Easy differential equations

for scientists, engineers, investors, musicians, etc.

As a physics TA, I've spent a lot of time trying to explain exponentials and oscillators without calculus. Students have to memorize many magic-looking equations, and each one only works in a few specific situations. I think it would be easier to just teach everyone a tiny bit of differential equations.

This idea might make you nervous if you haven't taken calculus – or if you've taken calculus but aren't sure you understood it. That's OK! For these examples, you don't need to actually *do* calculus. A vague intuition about "rate of change" is good enough for now. You need calculus to *verify*, *prove*, or *generalize* these results. You don't need calculus just to see how they are useful.

Velocity and acceleration

If you've taken a physics class, you've probably seen **Newton's Second Law** many times:

$$F = ma$$

This equation says the force on a thing equals the thing's mass times its acceleration.[1] Newton's Second Law makes a lot more sense if you have a working intuition about what *acceleration* means.

If x is some number which changes as time passes, then I'll use $\dot{x}$ to mean "the rate at which x is changing." This $\dot{x}$ is the *first time derivative* of x, also called the **velocity** of x.[2]

$$\text{velocity} = \dot{x} = \frac{dx}{dt}$$

Now let the symbol $\ddot{x}(t)$ denote the rate at which $\dot{x}(t)$ is changing. This $\ddot{x}$ is the *second time derivative* of x, also called the **acceleration** of x. It's "the rate of change of the rate of change of x."

$$\text{acceleration} = \ddot{x} = \frac{d^2x}{dt^2}$$

We could also define higher derivatives like $\dddot{x}$, but I won't do that here. Before we continue, I should warn people of a potential confusion: I'm assuming x is a **function** of time. That means x is a recipe that inputs a number t and outputs a number $x(t)$. Here is an example of a formula for $x(t)$:

$$x(t) = 4t + 5$$

Scientists and engineers will often say "the function $x(t)$ is $4t + 5$." Technically, this is bad language. We should say "a formula for the function x is $x(t) = 4t + 5$." The important things to remember are: x is a *function*, and $x(t)$ is the *number* output by the function x if you give it the number t as an input.

The picture below shows how I think about velocity and acceleration: I imagine I'm driving a car. At time t, $x(t)$ is how far I've gone, $\dot{x}(t)$ is how fast I'm going, and $\ddot{x}(t)$ is how hard I'm accelerating. If I use full throttle, then $\ddot{x}(t)$ will be large and positive. If I slam on the brakes, then $\ddot{x}(t)$ will be large and negative. (I don't use the steering wheel. If I did, I'd need to use at least two coordinates x, y to describe my position. For simplicity, this paper only describes one-dimensional differential equations.)

[1] If the thing can move in 2 or 3 dimensions, then F and a are *vectors*, not numbers. But vectors are not the issue here.

[2] $\dot{x}$, pronounced "x dot," is Newton's old notation. He called a *fluxion*.

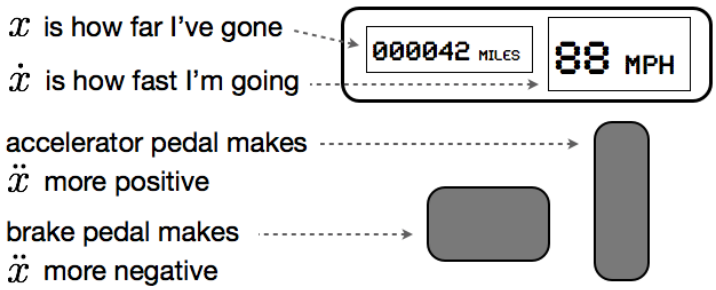

This analogy is simple, but it's not perfect. Real speedometers measure *speed*, which is the absolute value of velocity. Real odometers usually don't show negative numbers when the car moves in reverse. And in a real car, friction and air resistance cause the car's speed to decrease unless the engine fights back. Real drivers need to press the accelerator pedal just to keep the car at a constant velocity.

It's important to be careful about $+$ and $-$ signs. Negative $x(t)$ means I've gone backwards from wherever I was at time $t = 0$. Negative $\dot{x}(t)$ means I'm moving backwards right now. Negative $\ddot{x}(t)$ means I'm going forwards and losing speed OR I'm going backwards and gaining speed.

Differential equations

For our purposes, a **differential equation** is an equation with derivatives in it. These are often used in science and engineering to predict how some number changes in time. Here are three examples:

$$\begin{array}{rl} \text{Exponential growth} & \dot{x}(t) = 5x(t) \\ \text{Exponential decay} & \dot{x}(t) = -5x(t) \\ \text{Simple harmonic oscillation} & \ddot{x}(t) = -5x(t) \end{array}$$

These are all *linear ordinary differential equations* for one variable. They show up very often, they're relatively easy to solve, and they are often used as the building blocks of more complicated theories.

I should mention that I use the word "solve" in a way that might seem like cheating. When scientists and engineers say they've "solved" a differential equation, we usually mean something like this:

1. I already knew the solution.
2. I looked up a general solution and adjusted it to fit a specific problem.
3. I used a computer program to approximate the solution.

When I say a differential equation is "easy" to solve, I really mean the general solution is in a textbook. Choosing a unique solution for a specific problem requires some **initial conditions**: the value of x, and maybe some of its derivatives, must be known at a specific time t_0. In principle, *existence/uniqueness* theorems say when a differential equation has no solutions, one solution, or multiple solutions. In practice, scientists and engineers often assume a unique solution exists as long as we choose realistic initial conditions. If it doesn't, we call the equation *pathological* and warn people to stay away from it.

Exponential growth

Imagine a biologist has a dish with some stuff in it that bacteria can eat. Let $x(t)$ be the number of bacteria in the dish at time t. Assume the growth rate $\dot{x}(t)$ of the bacteria is proportional to the number of bacteria in the dish. In other words, $\dot{x}(t)$ is some constant multiplied by $x(t)$:

$$\dot{x}(t) = rx(t)$$

Here r is a positive real number which depends on the bacteria, food, and other details of the experiment. I looked up the general solution, and it is the **exponential growth** formula:

$$x(t) = Ae^{rt}$$

A is a constant which depends on the initial conditions of the system. Suppose that at time $t = 0$, our biologist puts a certain number of bacteria in the dish. Call that number $x(0)$. Then:

$$x(0) = Ae^0 = A \quad \Rightarrow \quad x(t) = x(0)e^{rt}$$

Apparently $A = x(0) =$ the number of bacteria at time $t = 0$. That number is the initial condition needed to choose a unique solution. Figure 1 plots two examples of exponential growth.

Figure 1: Exponential growth. Left: $\exp(t)$. Right: $\exp(3t)$.

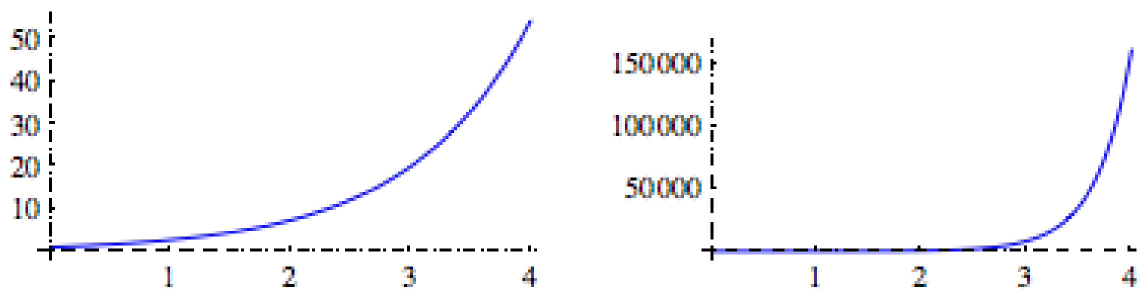

The famous number e is a weird *transcendental number* which can't be written down exactly as a fraction or decimal. (The number π is another famous transcendental.) The value of e is approximately:

$$e \approx 2.7182818284$$

Many people find "e to the rt power" confusing. I prefer to think of the **exponential map** $\exp(rt)$:

$$\exp(rt) = \sum_{n=1}^{\infty} \frac{(rt)^n}{n!} = 1 + rt + \frac{(rt)^2}{2} + \frac{(rt)^3}{3!} + \frac{(rt)^4}{4!} + \frac{(rt)^5}{5!} \cdots$$

The exclamation point is a **factorial**. It's defined like this:

$$0! = 1 \qquad 1! = 1 \qquad 2! = 2*1 \qquad 3! = 3*2*1 \qquad 4! = 4*3*2*1 \qquad \text{etc.}$$

In practice, we don't need to use all ∞ terms to get an accurate answer. Eventually the terms get very small and can be ignored.[3] I think of the exponential map $\exp(rt)$ as the *definition* of e^{rt}, and I define e as "the number you get when you calculate $\exp(1)$ very precisely."

If you remember the Chain Rule from calculus, you can prove that the time derivative of this series is $r\exp(rt)$. If not, then just think of $\exp(rt)$ as a magic formula for exponential growth.

Exponential growth is not a perfect model of bacteria in a dish. An obvious problem is that $x(t)$ is a real number, but the number of bacteria should really be an integer. Also if t gets large, then old bacteria might die, or maybe there won't be enough food for the population to keep growing exponentially.[4] The point is: many systems approximately grow exponentially, at least for a short time. Examples include disease propagation, chemical chain reactions, and the spread of internet memes.

[3]You can use a computer to check this if $|rt|$ is not too big. If $|rt| \gg 1$, then roundoff errors cause problems.

[4]The *logistic equation* is often a better model for population growth on longer time scales.

Exponential growth is extremely important to anyone studying finance or economics. In the 17th century, Jacob Bernoulli showed that the value of a unpaid loan with compound interest grows exponentially if the interest is compounded very often. If someone is making many small payments on the loan, the balance can be modeled by subtracting a constant from $\dot{x}$. And if r flickers randomly in time, then the resulting *stochastic differential equation* provides a model of a stock, bond, or other volatile investment.

Exponential decay

The **exponential decay** equation is a "backwards" version of exponential growth:

$$\dot{x}(t) = -rx(t)$$

This equation describes anything whose *decay rate* is proportional to how big it is. Examples include radioactive decay, voltage on a discharging capacitor, and the value of cash when inflation is significant. The solution looks just like exponential growth with a negative growth rate:

$$x(t) = x(0)e^{-rt}$$

Figure 2 plots two examples of exponential decay.

Figure 2: Exponential decay. Left: $\exp(-t)$. Right: $\exp(-3t)$.

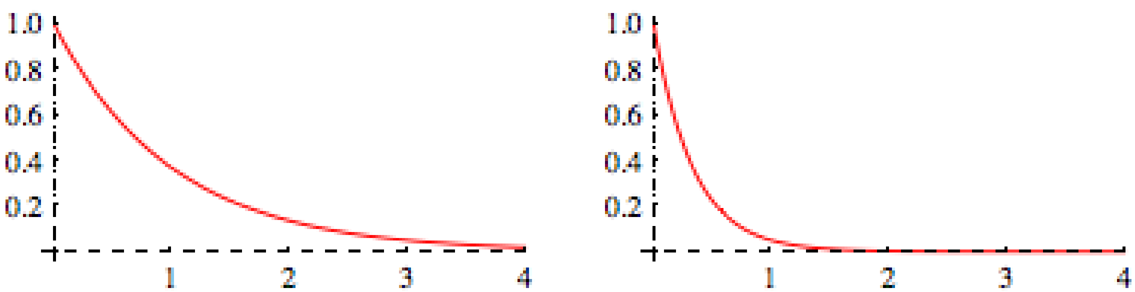

Some researchers prefer to use **half-life** t_h instead of the constant r. The half-life of an exponentially-decaying quantity x is the amount of time it takes to shrink to half of its original value. With a calculator, you can figure out r if you already know t_h, or vice versa:

$$x(t_h) = \tfrac{1}{2}x(0) \quad \Leftrightarrow \quad e^{-rt_h} = \tfrac{1}{2} \quad \Leftrightarrow \quad -rt_h = \ln\left(\tfrac{1}{2}\right) \quad \Leftrightarrow \quad rt_h = \ln(2)$$

The term "half-life" often makes people think of nuclear waste, but it can also be extremely important for medical doctors. The *biological half-life* of a drug is the time needed for its concentration in blood plasma to drop to half its peak value. For example, morphine has $t_{1/2} \approx 2.5$ hours. A patient given 10mg morphine at noon will typically have about 5mg left in his blood at 2:30pm. (Exponential decay is a good model for a *first-order elimination process*, but not all drugs work this way.)

Simple harmonic oscillators

Assume some box of stuff is connected to a spring. For simplicity, assume the box can only move in one direction, and there are no other forces like gravity or friction. Define x to be the position of the box. Choose coordinates such that $x = 0$ when the spring is at its equilibrium length. Here is a picture:

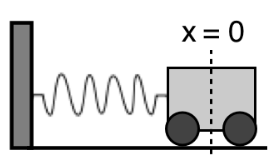

I drew wheels to mean "neglect friction." Suppose the force exerted by the spring on the box is:

$$F = -kx$$

This is *Hooke's Law*, though it should really be called *Hooke's linear approximation*. This equation is just a reasonable model; real springs are more complicated. In words, it says "the force the spring exerts on the cart is some constant k times how far the spring has been pushed (or pulled) from its equilibrium length." Newton's Second Law then leads us to the *spring equation*:

$$F = ma = m\ddot{x}(t) \qquad \text{and} \qquad F = -kx(t) \qquad \Rightarrow \qquad m\ddot{x}(t) = -kx(t)$$

This is a *second-order differential equation*: there's a second derivative in it. I looked up the general solution in a textbook, and I call it the **simple harmonic oscillator** formula:

$$x(t) = A\cos(\omega t) + B\sin(\omega t) \qquad \omega = \sqrt{\frac{k}{m}}$$

The symbol ω is an **angular frequency** in radians per second. It's conventional to define ω this way so we don't have to write a bunch of square-root signs. Here sin and cos are the *sine* and *cosine* functions from trigonometry. Figure 3 shows plots of these functions. If you teach yourself to draw Figure 3 from memory, it can save huge amounts of time and effort later.

Figure 3: Plots of $\sin(t)$ and $\cos(t)$ with $0 \le t \le 4\pi$.

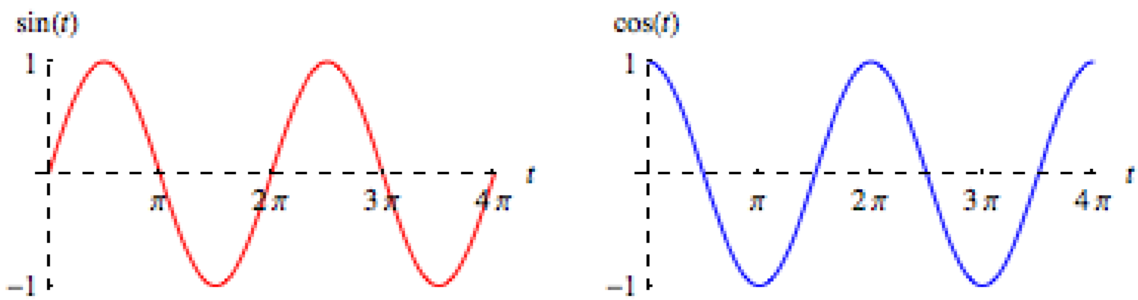

The numbers A and B are constants which depend on the initial conditions. The exponential growth solution had one arbitrary constant which we had to adjust to fit a specific problem. The simple harmonic oscillator solution has *two* arbitrary constants. (This is a general property of second-order differential equations.) It turns out that the constant A is just the initial value $x(0)$:

$$x(0) = A\cos(0) + B\sin(0) = A$$

To find a formula for $\dot{x}(t)$, I looked up the rules for taking derivatives of trig functions:

$$\dot{x}(t) = -\omega A\sin(\omega t) + \omega B\cos(\omega t)$$

This formula tells us that ωB is the initial velocity $\dot{x}(0)$:

$$\dot{x}(0) = -\omega A\sin(0) + \omega B\cos(0) = \omega B$$

So if we know the initial position $x(0)$ and initial velocity $\dot{x}(0)$ of the cart, then we can adjust A and B to fit our specific problem: just set $A = x(0)$ and $\omega B = \dot{x}(0)$. Figure 4 shows two examples.

If you have some practice taking derivatives of trig functions, you can check that

$$\ddot{x}(t) = -\omega^2 A\cos(\omega t) - \omega^2 B\sin(\omega t) = -\omega^2 x(t)$$

Figure 4: Example solutions. Left: $\cos(t) + \sin(t)$. Right: $\cos(5t) - \sin(5t)$.

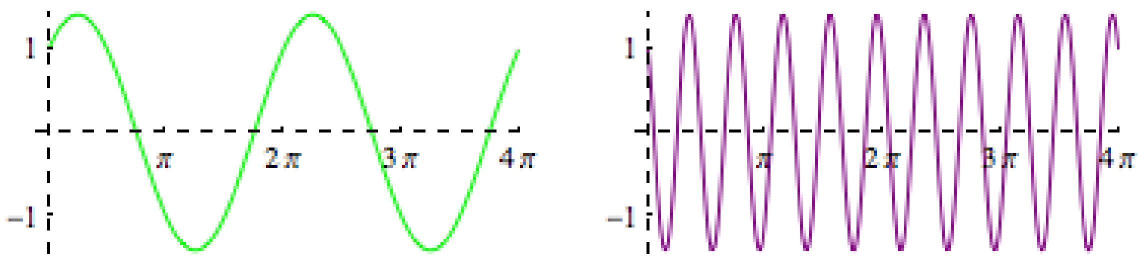

If a differential equation can be written in the form $\ddot{x}(t) = -\omega^2 x(t)$ for some positive real number ω, then the solution is $x(t) = A\cos(\omega t) + B\sin(\omega t)$ with $A = x(0)$ and $\omega B = \dot{x}(0)$.

Figures 3 and 4 show what simple harmonic motion looks like: the value of $x(t)$ wiggles back and forth forever. The wiggles have a smooth **sinusoidal** shape, which just means they look like a sin or cos function which may have been stretched or squished horizontally and/or vertically.

The **frequency** of an oscillation is the number of times $x(t)$ repeats itself per second. Figure 3 shows that $\sin(t)$ repeats itself once every 2π seconds, and so does $\cos(t)$. So the frequency of these oscillations is $f = 1/(2\pi)$. The **period** of an oscillation is the amount of time needed for the oscillation to repeat itself exactly once. By definition, frequency and period are each other's inverse: $T = 1/f$.

It's unfortunately easy to confuse *frequency* with *angular frequency*. I think of them as the same quantity written in different units: ω is in radians per second, and f is in cycles per second (also called Hz). To convert between units, I use the rule 1 cycle $= 2\pi$ radians.

$$\omega \frac{\text{radians}}{\text{second}} = f \frac{\cancel{\text{cycles}}}{\text{second}} * 2\pi \frac{\text{radians}}{\cancel{\text{cycle}}}$$

Figure 4 shows that $\cos(5t) - \sin(5t)$ repeats itself 5 times every 2π seconds. Its frequency is $f = 5/(2\pi)$ Hz, its period is $T = 1/f = (2\pi)/5$ seconds, and its angular frequency is $\omega = 2\pi f = 5$ radians/second.

The **amplitude** of an oscillation is a measure of how big the wiggles are. For a sinusoidal oscillation, amplitude is the difference between the maximum value and average value of $x(t)$. For example, the maximum value of $5\cos(\omega t)$ is 5, and the average value is 0, so the amplitude is 5. (This is true no matter what ω is.) The amplitude of $5\sin(\omega t)$ is also 5. The amplitude of $A\cos(\omega t) + B\sin(\omega t)$ is:

$$\sqrt{A^2 + B^2}$$

though you might need to do some trigonometry if you want to prove it.

Sinusoidal oscillations are called **simple** because more complicated oscillations can be described by adding many different sinusoidal functions together. The frequencies in a complicated oscillation are called its **spectrum**, and writing the oscillation as a combination of simple oscillators is called *Fourier analysis* or *spectral analysis*. This subject is very important for electrical engineering, mechanical engineering, quantum mechanics, optics, and the study of just about anything which oscillates.

Musicians often use an intuitive version of spectral decomposition. A sound is just an oscillating air pressure which is noticed by a listener. Sinusoidal sounds are called **pure tones**. (A tuning fork produces a nearly-pure tone.) More complicated sounds can be constructed by adding together many pure tones with different frequencies and amplitudes. I never understood spectral analysis until I realized it was just a mathematical formalism for something my ears had been doing automatically.

There are other ways to write the simple harmonic oscillator formula. Here are two popular choices:

$$x(t) = A\cos(\omega t + \phi) \qquad\qquad x(t) = B\sin(\omega t + \theta)$$

Each of these formulas can be adjusted to fit our initial conditions by choosing ϕ or θ as the *initial phase* of the system. It's also possible to write the formula using complex exponentials. But these are more advanced topics, and $x(t) = A\sin(\omega t) + B\cos(\omega t)$ is a perfectly good general solution.

Not-as-easy differential equations

This part gets complicated, so feel free to skip it if you don't need it.

So far, I've only mentioned exponential growth/decay and simple harmonic oscillators. If you have to solve a differential equation which doesn't fit into these categories, the previous formulas won't help. Fortunately, mathematicians have spent centuries finding general solutions. With some good books and websites, we can often look up the general solution and adjust it to fit our initial conditions.

For example, the **damped harmonic oscillator** equation is used to model automobile suspensions, electrical filters, and vibrations of percussion instruments. It looks like this:

$$\ddot{x}(t) = -2\beta\dot{x}(t) - \omega^2 x(t)$$

It's like the simple harmonic oscillator equation, except with a *damping force* which is proportional to velocity. Depending on how large β is relative to ω, the solutions look different:

Underdamped: $\beta < \omega$	$x(t) = e^{-\beta t}\Big[A\cos(\Omega t) + B\sin(\Omega t)\Big]$	$\Omega = \sqrt{\lvert\beta^2 - \omega^2\rvert}$
Critically damped: $\beta = 2\omega$	$x(t) = e^{-\beta t}\Big[A + Bt\Big]$	
Overdamped: $\beta > 4\omega$	$x(t) = e^{-\beta t}\Big[Ae^{\Omega t} + Be^{-\Omega t}\Big]$	$\Omega = \sqrt{\beta^2 - \omega^2}$

There is a slick way to find these solutions using complex exponentials – but for our purposes, it's good enough to look them up in a textbook, figure out which formula to use, and use the initial conditions to adjust A and B. (To do that, you might need just a little bit of calculus: take the derivative of the general formula for $x(t)$, plug in 0 for t, and set that equal to the initial velocity $\dot{x}(0)$.)

Figure 5 shows some examples of damped harmonic oscillators with initial conditions $x(0) = 0$ and $\dot{x}(0) = 1$. Notice that the underdamped solution oscillates, but the others do not. The critically-damped solution "gets back to zero" faster than the others. This would be important if we were building, for example, a tuned-mass damper designed to keep a skyscraper from shaking on windy days.

Figure 5: Example damped harmonic oscillators with $\omega = 1$ and initial conditions $x(0) = 0$, $\dot{x}(0) = 1$. Red: underdamped with $\beta = 0.25$. Green: critically damped $\beta = 1$. Blue: overdamped with $\beta = 4$.

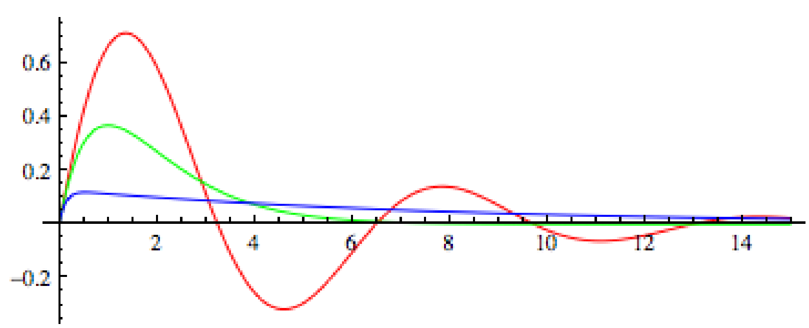

Before trying to solve any differential equation, it's usually a good idea to check a textbook for general solutions. Some complicated-looking equations show up so often that they have proper names:

Airy's equation:	$\ddot{x}(t) - tx(t) = 0$
Bessel's equation of order α :	$t^2\ddot{x}(t) + t\dot{x}(t) + (t^2 - \alpha^2)x(t) = 0$

The solutions to Airy's and Bessel's equations are called *Airy functions* and *Bessel functions*. Numerical software packages often have special algorithms for approximating **special functions** like these.

Some complicated differential equations are really just easy differential equations in disguise. New coordinates can sometimes transform an equation into something with a known solution. For example:

$$t^2\ddot{x}(t) + t\dot{x}(t) + 4x(t) = 0$$

Define "logarithmic time" $s(t) = \ln(t)$ and look for a function y such that $y(s(t)) = x(t)$. Use the Chain Rule to write $\dot{x}(t)$ and $\ddot{x}(t)$ in terms of $y(s(t))$ and its t derivatives:

$$\dot{x}(t) = \frac{1}{t}\dot{y}\big(s(t)\big) \qquad \ddot{x}(t) = \frac{1}{t^2}\Big[\ddot{y}\big(s(t)\big) - \dot{y}\big(s(t)\big)\Big]$$

Here $\dot{y}$ and $\ddot{y}$ are the first and second s derivatives of the unknown function y. Substituting these back into the original equation gives us a simple harmonic oscillator equation:

$$\ddot{y}(s) = -4y(s)$$

If you want to know more about this tactic, look up **Cauchy-Euler equations**.

Hard differential equations

The examples here are only the tip of the differential-equation iceberg. If multiple numbers $x(t), y(t), z(t)$ are changing in time, then we have to solve a *multivariable* equation. If the unknown functions have multiple inputs, we might have to solve a *partial differential equation*. *Nonlinear* differential equations can sometimes behave *chaotically*, which (roughly) means that two very similar initial conditions might lead to wildly different solutions. Many of these hard differential equations are impractical or impossible to solve with paper and pencil, so computer **numerical integrators** are used to approximate them.

Still, these few examples should be a good start. Hopefully they give people some idea of how calculus can be useful to nurses, virologists, motorcyclists, piano makers, bond traders, and so on.

About the author

Sam Kennerly is a physics PhD candidate, numerical programmer, and electronic musician.
