PROFESSOR: Last time, we talked about the representation of linear time-invariant systems through the convolution sum in the discrete-time case and the convolution integral in the continuous-time case. Now, although the derivation was relatively straightforward, in fact, the result is really kind of amazing because what it tells us is that for linear time-invariant systems, if we know the response of the system to a single impulse at t = 0, or in fact, at any other time, then we can determine from that its response to an arbitrary input through the use of convolution.

Furthermore, what we'll see as the course develops is that, in fact, the class of linear time-invariant systems is a very rich class. There are lots of systems that have that property. And in addition, there are lots of very interesting things that you can do with linear time-invariant systems.

In today's lecture, what I'd like to begin with is focusing on convolution as an algebraic operation. And we'll see that it has a number of algebraic properties that in turn have important implications for linear time-invariant systems. Then we'll turn to a discussion of what the characterization of linear time-invariant systems through convolution implies, in terms of the relationship, of various other system properties to the impulse response.

Let me begin by reminding you of the basic result that we developed last time, which is the convolution sum in discrete time, as I indicate here, and the convolution integral in continuous time. And what the convolution sum, or the convolution integral, tells us is how to relate the output to the input and to the system impulse response. And the expression looks basically the same in continuous time and discrete time.

And I remind you also that we talked about a graphical interpretation, where essentially, to graphically interpret convolution required, or was developed, in terms of taking the system impulse response, flipping it, sliding it past the input, and positioned appropriately, depending on the value of the independent variable for which we're computing the convolution, and then multiplying and summing in the discrete-time case, or integrating in the continuous-time case.

Now convolution, as an algebraic operation, has a number of important properties. One of the properties of convolution is that it is what is referred to as *commutative.* Commutative means that we can think either of convolving x with h, or h with x, and the order in which that's done doesn't affect the output result. So summarized here is what the commutative operation is in discrete time, or in continuous time. And it says, as I just indicated, that x[n] convolved with h[n] is equal to h[n] convolved with x[n]. Or the same, of course, in continuous time.

And in fact, in the lecture last time, we worked an example where we had, in discrete time, an impulse response, which was an exponential, and an input, which is a unit step. And in the text, what you'll find is the same example worked, except in that case, the input is the exponential. And the system impulse response is the step. And that corresponds to the example in the text, which is example 3.1.

And what happens in that example is that, in fact, what you'll see is that the same result occurs in example 3.1 as we generated in the lecture. And there's a similar comparison in continuous time. This example was worked in lecture. And this example is worked in the text. In other words, the text works the example where the system impulse response is a unit step. And the input is an exponential.

All right, now the commutative property, as I said, tells us that the order in which we do convolution doesn't affect the result of the convolution. The same is true for continuous time and discrete time. And in fact, for the other algebraic properties that I'll talk about, the results are exactly the same for continuous time and discrete time.

So in fact, what we can do is drop the independent variable as an argument so that we suppress any kind of difference between continuous and discrete time. And suppressing the independent variable, we then state the commutative property as I've rewritten it here. Just x convolved with h equals h convolved with x. The same in continuous time and discrete time.

Now, the derivation of the commutative property is, more or less, some algebra which you can follow through in the book. It involves some changes of variables and some things of that sort. What I'd like to focus on with that and other properties is not the derivation, which you can see in the text, but rather the interpretation.

So we have the commutative property, and now there are two other important algebraic properties. One being what is referred to as the associative property, which tells us that if we have x convolved with the result of convolving h1 with h2, that's exactly the same as x convolved with h1, and that result convolved with h2. And what this permits is for us to write, for example, x convolved with h1 convolved with h2 without any ambiguity because it doesn't matter from the associative property how we group the terms together.

The third important property is what is referred to as the distributive property, namely the fact that convolution distributes over addition. And what I mean by that is what I've indicated here on the slide: that if I think of x convolved with the sum of h1 and h2, that's identical to first convolving x with h1, also convolving x with h2, and then adding the two together. And that result will be the same as this result.

So convolution is commutative, associative, and it distributes over addition. Three very important algebraic properties. And by the way, there are other algebraic operations that have that same property. For example, multiplication of numbers is likewise commutative, associative, and distributive.

Now let's look at what these three properties imply specifically for linear time-invariant systems. And as we'll see, the implications are both very interesting and very important. Let's begin with the commutative property.

And consider, in particular, a system with an impulse response h. And I represent that by simply writing the h inside the box. An input x and an output, then, which is x * h. Now, since this operation is commutative, I can write instead of x * h, I can write h * x. And that would correspond to a system with impulse response x, and input h, and output then h * x.

So the commutative property tells us that for a linear time-invariant system, the system output is independent of which function we call the input and which function we call the impulse response. Kind of amazing actually. We can interchange the role of input and impulse response. And from an output point of view, the output or the system doesn't care.

Now furthermore, if we combine the commutative property with the associative property, we get another very interesting result. Namely that if we have two linear time-invariant systems in cascade, the overall system is independent of the order in which they're cascaded. And in fact, in either case, the cascade can be collapsed into a single system.

To see this, let's first consider the cascade of two systems, one with impulse response h1, the other with impulse response h2. And the output of the first system is then x * h1. And then that is the input to the second system. And so the output of that is that result convolved with h2.

So this is the result of cascading the two. And now we can use the associative property to rewrite this as x * (h1 * h2), where we group these two terms together. And so using the associative property, we now can collapse that into a single system with an input, which is x, and impulse response, which is h1 * h2. And the output is then x convolved with the result of those two convolved.

Next, we can apply the commutative property. And the commutative property says we could write this impulse response that way, or we could write it this way. And since convolution is commutative, the resulting output will be exactly the same. And so these resulting outputs will be exactly the same.

And now, once again we can use the associative property to group these two terms together. And x * h2 corresponds to putting x first through the system h2 and then that output through the system h1. And so finally applying the associative property again, as I just outlined, we can expand that system back into two systems in cascade with h2 first and h1 second,

OK, well that involves a certain amount of algebraic manipulation. And that is not the algebraic manipulation that is important. It's the result that it's important. And what the result says, to reiterate, is if I have two linear time-invariant systems in cascade, I can cascade them in any order, and the result is the same.

Now you might think, well gee, maybe that actually applies to systems in general, whether you put them this way or that way. But in fact, as we talked about last time, and I illustrated with an example, in general, if the systems are not linear and time-invariant, then the order in which they're cascaded is important to the interpretation of the overall system. For example, if one system took the square root and the other system doubled the input, taking the square root and then doubling gives us a different answer than doubling first and then taking the square root. And of course, one can construct much more elaborate examples than that.

So it's a property very particular to linear time-invariant systems. And also one that we will exploit many, many times as we go through this material.

The final property related to an interconnection of systems that I want to just indicate develops out of the distributive property. And what it applies to is an interpretation of the interconnection of systems in parallel. Recall that the parallel combination of systems corresponds, as I indicate here, to a system in which we simultaneously feed the input into h1 and h2, these representing the impulse responses.

And then, the outputs are summed to form the overall output. And using the fact that convolution distributes over addition, we can rewrite this as x * (h1 + h2). And when we do that then, we can recognize this as the output of a system with input x and impulse response, which is the sum of these two impulse responses. So for linear time-invariant systems in parallel, we can, if we choose, replace that interconnection by a single system whose impulse response is simply the sum of those impulse responses.

OK, now we have this very powerful representation for linear time-invariant systems in terms of convolution. And we've seen so far in this lecture how convolution and the representation through the impulse response leads to some important implications for system interconnections. What I'd like to turn to now are other system properties and see how, for linear time-invariant systems in particular, other system properties can be associated with particular properties or characteristics of the system impulse response.

And what we'll talk about are a variety of properties. We'll talk about the issue of memory, we'll talk about the issue of invertibility, and we'll talk about the issue of causality and also stability.

Well, let's begin with the issue of memory. And the question now is what are the implications for the system impulse response for a linear time-invariant system? Remember that we're always imposing that on the system. What are the implications on the impulse response if the system does or does not have memory?

Well, we can answer that by looking at the convolution property. And we have here, as a reminder, the convolution integral, which tells us how x(tau) and h(t - tau) are combined to give us y(t).

And what I've illustrated above is a general kind of example. Here is x(tau). Here is h(t - tau). And to compute the output at any time t, we would take these two, multiply them together, and integrate from -infinity to +infinity.

So the question then is what can we say about h(t), the impulse response in order to guarantee, let's say, that the output depends only on the input at time t. Well, it's pretty much obvious from looking at the graphs.

If we only want the output to depend on x(tau) at tau = t, then h(t - tau) better be non-zero only at tau = t. And so the implication then is that for the system to be memoryless, what we require is that h(t - tau) be non-zero only at tau = t.

So we want the impulse response to be non-zero at only one point. We want it to contribute something after we multiply and go through an integral. And in effect, what that says is the only thing that it can be and meet all those conditions is a scaled impulse.

So if the system is to be memoryless, then that requires that the impulse response be a scaled impulse. Any other impulse response then, in essence, requires that the system have memory, or implies that the system have memory. So for the continuous-time case then, memoryless would correspond only to the impulse response being proportional to an impulse. And in the discrete-time case, a similar statement, in which case, the output is just proportional to the input, again either in the continuous-time or in the discrete-time case.

All right. Now we can turn our attention to the issue of system invertibility. And recall that what is meant by invertibility of a system, or the inverse of a system. The inverse of a system is a system, which when we cascade it with the one that we're inquiring about, the overall cascade is the identity system. In other words, the output is equal to the input.

So let's consider a system with impulse response h, input is x. And let's say that the impulse response of the inverse system is h_i, and the output is y. Then, the output of this system is x * (h * h_i). And we want this to come out equal to x. And what that requires than is that this convolution just simply be equal to an impulse, either in the discrete-time case or in the continuous-time case.

And under those conditions then, h_i is equal to the inverse of h. Notationally, by the way, it's often convenient to write instead of h_i as the impulse response of the inverse, you'll find it convenient often and more typical to write as the inverse, instead of h_i, h^(-1). And we mean by that the inverse impulse response. And one has to be careful not to misinterpret this as the reciprocal of h(t) or h(n). What's meant in this notation is the inverse system.

Now, if h_i is the inverse of h, is h the inverse of h_i? Well, it seems like that ought to be plausible or perhaps make sense. The question, if you believe that the answer is yes, is how, in fact, do you verify that? And I'll leave it to you to think about it. The answer is yes,

that if h_i is the inverse of h, then h is the inverse of h_i. And the key to showing that is to exploit the fact that when we take these systems and cascade them, we can cascade them in either order.

All right now let's turn to another system property, the property of stability. And again, we can tie that property directly to issues related, in particular, to the system impulse response.

Now, stability is defined as we've chosen to define it and as I've defined it previously, as bounded-input bounded-output stability. In other words, for every bounded input is a bounded output.

What you can show-- and I won't go through the algebra here; it's gone through in the book-- is that a necessary and sufficient condition for a linear time-invariant system to be stable in the bounded-input bounded-output sense is that the impulse response be what is referred to as absolutely summable. In other words, if you take the absolute values and sum them over infinite limits, that's finite.

Or in the continuous-time case, that the impulse response is absolutely integrable. In other words, if you take the absolute values of h(t) and integrate, that's finite. And under those conditions, the system is stable. If those conditions are violated, then for sure, as you'll see in the text, the system is unstable. So stability can also be tied to the system impulse response.

Now, the next property that I want to talk about is the property of causality. And before I do, what I'd like to do is introduce a peripheral result that we'll then use-- when we talked about causality-- namely what's referred to as the zero input response of a linear system.

The basic result, which is a very interesting and useful one, is that for a linear system-- and in fact, it's whether it's time-invariant or not, that this applies-- if you put nothing into it, you get nothing out of it. So if we have an input x(t) which is 0 for all t, and if the output of that system is y(t), if the input is 0 for all time, then the output likewise is 0 for all time. That's true for continuous time, and it's also true for discrete time.

And in fact, to show that result is pretty much straightforward. We could do it either by using convolution, which would, of course, be associated with linearity and time invariance. But in fact, we can show that property relatively easily by simply using the fact that, for a linear system, what we know is that if we have an input x(t) with an output y(t), then if we scale the input, then the output scales accordingly.

Well, we can simply choose, as the scale factor, a = 0. And if we do that, it says put nothing in, you get nothing out. And what we'll see is that has some important implications in terms of causality.

It's important, though, while we're on it, to stress that not every system, obviously, has that property. That if you put nothing in, you get nothing out. A simple example is, let's say, a battery, let's say not connected to anything. The output is six volts no matter what the input is. And it of course then doesn't have this zero response to a zero input. It's very particular to linear systems.

All right, well now let's see what this means for causality. To remind you, causality says, in effect, that the system can't anticipate the input. That's what, basically, causality means.

When we talked about it previously, we defined it in a variety of ways, one of which was the statement that if two inputs are identical up until some time, then the outputs must be identical up until the same time. The reason, kind of intuitively, is that if the system is causal-- so it can't anticipate the future-- it can't anticipate whether these two identical inputs are sometime later going to change from each other or not.

So causality, in general, is simply this statement, either continuous-time or discrete-time. And now, so let's look at what that means for a linear system. For a linear system, what that corresponds to or could be translated to is a statement that says that if x(t) is 0, for t less than t_0, then y(t) must be 0 for t less than t_0 also.

And so what that, in effect, says, is that the system-- for a linear system to be causal, it must have the property sometimes referred to as initial rest, meaning it doesn't respond until there's some input that happens. That it's initially at rest until the input becomes non-zero.

Now, why is this true? Why is this a consequence of causality for linear systems? Well, the reason is we know that if we put nothing in, we get nothing out. If we have an input that's 0 for t less than t_0, and the system can't anticipate whether that input is going to change from 0 or not, then the system must generate an output that's 0 up until that time, following the principle that if two inputs are identical up until some time, the outputs must be identical up until the same time.

So this basic result for linear systems is essentially a consequence of the statement that for a linear system, zero in gives us zero out. Now, that tells us how to interpret causality for linear systems. Now, let's proceed to linear time-invariant systems.

And in fact, we can carry the point one step further. In particular, a necessary and sufficient condition for causality in the case of linear time-invariant systems is that the impulse response be 0, for t less than 0 in the continuous-time case, or for n less than 0 in the discrete-time case. So for linear time-invariant systems, causality is equivalent to the impulse response being 0 up until t or n equal to 0.

Now, to show this essentially follows by first considering why causality would imply that this is true. And that follows because of the straightforward fact that the impulse itself is 0 for t less than 0. And what we saw before is that for any linear system, causality requires that if the input is 0 up until some time, the output must be 0 up until the same time. And so that's showing the result in one direction.

To show the result in the other direction, namely to show that if, in fact, the impulse response satisfies that condition, then the system is causal. While I won't work through it in detail, it essentially boils down to recognizing that in the convolution sum, or in the convolution integral, if, in fact, that condition is satisfied on the impulse response, then the upper limit on the sum, in the discrete-time case, changes to n. And in the continuous-time case, changes to t.

And that, in effect, says that values of the input only from -infinity up to time n are used in computing y[n]. And a similar kind of result for the continuous-time case y(t).

OK, so we've seen how the impulse response, or rather how certain system properties in the linear time-invariant case can, be converted into various requirements on the impulse response of a linear time-invariant system, the impulse response being a complete characterization. Let's look at some particular examples just to kind of cement the ideas further.

And let's begin with a system that you've seen previously, which is an accumulator. An accumulator, as you recall, has an output y[n], which is the accumulated value of the input from -infinity up to n. Now, you've seen in the impulse in a previous lecture, or rather in the video course manual for a previous lecture, that an accumulator is a linear time-invariant system. And in fact, its impulse response is the accumulated values of an impulse. Namely, the impulse response is equal to a step.

So what we want to answer is, knowing what that impulse response is, what some properties are of the accumulator. And let's first ask about memory. Well, we recognize that the impulse response is not simply an impulse. In fact, it's a step. And so this implies what? Well, it implies that the system has memory.

Second, the impulse response is 0 for n less than 0. That implies that the system is causal. And finally, if we look at the sum of the absolute values of the impulse response from -infinity to +infinity, this is a step. If we accumulate those values over infinite limits, then that in fact comes out to be infinite. And so what that implies, then, is that the accumulator is not stable in the bounded-input bounded-output sense.

Now I want to turn to some other systems. But while we're on the accumulator, I just want to draw your attention to the fact, which will kind of come up in a variety of ways again later, that we can rewrite the equation for an accumulator, the difference equation, by recognizing that we could, in fact, write the output as the accumulated values up to time n - 1 and then add on the last value. And in fact, if we do that, this corresponds to y[n-1].

And so we could rewrite this difference equation as y[n] = y[n-1] + x[n]. So the output is the previously-computed output plus the input. Expressed that way, what that corresponds to is what is called a *recursive difference equation.* And different equations will be a topic of considerable emphasis in the next lecture.

Now, does an accumulator have an inverse? Well, the answer is, in fact, yes. And let's look at what the inverse of the accumulator is. The impulse response of the accumulator is a step. To inquire about its inverse, we inquire about whether there's a system, which when we cascade the accumulator with that system, which I'm calling its inverse, we get an impulse out.

Well, let's see. The impulse response of the accumulator is a step. We want to put the step into something and get out an impulse. And in fact, what you recall from the lecture in which we introduced steps and impulses, the impulse is, in fact, the first difference of the units step. So we have a difference equation that describes for us how the impulse is related to the step.

And so if this system does this, the output will be that, an impulse. And so if we think of x_2[n] as the input and y_2[n] as the output, then the difference equation for the inverse system is what I've indicated here. And if we want to look at the impulse response of that, we can then inquire as to what the response is with an impulse in.

And what develops in a straightforward way then is delta[n], which is our impulse input, minus delta[n-1] is equal to the impulse response of the inverse system. So I'll write that as h^(-1)[n] *(h-inverse of n).*

Now, we have then that the accumulator has an inverse. And this is the inverse. And you can examine issues of memory, stability, causality, et cetera. What you'll find is that the system has memory, the inverse accumulator. It's stable, and it's causal.

And it's interesting to note, by the way, that although the accumulator was an unstable system, the inverse of the accumulator is a stable system. In general, if the system is stable, its inverse does not have to be stable or vice versa. And the same thing with causality.

OK now, there are a number of other examples, which, of course, we could discuss.

And let me just quickly point to one example, which is a difference equation, as I've indicated here. And as we'll talk about in more detail in our next lecture, where we'll get involved in a fairly detailed discussion of linear constant-coefficient difference and differential equations, this falls into that category.

And under the imposition of what's referred to as initial rest, which corresponds to the response being 0 up until the time that the input becomes non-zero, the impulse response is a^n times u[n]. And something that you'll be asked to think about in the video course manual is whether that system has memory, whether it's causal, and whether it's stable.

And likewise, for a linear constant coefficient differential equation, the specific one that I've indicated here, under the assumption of initial rest, the impulse response is e^(-2t) times u(t). And in the video course manual again, you'll be asked to examine whether the system has memory, whether it's causal, and whether it's stable.

OK well, as I've indicated, in the next lecture we'll return to a much more detailed discussion of linear constant-coefficient differential and difference equations. Now, what I'd like to finally do in this lecture is use the notion of convolution in a much different way to help us with a problem that I alluded to earlier. In particular, the issue of how to deal with some of the mathematical difficulties associated with impulses and steps.

Now, let me begin by illustrating kind of what the problem is and an example of the kind of paradox that you sort of run into when dealing with impulse functions and step functions. Let's consider, first of all, a system, which is the identity system. And so the output is, of course, equal to the input. And again, we can talk about that either in continuous time or in discrete time.

Well, we know that the function that you convolve with a signal to retain the signal is an impulse. And so that means that the impulse response of an identity system is an impulse. Makes logical sense.

Furthermore, if I take two identity systems and cascade them, I put in an input, get the same thing out of the first system. That goes into the second system. Get the same thing out of the second. In other words, if I have two identity systems in cascade, the cascade, likewise, is an identity system. In other words, this overall system is also an identity system.

And the implication there is that the impulse response of this is an impulse. The impulse response of this is an impulse. And the convolution of those two is also an impulse. So for continuous time, we require, then, that an impulse convolved with itself is an impulse. And the same thing for discrete time.

Now, in discrete time, we don't have any particular problem with that. If you think about convolving these together, it's a straightforward mathematical operation since the impulse in discrete time is very nicely defined. However, in continuous time, we have to be somewhat careful about the definition of the impulse because it was the derivative of a step. A step has a discontinuity. You can't really differentiate at a discontinuity.

And the way that we dealt with that was to expand out the discontinuity so that it had some finite time region in which it happened. When we did that, we ended up with a definition for the impulse, which was the limiting form of this function, which is a rectangle of width Delta, and height 1 / Delta, and an area equal to 1.

Now, if we think of convolving this signal with itself, the impulse being the limiting form of this, then the convolution of this with itself is a triangle of width 2 Delta, height 1 / Delta, and area 1. In other words, this triangular function is this approximation delta_Delta(t) convolved with delta_Delta(t).

And since the limit of this would correspond to the impulse response of the identity system convolved with itself, the implication is that not only should the top function, this one, correspond in its limiting form to an impulse, but also this should correspond in its limiting form to an impulse. So one could wonder well, what is an impulse? Is it this one in the limit? Or is it this one in the limit?

Now, beyond that-- so kind of what this suggests is that in the limiting form, you kind of run into a contradiction unless you don't try to distinguish between this rectangle and the triangle. Things get even worse when you think about what happens when you put an impulse into a differentiator. And a differentiator is a very commonly occurring system.

In particular, suppose we had a system for which the output was the derivative of the input. So if we put in x(t), we got out dx(t) / dt. If I put in an impulse, or if I talked about the impulse response, what is that?

And of course, the problem is that if you think that the impulse itself is very badly behaved, then what about its derivative, which is not only infinitely big, but there's a positive-going one, and a negative-going one, and the difference between there has some area. And you end up in a lot of difficulty.

Well, the way around this, formally, is through a set of mathematics referred to as generalized functions. We won't be quite that formal. But I'd like to, at least, suggest what the essence of that formality is. And it really helps us in interpreting the impulses in steps and functions of that type.

And what it is is an operational definition of steps, impulses, and their derivatives in the following sense. Usually when we talk about a function, we talk about what the value of the function is at any instant of time. And of course, the trouble with an impulse is it's infinitely big, in zero width, and has some area, et cetera.

What we can turn to is what is referred to as an operational definition where the operational definition is related not to what the impulse is, but to what the impulse does under the operation of convolution. So what is an impulse? An impulse is something, which under convolution, retains the function. And that then can serve as a definition of the impulse.

Well, let's see where that gets us. Suppose that we now want to talk about the derivative of the impulse. Well, what we ask about is what it is operationally. And so if we have a system, which is a differentiator, and we inquire about its impulse response, which let's say we define notationally as u_1(t).

What's important about this function u_1(t) is not what it is at each value of time but what it does under convolution. What does it do under convolution? Well, the output of the differentiator is the convolution of the input with the impulse response. And so what u_1(t) does under convolution is to differentiate. And that is the operational definition.

And now, of course, we can think of extending that. Not only would we want to think about differentiating an impulse, but we would also want to think about differentiating the derivative of an impulse. We'll define that as a function u_2(t). u_2(t)-- because we have this impulse response convolved with this one is u_1(t) * u_1(t).

And what is u_2(t) operationally? It is the operation such that when you convolve that with x(t), what you get is the second derivative.

OK now, we can carry this further and, in fact, talk about the result of convolving u_1(t) with itself more times. In fact, if we think of the convulution of u_1(t) with itself k times, then logically we would define that as u_k(t). Again, we would interpret that operationally.

And the operational definition is through convolution, where this corresponds to u_k(t) being the impulse response of k differentiators in cascade. So what is the operational definition? Well, it's simply that x(t) * u_k(t) is the k derivative of x(t).

And this now gives us a set of what are referred to as singularity functions. Very badly behaved mathematically in a sense, but as we've seen, reasonably well defined under an operational definition. With k = 0, incidentally, that's the same as what we have referred to previously as the impulse. So with k 0, that's just delta(t).

Now to be complete, we can also go the other way and talk about the impulse response of a string of integrators instead of a string of differentiators. Of course, the impulse response of a single integrator is a unit step. Two integrators together is the integral of a unit step, et cetera.

And that, likewise, corresponds to a set of what are called singularity functions. In particular, if I take a string of m integrators in cascade, then the impulse response of that is denoted as u sub minus m of t. And for example, with a single integrator, u sub minus 1 of t corresponds to our unit step as we talked about previously. u sub minus 2 of t corresponds to a unit ramp, et cetera.

And there is, in fact, a reason for choosing negative values of the argument when going in one direction near integration as compared with positive values of the argument when going in the other direction, namely differentiation. In particular, we know that with u sub minus m of t, the operational definition is the mth running integral.

And likewise, u_k(t)-- so with a positive sub script-- has an operational definition, which is the derivative. So it's the kth derivative of x(t). And partly as a consequence of that, if we take u_k(t) and convolve it with u_l(t), the result is the singularity function with the subscript, which is the sum of k and l. And that holds whether this is positive values of the subscript or negative values of the subscript.

So just to summarize this last discussion, we've used an operational definition to talk about derivatives of impulses and integrals of impulses. This led to a set of singularity functions-- what I've called singularity functions-- of which the impulse and the step are two examples. But using an operational definition through convolution allows us to define, at least in an operational sense, these functions that otherwise are very badly behaved.

OK now, in this lecture and previous lectures, for the most part, our discussion has been about linear time-invariant systems in fairly general terms. And we've seen a variety of properties, representation through convolution, and properties as they can be associated with the impulse response.

In the next lecture, we'll turn our attention to a very important subclass of those systems, namely systems that are describable by linear constant-coefficient difference equations in the discrete-time case, and linear constant-coefficient differential equations in the continuous-time case. Those classes, while not forming all of the class of linear time-invariant systems, are a very important sub class. And we'll focus in on those specifically next time. Thank you.