — contribute to the overall structure of the signal or image.
The formula here tells us how to compute the 2D Fourier transform. The function f of x and y is transformed into capital F of u and v, where u and v represent the spatial frequencies in the horizontal and vertical directions.
And just like before, there’s an inverse formula that lets us go back from the frequency domain to the original spatial domain.
Now, why is this so useful?
Because in many real-world applications — like medical imaging, computer vision, or signal processing — we deal with 2D data. Think of an image, for example. Each pixel represents a value at some x and y location. When we apply the 2D Fourier transform, we can analyze the texture, orientation, and frequency content of that image.
On the left side of the slide, we see a simple geometric object — like a bar. On the right is its 2D Fourier spectrum. Notice how the orientation of the object affects the direction of the frequency response.
So just like in one dimension, we’re decomposing a complex signal into simpler, sinusoidal components — but now in two dimensions. And that opens up a whole new world of possibilities.

slide36:
Let’s now look at a powerful application of the 2D Fourier transform — noise suppression.
Take a look at the image on the left. It’s a noisy version of a familiar test image — lots of random graininess, especially in the background and darker areas.
Now, if we perform a 2D Fourier transform — which we indicate here by "F T" — we move from the spatial domain to the frequency domain. The result is this bottom-left image, which shows the frequency spectrum of the noisy image.
Notice something important: most of the meaningful image information is concentrated around the center, which corresponds to the low-frequency components. But the noise is spread out — especially toward the edges — as high-frequency speckles.
And this gives us an idea.
What if we simply suppress or remove those high-frequency components? That’s what we’re doing on the bottom-right: we apply a mask — setting the noisy high-frequency regions to zero, while keeping only the central, low-frequency region.
Then we perform the inverse Fourier transform — marked here as "I F T" — to convert back to the spatial domain.
And just like that, we get a much cleaner version of the original image. Most of the noise is gone, and the essential structure is preserved.
This is one of the key strengths of working in the frequency domain. Certain operations — like filtering or noise removal — can be done more easily, more effectively, and sometimes more intuitively after transformation.

slide37:
Let’s take this a step further and look at low-pass and high-pass filtering using the Fourier transform.
Start with the image on the far left — this is our original image of a building. Below it, you see the corresponding frequency spectrum. It contains both low-frequency components, which encode smooth variations like brightness and shading, and high-frequency components, which capture edges and fine details.
Now look at the middle column.
Here, we’ve applied a low-pass filter. That means we kept only the low-frequency components — those near the center of the Fourier spectrum — and removed the rest. You can see the filtered spectrum just below. The result, shown in the middle image above, is a smoothed version of the original — the fine details are gone, and only the broad, soft structures remain.
Next, we move to the rightmost column.
This time we’ve done the opposite. We removed the low frequencies and kept only a selected band of high-frequency components. You can see that in the frequency plot — a ring where the center is zeroed out. The corresponding spatial image

above now reveals only the edges — the sharp transitions in intensity.
So by simply choosing which frequency bands to preserve or remove, we can control the kind of information we keep — soft versus sharp, background versus boundary.
This is another powerful reason why we often work in the frequency domain — filtering becomes intuitive and flexible.

slide38:
Let’s now walk through an important example: the two-dimensional rectangle function, centered at the origin, with side lengths X and Y.
This function is quite simple in the spatial domain — it's just a bright rectangular block surrounded by zeros. But when we take its 2D Fourier transform, something fascinating happens.
The resulting function in the frequency domain is a product of two sinc functions — one in the u direction and one in the v direction. And remember, sinc functions come from the Fourier transform of a box function in one dimension. So when we go to two dimensions, the result is just a multiplication of two of them — one along each axis.
You can see this result visualized in the bottom right. There's a sharp peak in the center — that’s the low-frequency content — and then it decays with oscillations outward, characteristic of sinc behavior.
Now here’s something more interesting: when we move into higher dimensions, something new becomes possible — rotation. In 1D, you can only flip a signal left or right, but there's no true concept of rotation. However, in 2D or 3D, you can rotate the object, and that rotation affects the frequency domain in a meaningful way.
This rotational property is one of the key advantages of analyzing signals in higher dimensions. It allows us to understand and manipulate how orientation in space translates into patterns in frequency.
So this simple example — a rectangular function — gives us a powerful insight into the structure of Fourier transforms in two dimensions.

slide39:
Let’s talk about something elegant and powerful — the rotation property of the Fourier transform.
Here’s the idea: if you rotate a two-dimensional function in space, then its Fourier transform also rotates — by the exact same angle and in the same direction. That means the structure of the frequency content preserves the orientation of the original signal.
For instance, if you rotate an image by 30 degrees counterclockwise, the entire frequency spectrum also rotates by 30 degrees counterclockwise. The shape and the distribution of frequencies remain the same, just oriented differently. This behavior is not just limited to 2D — it holds in any number of dimensions.
Why does this happen? Well, it comes from the mathematics of how the Fourier transform is defined. If we apply a rotation matrix — let’s call it R theta — to the spatial variable, we see that the frequency variable ends up rotated in exactly the same way. It’s a symmetry property built into the transform itself.
Now, the proof is shown here, and you’re welcome to follow through it if you like. It involves a change of variables and some linear algebra, but the takeaway is beautifully simple: rotate the function, and its spectrum rotates too.
This rotational invariance is one of the reasons Fourier analysis is so useful in imaging, especially in applications like object detection, texture analysis, and pattern recognition — where orientation should not change the essential features.

slide40:
Let me now give you a more intuitive, geometric explanation for the rotation property of the Fourier transform.
Suppose you have a 2D function — like the white rectangular bar you see on the top-left image. This function has a certain frequency structure, and that’s captured in its Fourier transform, shown to the right.
Now, here’s the key idea: each point in the Fourier spectrum corresponds to a sinusoidal wave component — a wave moving in a specific direction and with a

certain frequency. When you sum up all these wave components, you reconstruct the original image.
So what happens when we rotate the original function? Say we rotate it 45 degrees counterclockwise. Well, to reconstruct this rotated function using the same wave-based approach, each of those original wave components must also rotate by the same 45 degrees. That’s the only way they’ll still combine to match the rotated shape.
As you can see in the second row, the rotated spatial function leads to a rotated frequency spectrum. Both are rotated by the same angle, preserving the overall structure. This is what we mean by the rotation property of the Fourier transform.
You can verify this rigorously with mathematics, but this geometric reasoning gives us an intuitive, visual understanding of what’s happening.
This is part of a much deeper concept: the duality of information. You can think of an image as made of individual points — pixels or voxels — or as made of sinusoids, each with different directions and frequencies. These two views are completely interchangeable, and Fourier analysis is the bridge between them.
That’s why this topic is so foundational. It’s not just about imaging. Fourier analysis underlies many areas — in physics, engineering, and even mathematics — wherever we want to understand structure in terms of frequency.
Now let’s take a moment to summarize what we’ve learned so far.

slide41:
Maybe you'll enjoy this little logo — it’s a compact way to capture the spirit of what we've been discussing.
On the left, we see the Greek letter delta. That symbolizes the delta function, which represents the particle-like, pointwise nature of signals — sharp, localized, discrete. It's our way of representing structure in the spatial or time domain.
Next to it is e to the power i-theta, a complex exponential. This captures the wave nature — smooth, oscillating, continuous. It represents sinusoids, and it's central to how we describe frequency, phase, rotation, and oscilation in the complex plane, describing wave propagation.
So together, delta and e to the i-theta — one representing particles, the other waves — form a kind of duality. And this duality runs throughout everything we’ve covered: convolution, linear systems, Fourier series, and Fourier transforms.
The real power of Fourier analysis is in how it unifies these two views. A signal can be described by where things happen — using delta functions — or by how things oscillate — using sinusoids. This interplay is what makes the field so rich and so widely applicable.
That’s why I like this little symbol. It’s simple, but it reflects deep ideas that are at the heart of signal processing and medical imaging. We will understand this logo more and more as we unravel its meaning in this and next several lectures.

slide42:
Alright, to wrap up today’s session, here’s your homework assignment.
First, I’d like you to read about the uncertainty principle of the Fourier transform. This is a fascinating and important concept. Summarize it in your own words, but keep it concise — no more than three sentences.
Second, I want you to analytically compute the Fourier transform of this function: the exponential of b times t, multiplied by the step function u of negative t. Here, b is a positive constant, and u of t is defined as 1 for positive time and 0 otherwise.
The due date is one week from now, by midnight next Friday. Please make sure to upload your report to MLS.
And before we end, I’ve heard from a few students that the Fourier series material was a bit confusing. If that’s you, don’t hesitate — come talk to me now or reach out later. I’m happy to help clarify anything.
That’s all for today — thank you!