MATH 228 lecture notes for November 1, 3, and 5

Russell Milne

October 2021

1 November 1: Using Laplace transforms to solve ODEs

Before I begin, I would like to show one more example of taking the Laplace transform. Suppose you have a piecewise continuous function, or in other words a function that is continuous everywhere except for a finite number of jump discontinuities. Remember that the Laplace transform consists of an improper integral, with the upper bound of the integral being infinity:

$$L[f(t)] = F(s) = \int_0^\infty e^{-st} f(t)\ dt \tag{1}$$

If you have a piecewise function, it can be written as a bunch of different "pieces":

$$f(t) = \begin{cases} f_1(t),\ 0 \leq t < t_1 \\ f_2(t),\ t_1 \leq t < t_2 \\ \dots \\ f_n(t),\ t_{n-1} \leq t < \infty \end{cases} \tag{2}$$

Each of these different pieces can be considered to be zero everywhere outside its particular interval (i.e. $t < t_{n-1}$ and $t \geq t_n$). If a function is zero over some interval, then integrating it over that interval will produce a result of zero. Hence, the Laplace transform of each piece is the integral over some subset of $(0, \infty)$, as follows:

$$L[f] = \int_0^{t_1} e^{-st} f_1(t)\ dt + \ldots + \int_{t_{n-1}}^\infty e^{-st} f_n(t)\ dt \tag{3}$$

The Second Shift Theorem can be thought of as a special case of this, in which there are two pieces and one of them is the zero function. One example of this in action is the Laplace transform of the unit step function, defined as $g(t) = 1$ for $0 \leq t \leq 1$ and $g(t) = 0$ everywhere else:

$$L[g] = \int_0^1 e^{-st}\ dt = \frac{1}{s}\left(1 - e^{-s}\right) \tag{4}$$

Anyway, since this is a course on ODEs, the main reason for learning about Laplace transforms in this course is to use them to solve ODEs. We have previously seen that when you take the Laplace transform of the derivative of a function, the formula reduces to something involving the Laplace transform of the original function:

$$L[f'] = sL[f] - f(0) \tag{5}$$

We can use this as a recurrence relation to get the Laplace transform of any derivative in terms of the Laplace transform of the original function. However, since we have only solved ODEs up to second order in this course, we are mostly concerned with the second derivative:

$$L[f''] = sL[f'] - f(0) = s^2L[f] - sf(0) - f'(0) \tag{6}$$

We now have all of the tools we need to solve ODEs with this. Let's start with one that we already know the solution for, just to make sure that it works. Suppose we have the initial value problem $y'' + y = 0$, $y(0) = 1$, $y'(0) = 2$. We know that this will have a solution of $y(t) = C_1 \cos t + C_2 \sin t$, and by applying the initial conditions, we get that $C_1 = 1$ and $C_2 = 2$. If we wanted to solve this using Laplace transforms, we would just take the Laplace transform of both sides in order to get the y'' term into a form with just y:

$$L[y''] + L[y] = 0 \tag{7}$$

Note that we can separate the two terms inside the Laplace operator, because it is a linear operator. Note also that the Laplace transform of 0 is 0, as mentioned above. Using our formula for $L[y'']$, and using the notation Y for $L[y]$, we get the following:

$$s^2Y - sy(0) - y'(0) + Y = 0 \tag{8}$$

The goal now becomes to find Y, as it might be the Laplace transform of some function that we know. To do this, we can substitute in our initial conditions and group like terms:

$$Y\left(s^2 + 1\right) - s - 2 = 0 \implies Y = \frac{s+2}{s^2+1} \tag{9}$$

Remember that $L[\cos \omega t] = \frac{s}{s^2+\omega^2}$ and $L[\sin \omega t] = \frac{\omega}{s^2+\omega^2}$. Using 1 for ω and breaking apart our fraction, we get the following:

$$Y = \frac{s}{s^2+1} + 2\frac{1}{s^2+1} = L[\cos t + 2\sin t] \tag{10}$$

When we "invert" the Laplace transform that we ended up with, we get exactly the solution that we were looking for.

We can also use this method to solve ODEs with forcing functions. For instance, let's use the example that we just worked through, but with a forcing function of $7e^{2t}$ instead of 0. For this DE, the complementary solution is still

$y_c = C_1 \cos t + C_2 \sin t$. The particular solution can be obtained using methods that we already know, and it works out to be $y_p = \frac{7}{5}e^{2t}$. Because we have this extra term, C_1 and C_2 turn out to be slightly different, namely $\frac{-2}{5}$ and $\frac{-4}{5}$. Let's see what happens when we use the Laplace transform on this DE:

$$L[y''] + L[y] = Y\left(s^2+1\right) - s - 2 = L[7e^{2t}] = \frac{7}{s-2} \tag{11}$$

When we isolate Y, we get two terms on the right-hand side. One of these is the same as what we had earlier (for the version of this ODE without a forcing function), but the other one is new:

$$Y = \frac{s+2}{s^2+1} + \frac{7}{(s-2)(s^2+1)} \tag{12}$$

Since the Laplace transform is linear, we might be tempted to consider each term separately and claim that part of the solution is $\cos t + 2\sin t$, as before. However, the other term isn't the Laplace transform of any function that we know of, so we can't actually do this. Instead, we'll need to break it apart by using partial fractions. In the interest of speed, this part of the derivation will be skipped. The result of the partial fractions is the following:

$$Y = \frac{7}{5}\frac{1}{s-2} - \frac{2}{5}\frac{s}{s^2+1} - \frac{4}{5}\frac{1}{s^2+1} \tag{13}$$

We know what each of these terms is the Laplace transform of, so we can easily verify that the solution we got using this method is the same as the expected one.

Using the same techniques, we can apply the method of Laplace transforms to general second-order linear ODEs with constant coefficients. If we have the ODE $ay'' + by' + cy = p(t)$ with initial conditions $y(0)$ and $y'(0)$ defined, then we get the following, for Y and P the Laplace transforms of y and p:

$$Y = \frac{as + ay'(0) + by(0)}{as^2 + bs + c} + \frac{P(s)}{as^2 + bs + c} \tag{14}$$

Getting these into forms that we can use will typically involve partial fractions, and may also involve the Shift Theorems. The first term on the right-hand side usually isn't that hard to deal with. There are three separate cases for what you'll get when you attempt to use partial fractions to get it into a usable form. The first case is if as^2+bs+c has two real roots, so $as^2+bs+c = a(s-s_1)(s-s_2)$:

$$\frac{as + ay'(0) + by(0)}{as^2 + bs + c} = \frac{k_1}{s-s_1} + \frac{k_2}{s-s_2} \tag{15}$$

The second case is if $as^2 + bs + c$ has a repeated real root, so $as^2 + bs + c = a(s-s_1)^2$:

$$\frac{as + ay'(0) + by(0)}{as^2 + bs + c} = \frac{k_1}{s-s_1} + \frac{k_2 s + k_3}{(s-s_1)^2} \tag{16}$$

The third case is if $as^2 + bs + c$ has complex conjugate roots:

$$\frac{as + ay'(0) + by(0)}{as^2 + bs + c} = \frac{k_1 s + k_2}{(s - \frac{b}{2a})^2 + (\frac{c}{a} - \frac{b^2}{4a^2})} \tag{17}$$

The term with $P(s)$ is usually the more challenging one. A good first effort would be to use partial fractions and (if need be) the Shift Theorems on it. If that doesn't work, you could also use the Convolution Theorem. Suppose you have some function $H(s)$ that can be expressed as $F(s)G(s)$ for some F and G. You might think that if $H(s) = L[h(t)]$ and so on, then $h(t)$ would be equal to $f(t)g(t)$. However, this is not necessarily true. Instead of being the product of f and g, we instead get that h is the convolution of f and g. If you haven't seen convolution before, then the convolution of two functions f and g, denoted $f * g$, is a certain kind of integral:

$$f * g = \int_0^t f(u)g(t-u)\ du = \int_0^t f(t-u)g(u)\ du = g * f \tag{18}$$

Note that taking the convolution is commutative. It can also be thought of as a "reverse inner product" over function space. With the standard inner product of two functions f and g, we take the integral over all points of the product of f and g evaluated at the same point. With the convolution $f * g$, we're also multiplying a value taken by f with a value taken by g, but the points of f start from the left-hand side of the domain and move right, while the points of g start from the right-hand side of the domain and move left. (This is easiest to visualize when the convolution integral has finite bounds, due to one or both of the functions having a compact support.)

2 November 3: More on using Laplace transforms to solve ODEs, including ODEs with piecewise forcing terms

Previously, we saw a way to use Laplace transforms to solve general second-order ODEs with constant coefficients. Depending on the problem, you may find that this method requires more or less work than others (such as variation of parameters). However, there is one class of ODEs that Laplace transforms are particularly useful for solving, namely those with discontinuous forcing functions. This is because the Laplace transform of the entire forcing function can be broken up into multiple pieces (as we saw previously), and we can deal with the pieces using the Shift Theorems.

Let's see a very simple example of this in action. Suppose that we have a first-order ODE with a piecewise continuous forcing function. For instance, we may be interested in calculating the velocity of some object if we apply some constant level of force to it up until a specified time T, and then a different level of force to it for $t > T$. We can get the general form of this ODE from Newton's law:

$$F = ma = m\frac{dv}{dt} \implies \frac{dv}{dt} = \begin{cases} \frac{1}{m}F_1, & 0 \le t \le T \\ \frac{1}{m}F_2, & t > T \end{cases} \tag{19}$$

Because the integral of a constant is just a linear function, we can determine the solution of this ODE without any advanced techniques. v will increase linearly at a rate $\frac{F_1}{m}$ from time 0 to time T, and then at a rate $\frac{F_2}{m}$ thereafter. This can be written in the following way, assuming we know the initial condition $v(0)$:

$$v(t) = \frac{F_1}{m}t + \mathrm{H}(t-T)\frac{F_2 - F_1}{m}(t-T) + v(0) \tag{20}$$

Here, H represents the Heaviside step function. This function, named after the English applied mathematician Oliver Heaviside, is equal to zero for $t < 0$ and one for $t > 0$; in this particular example, it has been shifted to the right by T units. This has the effect of "turning off" the second term in the equation when $t < T$, making $v(t)$ the much simpler function of $v(t) = \frac{F_1}{m}t + v(0)$. When $t > T$, this formulation for v means that we subtract the $\frac{F_1}{m}t$ term and are left with a function that grows linearly at a rate $\frac{F_2}{m}$. In the second term, we use $t - T$ for time dependence rather than just t, because the line $\frac{F_2 - F_1}{m}t$ passes through the origin instead of the point $(T, v(0) + \frac{F_1}{m}T)$.

Now, let's solve this using Laplace transforms. Before we do this, it's best to get the piecewise forcing function expressed in terms of Heaviside step functions. This can be done like so:

$$\frac{F_1}{m} + \frac{F_2 - F_1}{m}\mathrm{H}(t-T) = \begin{cases} \frac{1}{m}F_1, & 0 \le t \le T \\ \frac{1}{m}F_2, & t > T \end{cases} \tag{21}$$

Just like before (when we solved this ODE by inspection), we use a Heaviside step function to "turn on" F_2 and "turn off" F_1 at $t = T$. Now, we can take the Laplace transform of both sides, keeping in mind that it is a linear operator and hence we can pull out constants:

$$L[v'] = sL[v] - v(0) = \frac{F_1}{m}L[1] + \frac{F_2 - F_1}{m}L[\mathrm{H}(t-T)\cdot 1] \tag{22}$$

The first term on the right-hand side is easy to evaluate; remember that $L[1] = L[t^0]$, which we have a formula for. In the second term on the right-hand side, we are taking the Laplace transform of a shifted version of the constant function $f(t) = 1$. Thus, we can use the Second Shift Theorem to evaluate that term. We end up with the following, for $V = L[v]$:

$$sV - v(0) = \frac{F_1}{m}\frac{1}{s} + \frac{F_2 - F_1}{m}\frac{1}{s}e^{-Ts} \tag{23}$$

Isolating V gives us the following:

$$V = \frac{F_1}{ms^2} + \frac{F_2 - F_1}{ms^2}e^{-Ts} + \frac{v(0)}{s} \tag{24}$$

We know that $\frac{1}{s^2}$ is the Laplace transform of $t^1 = t$. Additionally, because of the e^{-Ts} in the middle term, we know that everything in that term will be shifted by T. (We will once again use the Heaviside step function to represent the outcome of this shift, so that the term is only nonzero for values of t larger than whatever the origin was shifted to, in this case $t = T$.) We can therefore find v, which takes exactly the same form as our expected answer:

$$v = \frac{F_1}{m}t + \frac{F_2 - F_1}{m}(t - T)\mathrm{H}(t - T) + v(0) \tag{25}$$

If we wanted to, we could also integrate another time to find the position function. The math for this is mostly the same, although we get another factor of s that we need to divide through by (from the relation that $L[x'] = sL[x]$). This gives us the following:

$$X = \frac{F_1}{ms^3} + \frac{F_2 - F_1}{ms^3}e^{-Ts} + \frac{v(0)}{s^2} + \frac{x(0)}{s} \tag{26}$$

When solving this, remember that $L[t^n] = \frac{n!}{s^{n+1}}$. This means that $\frac{1}{s^3}$ is the Laplace transform of $\frac{1}{2}x^2$, not x^2, so we get an extra factor of $\frac{1}{2}$ when inverting the Laplace transform. This yields the following result, as expected:

$$x = \frac{F_1}{2m}t^2 + \frac{F_2 - F_1}{2m}(t - T)^2\mathrm{H}(t - T) + v(0)t + x(0) \tag{27}$$

Before we dive into another example of solving ODEs using Laplace transforms, I'll take a quick intermission to explain three very important functions in applied math, which are linked by being the derivatives of each other. The first of these is the ramp function. This is defined as follows:

$$\mathrm{R}(t) = \max(0, t) = \begin{cases} 0, & t < 0 \\ t, & t \geq 0 \end{cases} \tag{28}$$

For positive input, the ramp function increases linearly, while for negative input it is flat. This makes its graph resemble a ramp coming out of the ground, leading to its name. Many of the applications of the ramp function come in machine learning. There, it is referred to as the ReLU function (for "rectified linear unit"), and within a neural network it serves to determine which inputs are propagated forward through the network and which are ignored. In this way, it is responsible for the "learning" aspect of machine learning.

The Heaviside step function is the derivative of the ramp function. It was originally developed in electrical engineering to represent the process of switching on a circuit, and can be used more generally in the modelling of any signal that is active at some times and inactive at others. (By scaling and combining Heaviside functions, we can specify the strength of the signal as well as exactly which times it is on and off.) It is formally defined as follows:

$$\mathrm{H}(t) = \begin{cases} 0, & t < 0 \\ 1, & t > 0 \end{cases} \tag{29}$$

Note that the value that the Heaviside step function takes at 0 is not defined. This is because H(0) can take a few different values, depending on whether any specific property of H is desired. For instance, if we want H to be equal to its limit at 0 from above or from below, we would take H(0) to be 1 or 0, respectively. $\frac{1}{2}$ is also a common choice, as it makes H an odd function, as well as serving as an analogue to the half-saturation constant in a logistic function (which H resembles).

Regardless of which value is chosen, the Heaviside step function is vertical at $t = 0$, so it has an infinite slope there. However, it is flat everywhere else. This means that taking its derivative would result in a function defined like this:

$$\delta(t) = \begin{cases} \infty, & t = 0 \\ 0, & t \neq 0 \end{cases} \tag{30}$$

Alternatively, we could say that this function is undefined at 0. We additionally have the constraint that the derivative of the Heaviside step function integrates to 1, as $\mathrm{H}(t) = 1 \ \forall t > 0$:

$$\int_{-\infty}^{\infty} \delta(t) \ dt = 1 \tag{31}$$

The function that satisfies both of these constraints is the Dirac delta function. It is named after Paul Dirac, who introduced it in his 1930 textbook "The Principles of Quantum Mechanics", one of the first on that subject. On the surface, the Dirac delta function appears rather abstract. So, I'll show you an example of how it can be derived, as well as a problem in ODEs where it can be used. Suppose that we have some amount of a radioactive chemical in a beaker, which will decay at a rate k. This can be modelled as $\frac{dx}{dt} = -kx$ for x the amount of chemical present, as we have seen before. Now, suppose that for some length of time starting at $t = a$ and continuing for Δ units of time, we add more of the chemical into the beaker. Suppose also that we add the chemical at a constant rate, and the total amount of chemical that we add is b. This means that our new ODE describing the amount of chemical in the beaker is as follows:

$$\frac{dx}{dt} = \begin{cases} -kx, & t < a \\ -kx + \frac{b}{\Delta}, & a \leq t \leq a + \Delta \\ -kx, & t > a + \Delta \end{cases} \tag{32}$$

At any given point in time during the interval $[a, \Delta]$, the rate that we add the chemical is $\frac{b}{\Delta}$. This is because the total amount added is b, so if we integrate just the amount added from a to Δ (ignoring the decay for a moment), we must get b. This can be summarized by the following equation:

$$\int_{a}^{a+\Delta} \frac{b}{\Delta} \ dt = \frac{b}{\Delta} \left(a + \Delta - a\right) = b \tag{33}$$

Within the context of the original ODE, we can use Heaviside step functions to express the process of adding the chemical as one that gets turned on and then back off. Specifically, we can use a positive Heaviside step function to represent the start of adding the chemical, and a negative one to represent the end of it. Using f to denote the rate at which the chemical is added, we get the following:

$$f(t) = \frac{b}{\Delta}\left(\mathrm{H}(t-a) - \mathrm{H}(t-(a+\Delta))\right) \tag{34}$$

We now have $\frac{dx}{dt} = -kx + f(t)$, and we can solve this using Laplace transforms. If we take the Laplace transform of both sides and isolate $X = L[x]$, then we get the following (for $F = L[f]$):

$$X = \frac{x(0)}{s+k} + \frac{F}{s+k} \tag{35}$$

The first term is relatively easy to invert, since $\frac{1}{s+k}$ is the Laplace transform of e^{-kt}. The second one needs to be handled using the Convolution Theorem, since it is of the form $H(s) = F(s)\frac{1}{s+k}$. We therefore get the following convolution integral that we need to evaluate in order to invert the second term:

$$f * e^{-kt} = \int_0^t f(u)e^{-k(t-u)}\ du \tag{36}$$

This needs to be broken up into two different parts, on account of the two Heaviside step functions that make up f. After some integration, we get the following for our solution x:

$$x = x(0)e^{-kt} + \frac{b}{k\Delta}\left(1 - e^{-k(t-a)}\right)\mathrm{H}(t-a) + \frac{b}{k\Delta}\left(1 - e^{-k(t-(a+\Delta))}\right)\mathrm{H}(t-(a+\Delta)) \tag{37}$$

However, what happens when we add the chemical over a smaller and smaller interval? In other words, what happens when we decrease Δ so that it goes to 0? As Δ shrinks, the rate of input ($\frac{b}{\Delta}$) gets arbitrarily large. However, we are still adding b units of the chemical no matter how small the interval gets, as the integral describing how much chemical is added does not change at all and still equals b. In other words, the limit as Δ approaches zero of the rate of input is undefined, but the limit as Δ approaches zero of how much chemical is added is b. This results in f becoming a rather paradoxical mathematical object that is undefined at a single point (escaping to $+\infty$ there) and zero everywhere else, but whose definite integral across its entire domain is equal to the finite value b. This is the Dirac delta function. If a is taken to be 0 and b is taken to be 1, then it becomes clear that the integral of this function is the Heaviside step function $\mathrm{H}(t)$. Choosing different values for a and b causes the integral to be scaled and shifted, with the result of $b\mathrm{H}(t-a)$. A method of defining the Dirac delta function that reflects its status as a limit is shown below:

$$I_\varepsilon(t) = \begin{cases} \frac{1}{\varepsilon},\ 0 \le t \le \varepsilon \\ 0,\ t > \varepsilon \end{cases} \implies \delta(t) = \lim_{\varepsilon \to 0} I_\varepsilon(t) \tag{38}$$

So, what is the Laplace transform of the Dirac delta function? Well, for any continuous function $g(t)$ as well as $I_\varepsilon(t)$ as defined above, we have the following result:

$$\int_{-\infty}^{\infty} I_\varepsilon(t) g(t)\ dt = \int_0^{\varepsilon} I_\varepsilon(t) g(t)\ dt = \frac{1}{\varepsilon} \int_0^{\varepsilon} g(t)\ dt \tag{39}$$

By the Mean Value Theorem for Integrals, we also have that within the interval $[0, \varepsilon]$ there exists some c such that the following holds:

$$\int_0^{\varepsilon} g(t)\ dt = \varepsilon \cdot g(c) \implies \int_0^{\varepsilon} I_\varepsilon(t) g(t)\ dt = g(c) \tag{40}$$

If we take the limit as ε goes to 0 (approaching the Dirac delta function), we get that c also goes to 0, since 0 eventually becomes the only value in the interval $[0, \varepsilon]$. Therefore, we get the following:

$$\lim_{\varepsilon \to 0} \int_0^{\varepsilon} I_\varepsilon(t) g(t)\ dt = g(0) \tag{41}$$

This can, of course, be shifted to $g(a)$ for any a, if the Dirac delta function is centered there instead of at 0. Since exponential functions are continuous, we can combine a few of the above results to get a statement for the Laplace transform of the Dirac delta function:

$$L[\delta(t)] = \int_0^{\infty} e^{-st} \delta(t)\ dt = e^{-s \cdot 0} = 1 \tag{42}$$

For a Dirac delta function centered at some $a \neq 0$, we get the following result instead:

$$L[\delta(t)] = \int_0^{\infty} e^{-st} \delta(t - a)\ dt = e^{-as} \tag{43}$$

This is consistent with the Second Shift Theorem, since the result can also be thought of as e^{-as} times the Laplace transform of the "unshifted" Dirac delta function (which is 1). Now that we know this, we can solve the ODE modelling the chemical in a beaker that we talked about earlier with an instantaneous addition of the chemical. The ODE now takes the following form:

$$\frac{dx}{dt} = -kx + b\delta(t - a) \tag{44}$$

Taking the Laplace transform of both sides gives us this:

$$sX - x(0) = -kX + be^{-as} \tag{45}$$

Isolating X produces the following:

$$X = \frac{x(0)}{s+k} + \frac{b}{s+k}e^{-as} \tag{46}$$

The first term is the Laplace transform of an exponential, while the second one is an exponential that has been shifted. We can invert and get the following solution for x:

$$x(t) = x(0)e^{-kt} + be^{-k(t-a)}\mathrm{H}(t-a) \tag{47}$$

We end up with two exponential decay curves that produce a jump discontinuity when graphed, as we might expect given how the problem was formulated.

3 November 5: Introduction to systems of differential equations, including a review of linear algebra

Previously, we've focused on solving single differential equations. However, in practice, many things that you may want to model using differential equations will interact with each other. This means that we will need to consider solving multiple differential equations at the same time, with each of the differential equations in question representing the rate of change of some quantity that we are interested in. For instance, suppose that we are modelling a chemical reaction where one reactant is converted into one product, which then decays:

$$A \longrightarrow B \longrightarrow \tag{48}$$

From this reaction schematic, we can determine a simple differential equation model of the concentrations of both A and B. We assume that A is being converted into B at a linear rate (as this is a first-order reaction), and that it is not being produced by any mechanism. This leads to the following dynamics, which we have seen a few times in this course:

$$\frac{dA}{dt} = -k_1 A(t) \tag{49}$$

On the other hand, there are two rates associated with B, namely the rate at which it is produced and the rate at which it decays. The decay rate of B is easily modelled, using the same techniques as we've been through before. However, the production rate of B depends on how much A is available, rather than being a function of B. More specifically, it is the opposite of the corresponding term in $\frac{dA}{dt}$, due to the reaction stoichiometry (where one molecule of B is produced for every one molecule of A consumed). This leads to the following differential equation for B:

$$\frac{dB}{dt} = k_1 A(t) - k_2 B(t) \tag{50}$$

Note that integrating $\frac{dB}{dt}$ requires knowledge of $A(t)$, so we can't solve for B on its own. Instead, we need to consider both $\frac{dA}{dt}$ and $\frac{dB}{dt}$ together, in what is called a system of differential equations. You may see these written like this:

$$\begin{cases} \frac{dA}{dt} = -k_1 A \\ \frac{dB}{dt} = k_1 A - k_2 B \end{cases} \tag{51}$$

Note that this is written in a slightly different form than what you have been seeing in other parts of this course, with the time derivative isolated on the left-hand side rather than the forcing function (if any) isolated on the right-hand side. This is because systems of differential equations often come up in mathematical models, where we know the rate of change of some quantity and are interested in determining how that quantity evolves over time.

Systems of differential equations can arise in other ways besides the interactions of several different variables. Another way that you can get a system is by taking a single differential equation of higher than first order and defining extra variables to make it easier to solve (or simulate). You can do this with differential equations of any order, although in practice, third-order ODEs will usually be the highest that you'll see. One example where you might see a third derivative is the process of controlling an elevator, or any other object that starts at rest, accelerates to some velocity, then returns to rest upon reaching its destination. As the acceleration changes over time, its derivative must be nonzero. In fact, limiting this third derivative of position (referred to as jerk) is very important, since high values of jerk may cause passengers to feel uncomfortable.

For a straightforward example of turning a higher-order ODE into a system of first-order ones, suppose we have a third-order ODE with constant coefficients:

$$x''' + ax'' + bx' + cx = 0 \tag{52}$$

We haven't learned how to solve these using our previous methods. However, we can define the quantities $y = x'$ and $z = y' = x''$ in order to get three first-order differential equations, like this:

$$\begin{cases} x' = y \\ y' = z \\ z' = -az - by - cx \end{cases} \tag{53}$$

This is actually quite straightforward to solve, which you can do using linear algebra. Before we learn how to do this, I'll do a quick review of some linear algebra basics. First, there's the concept of linear independence, which you have seen earlier in this course with regards to finding solutions to second-order ODEs. In the context of linear algebra, linear independence is similar. If we have vectors $v_1, \ldots, v_n$, then these vectors are linearly dependent if at least one of them can be written as a scalar multiple of the others, and linearly independent if this cannot be done. In other words, $v_1, \ldots, v_n$ are linearly independent if the following holds:

$$a_1 v_1 + \ldots + a_n v_n = 0 \implies a_i = 0 \; \forall i \tag{54}$$

For some easy examples, $v_1 = [1\ 0]^T$ and $v_2 = [1\ 1]^T$ are linearly independent, whereas $w_1 = [1\ 0\ 1]^T$, $w_2 = [1\ 1\ 0]^T$ and $w_3 = [0\ -1\ 1]^T$ are linearly dependent because w_3 can be expressed as $w_1 - w_2$.

If you have a set of vectors that make up the rows or columns of a matrix, then one way to determine if they are linearly independent is to take the determinant of the matrix. The determinant of a matrix is defined recursively. For a 1x1 matrix, it's just the single value in the matrix. For a matrix with higher dimensions, calculating the determinant is done by going along one particular row (or down one particular column) and taking the sum of the entries in that row or column multiplied by what are called the cofactors in that row or column. For each entry a_{ij} in a matrix A, the cofactor C_{ij} associated with that entry is defined in the following way, where M_{ij} is the matrix formed by deleting row i and column j from A (also known as a minor):

$$C_{ij} = (-1)^{i+j} \det M_{ij} \tag{55}$$

So, the determinant is just $a_{1,1}C_{1,1} + a_{1,2}C_{1,2} + \ldots + a_{1,n}C_{1,n}$ for an $n \times n$ matrix, or equivalently $a_{2,1}C_{2,1} + a_{2,2}C_{2,2} + \ldots + a_{2,n}C_{2,n}$, or $a_{1,1}C_{1,1} + a_{2,1}C_{2,1} + \ldots + a_{n,1}C_{n,1}$ if expanding along the first column, et cetera. One trick to doing this is to expand along a row or column that has 0 for some of its entries, as that makes the computation easier. If this determinant is nonzero, then the row vectors or column vectors making up the matrix are linearly independent.

Another important concept is invertibility of a matrix. If a matrix A is invertible, then there exists some other matrix A^{-1} for which A times A^{-1} is equal to the identity matrix I. (Remember that the identity matrix I has the property that $AI = A$ for any other matrix A, and is defined by having ones on its diagonal and zeros everywhere else.) If the determinant of a matrix is nonzero, then it is invertible. If the determinant is zero, then its inverse does not exist, and it is referred to as singular.

In general, multiplying a matrix by a vector results in another vector. In some cases, this operation may be equivalent to multiplying the original vector by a constant (in other words, scaling each element of the vector by the same amount). For a given matrix A, the vectors x that can be operated on in this manner satisfy the following property, for some constant λ:

$$Ax = \lambda x \tag{56}$$

In this case, x is called an eigenvector of A, and λ is its associated eigenvalue. We can find these by manipulating the above equation to get the following:

$$(A - \lambda I)x = 0 \tag{57}$$

Note that $(A - \lambda I)$ is a matrix, and that it needs to have a determinant of zero in order for $(A - \lambda I)x = 0$ to have nonzero solutions for x. This is because the kernel of $(A - \lambda I)$ (i.e. the set of vectors that $(A - \lambda I)$ maps to the

zero vector) needs to be non-trivial, which cannot happen if all of the rows of $(A - \lambda I)$ are linearly independent as a consequence of the rank-nullity theorem. Therefore, we need to find λ that satisfies the following equation:

$$\det(A - \lambda I) = 0 \tag{58}$$

In order to do this, remember that λI is just the identity matrix multiplied by the scalar λ, so $(A - \lambda I)$ will necessarily have diagonal entries of the form $a_{ii} - \lambda$. Since λ is unknown, this means that the determinant of $(A - \lambda I)$ will be a polynomial in λ rather than a single constant. This polynomial is called the characteristic equation of A, and solving it for λ produces the eigenvalues of A. Once the eigenvalues have been found, finding the corresponding eigenvectors can be done by substituting each eigenvalue back into the matrix equation $(A - \lambda I)x = 0$ and solving the resulting equation for x. (Since finding the determinant is easier if you expand along a row or column that is mostly zeros, you might expect there to be some trick involving doing that when finding eigenvalues, and indeed there is. The eigenvalues of a triangular matrix, one where all of the entries above or below the main diagonal are zero, are equal to the diagonal entries of that matrix.)

Here's an example of how to find eigenvalues and eigenvectors. Suppose that we have the following matrix:

$$A = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix} \tag{59}$$

Subtracting λI from this matrix, i.e. subtracting λ from each of the diagonal entries, gives us the following:

$$A = \begin{bmatrix} 1 - \lambda & 2 \\ 2 & 1 - \lambda \end{bmatrix} \tag{60}$$

We will now expand along the first row to obtain the determinant of this matrix. This produces a quadratic polynomial for λ, as this is a 2×2 matrix:

$$\det A - \lambda I = (1 - \lambda)(1 - \lambda) - 4 = \lambda^2 - 2\lambda - 3 = (\lambda - 3)(\lambda + 1) \tag{61}$$

This factors rather nicely, and gives us the eigenvalues of $\lambda = 3$ and $\lambda = -1$. Now, it remains to find their corresponding eigenvectors. Consider the equation $(A - \lambda I)x = 0$ for the case where $\lambda = 3$:

$$\begin{bmatrix} -2 & 2 \\ 2 & -2 \end{bmatrix} x = 0 \tag{62}$$

If we assume that x is a vector with two entries, or $x = [x_1 \ x_2]^T$, then both the first and second rows of the matrix in this matrix-vector equation imply that $x_1 = x_2$. We can thus assume that the eigenvector associated with $\lambda = 3$ is any vector satisfying these requirements. (For instance, we can go with $[1 \ 1]^T$.) Considering the case where $\lambda = -1$ gives us a different equation to solve:

$$\begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix} x = 0 \tag{63}$$

Here, we get the condition that $x_1 = -x_2$. Therefore, we can choose $[-1\ 1]^T$ for the eigenvector associated with $\lambda = -1$, although once again any multiple of this is also an eigenvector. Note that in this case, the eigenvalues were distinct. However, it is possible to have repeated roots of the characteristic polynomial, and therefore repeated eigenvalues. One very simple example of this is the identity matrix I. When we subtract λ from the diagonal and take the determinant, we get the following:

$$\begin{vmatrix} 1-\lambda & 0 \\ 0 & 1-\lambda \end{vmatrix} = (1-\lambda)^2 \implies \lambda = 1 \tag{64}$$

We call $\lambda = 1$ an eigenvalue of multiplicity 2. Note however that I still has two distinct eigenvectors, which we can once again get by solving $(A-\lambda I)x = 0$ for $\lambda = 1$. When doing this, we get the equation $0x = 0$ and can therefore choose any vectors for x. The canonical ones are $[1\ 0]^T$ and $[0\ 1]^T$. These together span all of $\mathbb{R}^2$, which is intuitive as I has the property that $Ix = x$ for any vector x (and hence all vectors in $\mathbb{R}^2$ are eigenvectors of I).