

Math 221 Section 5
Topics for the first exam

Background

A differential equation is an equation involving an (unknown) function y and some of its derivatives. The basic goal is to *solve* the equation, i.e., to determine which function or functions satisfy the equation. Differential equations come in several types, and our techniques for solving them will differ depending on the type.

Ordinary vs. partial: If y is a function of only one variable t , then our differential equation will involve only derivatives w.r.t. t , and we will call the equation an ordinary differential equation. If y is a function of more than one variable, then our differential equation will involve partial derivatives, and we will call it a *partial differential equation*. We will deal almost exclusively with ordinary differential equations in this class.

Systems: Sometimes the rates of change of several functions are inter-related, as with the populations of a predator $y(t)$ and its prey $x(t)$, where $x' = ax - \alpha xy$ and $y' = \gamma xy - cy$. We call this a *system* of differential equations, and its solution would involve finding both $x(t)$ and $y(t)$.

Order: Techniques for solving differential equations differ depending upon how many derivatives of our unknown function are involved. The *order* of a differential equation is the order of the highest derivative appearing in the equation. The Implicit Function Theorem tells us that we can rewrite our equation so that it equates the highest order derivative with an expression involving lower order terms:

$$y^{(n)} = F(t, y, y', \dots, y^{(n-1)})$$

Linear vs. non-linear: A differential equation is *linear* if it can be written as

$$a_0(t)y^{(n)} + \dots + a_{n-1}(t)y' + a_n(t)y = g(t)$$

(i.e., the function F is linear in the variables $y, y', \dots, y^{(n-1)}$, although it **need not** be linear in t). A differential equation is non-linear if it isn't linear! E.g.,

$$y' = y^2$$

is non-linear, while

$$y' = (\sin t)y/(1 + t^2) - \cos(\cos t)$$

is linear.

Solutions and Initial Value Problems

Solving a differential equation means to determine which function or functions satisfy the equation. Our solutions come in two flavors: *explicit* solutions $y = y(t)$ which provide a function of t which satisfies the equation, and *implicit* solutions which provide an equation $g(y, t) = 0$ which any explicit solution would have to satisfy. The idea is that we can treat $g(y, t) = 0$ as implicitly defining y as a function of t ; given a specific value $t = c$ for t , we solve (numerically?) $g(y, c) = 0$ for y to determine the value of the solution to the differential equation at c .

In general, a differential equation $y' = f(t, y)$ will have many solutions; but typically one particular solution can be specified by requiring one additional condition be met; that y take a specific value y_0 at a specific point t_0 . If we think of the time t_0 as the time at which we “start” our solution, then we call the pair of equations

$$y' = f(t, y) \quad y(t_0) = y_0$$

an *initial value problem* (or *IVP*). There is a general result which gives conditions guaranteeing that an IVP has a solution:

If $y' = f(t, y)$ is a differential equation with both f and $\frac{\partial f}{\partial y}$ continuous for $a < t < b$ and $\alpha < y < \beta$, and $t_0 \in (a, b)$ and $y_0 \in (\alpha, \beta)$, then for some $h > 0$, the initial value problem

$$y' = f(t, y) \quad , \quad y(t_0) = y_0$$

has a unique solution for $t \in (t_0 - h, t_0 + h)$.

If only f is continuous, then we can guarantee existence, but not necessarily uniqueness. In general, however, the size of the interval where we can guarantee existence (and uniqueness) can be very small, and often depends on the choice of initial value! For example, for the equation

$$y' = y^2$$

the righthand side is continuous everywhere (as is the partial derivative), but the interval we can choose for the solutions $y = -1/(t+c)$ depends on c , which will depend on the initial condition! And it can *never* be chosen to be the entire real line.

Failure to satisfy the hypotheses of the result can easily kill both existence and uniqueness. For example, the equation

$$y' = y^{1/3}$$

has many solutions with the initial condition $y(0) = 0$, such as $y = 0$ and $y = (2t/3)^{3/2}$.

Direction Fields

In many cases, especially for first order differential equations, we can ‘see’ what a solution should look like without actually finding the solution. For first order equations, $y' = f(t, y)$, a solution $y(t)$ will satisfy $y'(t) = f(t, y(t))$, and so we can think of $f(t, y)$ as giving the *slope* of the tangent line to the graph of $y(t)$ at the point $(t, y(t))$. But since the function f is already known, we can draw small line segments at ‘every’ point of the t - y plane with slope $f(t, y)$ at the point (t, y) ; this is called the *direction field* for our differential equation. A solution to our differential equation is simply a function whose graph is tangent to each of these line segments at every point along the graph. Thinking of the direction field as a velocity vector field (always pointing to the right), our solution is then the path of a particle being pushed along by the velocity vector field. From this point of view it is not hard to believe that *every* (first order ordinary) differential equation has a solution, in fact many solutions; you just drop a particle in and watch where it goes. *Where* you drop it is important (it changes where it goes), which really is what gives rise to the notion of an *initial value problem*; we seek to find the specific solution with the additional *initial value* $y(t_0) = y_0$.

Sketching the slope field for the differential equation $y' = f(t, y)$ can be aided by using *nullclines* (= the points where $f(t, y) = 0$; we draw horizontal lines there) and *isoclines* (= the points where $f(t, y) = c$ for a fixed constant c). See the handout from class for more specifics of this technique.

Separable Equations

There is a class of first order equations for which we can readily find solutions by integration; there are the *separable* equations. A differential equation is separable if it can be written as

$$y' = A(t)B(y)$$

This allows us to ‘separate the variables’ and integrate with respect to dy and dt to get a solution:

$$\frac{1}{B(y)} dy = A(t) dt ; \text{ integrate both sides}$$

In the end, our solutions look like $F(y) = G(t) + c$, so it defines y *implicitly* as a function of t , rather than explicitly. In some cases we can invert F to get an explicit solution, but often we cannot.

For example, the separable equation $y' = ty^2$, $y(1) = 2$ has solution

$$\int \frac{dy}{y^2} = \int t dt + c$$

so solving the integrals we get $(-1/y) = (t^2/2) + c$, or $y = -2/(t^2 + 2c)$; setting $y = 2$ when $t = 1$ gives $c = -1$.

Linear Equations

Perhaps the most straightforward sort of differential equation to solve is the *first order linear ordinary differential equation*

$$y' = a(t)y + b(t)$$

We will typically (following tradition) write such equations in *standard form* as

$$y' + p(t)y = g(t) \quad (**)$$

For example, near the earth and in the presence of air resistance, the velocity v of a falling object obeys the differential equation $v' = -g - kv$, where g and k are (positive) constants. There is a general technique for solving such equations, by trying to think of the left-hand side of the equation as the derivative of a single function. In form it looks like the derivative of a product, and by introducing an *integrating factor* $\mu(t)$, we can actually arrange this. Writing

$$(\mu(t)y)' = \mu(t)(y' + p(t)y) = \mu(t)g(t)$$

we find that (where $\exp(\text{blah})$ means e raised to the power 'blah')

$$\mu(t) = \exp(\int p(t) dt)$$

and so

$$\mu(t)y = \int \mu(t)g(t) dt = \int (\exp(\int p(t) dt)) g(t) dt + c$$

which we can then solve for y .

Putting this all together, we find that the solutions to (**) are given by

$$y = e^{-\int p(t) dt} (\int e^{\int p(t) dt} g(t) dt + c)$$

For example, the differential equation $ty' - y = t^2 + 1$, after being rewritten in standard form as $y' - (1/t)y = t + (1/t)$, has homogeneous solution

$$y_h = \exp(\int 1/t dt) = \exp(\ln t) = t$$

so we have

$$y = t(\int 1 + 1/t^2 dt) = t(t - (1/t) + c)$$

and so our solutions are $y = t^2 - 1 + ct$, where c is a constant.

But what is c ? Our solution is actually a *family* of solutions; a particular solution (i.e., a particular value for c) can be found from an initial value $y(t_0) = y_0$. For example, if we wished to solve the initial value problem

$$ty' - y = t^2 + 1, y(2) = 5$$

we can plug $t = 2$ and $y = 5$ into our general solution to obtain $c = 1$.

Application: Mixing Problems

In many instances, the rate of change of a quantity can be best analyzed by treating the factors that make the quantity go up separately from those that make it go down; each can often be readily understood in isolation. We can then build a differential equation modeling the behavior of the quantity $y = y(t)$ as

$$y' = (\text{things that make } y \text{ go up}) - (\text{things that make } y \text{ go down})$$

As a basic example, we have mixing problems. The basic setup has a solution of a known concentration mixing at a known rate with a solution in a vat, while the mixed solution is poured off at a known rate. The problem is to find the function which gives concentration in the vat at time t . It turns out that it is much easier to find a differential equation which describes the amount of solute (e.g., salt) in the solution (e.g., water), rather than the concentration.

If the concentration pouring in is A , at a rate of N , while the solution is pouring out at rate M with concentration $A(t) = x(t)/V(t)$, then if the initial volume is V_0 , we can compute $V(t) = V_0 + (N - M)t$. The change in the amount $x(t)$ of solute can be computed as (rate falling in) – (rate falling out), which is

$$x' = AN - A(t)M = AN - \frac{x}{V_0 + (N - M)t}M$$

This is a linear equation, and so we can solve it using our techniques above.

We can also deal with a succession of mixing problems, the output of one becoming the input of the next, by treating them one at a time; the only change in the setup above is that the incoming concentration for the next vat (to solve for $x_{i+1}(t)$) would be the concentration $x_i(t)/V_i(t)$ found by solving the equation for the previous vat.

Population Models

Another situation where this kind of analysis proves successful is in modeling population growth. The idea is that if y is the population at time t , then

$$y' = (\text{birth rate}) - (\text{death rate})$$

Typically, the birth rate is proportional to the population, i.e. is ry , while the death rate is either modeled as being proportional to the population (Malthusian model) or is a sum (logistic model); one part is proportional to the population (death by “natural causes”), the other is proportional to the square of the population (this typically represents contact between individuals, arising from competition for food, overcrowding, etc.), i.e., is ky^2 . Put together, and combining the two terms proportional to population, we obtain

$$\begin{aligned} y' &= ry && \text{for the Malthusian model, and} \\ y' &= ry - ky^2 && \text{for the logistic model} \end{aligned}$$

Both equations are separable, and so we can use phase diagrams to understand their long-term behavior (see below), as well as finding explicit solutions (using partial fractions, for the logistic equation).

Heating and Cooling

Newton's Law of Cooling: This states that the rate of change of the temperature $T(t)$ of an object is proportional to the difference between its temperature and the ambient temperature of the air around it. The constant of proportionality depends upon the particular object (and the medium, e.g., air or water) it is in. In other words,

$$T' = k(A - T)$$

Since a cold object will warm up, and a warm object will cool down, this means that the constant k should be positive. Writing the equation as

$$T' + kT = kA$$

we find the solution (after solving the IVP)

$$T(t) = A + (T(0) - A)e^{-kt}$$

Typically, k is not given, but can be determined by knowing the temperature at some other time t_1 , by plugging into the equation above and solving for k .

Autonomous Equations and the Phase Diagram

A differential equation is called *autonomous* if the function $f(t, y)$ is really a function $f(y)$ only of the variable y . We will learn how to solve such equations below; but we can learn a lot about the solutions to such an equation simply by understanding the graph of $f(y)$.

One feature of the solutions is that we can translate in time and get another solution; if $y(t)$ is a solution to $y' = f(y)$, then so is $z(t) = y(t + c)$ for any constant c , as can be verified by plugging z into the differential equation. This can also be verified geometrically, using the direction field approach. For an autonomous equation, the slope of the direction field is always the same along horizontal lines (since it depends on y , not t), and so if we pick up a solution curve, tangent to the direction field, and translate it in the horizontal direction, it will *still* be everywhere tangent to the direction field, and so is also a solution.

The key to understanding solutions to such equations $y' = f(y)$ is to find *equilibrium solutions*, that is, solutions $y = \text{constant} = c$. Such solutions have derivative 0, and so for such solutions we must have $f(c) = 0$. The basic idea is that these equilibrium solutions tell us a great deal about the behavior of *every* solution to the differential equation.

If the function $f(y)$ is continuous, then between the zeroes of f (i.e., the equilibrium solutions of the differential equation) f has all the same sign, and so for the solutions, y' has the same sign, so $y(t)$ is either always increasing or always decreasing. It cannot *cross* the equilibrium solutions, since this would violate the uniqueness of solutions to our differential equation. (Here we assume that the derivative of f is also continuous.) If a solution curve becomes asymptotic to a horizontal line, that line must *be* an equilibrium solution, because the tangent lines along our solution must be becoming horizontal, i.e., $f(y(t)) = f(y)$ is approaching $0 = f(\text{limit of } y(t))$.

Therefore, the structure of the solutions is very simple; between consecutive equilibrium solutions, the solutions increase or decrease monotonically from one equilibrium to the other. This allows us to classify equilibrium solutions as one of three kinds: *stable equilibria*, where nearby solutions all converge *back* to the equilibrium, *unstable equilibria*, where nearby solutions all diverge away from the equilibrium, and *semistable equilibria* or *nodes*, where on one side the solutions converge back, and on the other they diverge away.

The easiest way to assemble this data is to plot the roots of f on a number line (called the *phase diagram*, and then determine the sign of f in the intervals in between. Where it is positive, solutions move to the right (i.e., up), while where it is negative they move left. Marking these as arrows, a stable equilibrium has arrows on both sides pointing towards it, and an unstable equilibrium has both arrows pointing away.

Newtonian Mechanics: Velocity/Acceleration Models

If we wish to model the motion of an object, whose position at time t is given by $x(t)$, then (setting $v(t) = x'(t)$) Newton's Second Law of Motion tells us that

$$mv' = \text{the sum of the individual forces acting on the object}$$

When we can understand these forces, in terms of t and v , we can build a first order differential equation, which we can then bring our techniques to bear to solve. Typical forces include:

gravity: $F_g = mg$ or $F_g = -mg$, depending upon whether we think of the positive direction as down (giving $+$) or up (giving $-$). $g = 9.8 \text{ m/sec}^2 = 32 \text{ ft/sec}^2$ (approximately)

air resistance: this is typically modeled either as $F_a = -kv$ (for smallish velocities) or $F_a = -kv|v|$ (for large velocities). It always acts to push our velocity towards 0, hence the $-$ sign.

external force: $F_e = g(t)$; this represents a force that “follows along” the object and tries to push it in a direction that is “pre-programmed” in time.

With these sorts of forces, we get a general equation

$$mv' = \pm mg - kv + g(t)$$

which we can solve by the methods we have developed. For example, ignoring external forces and assume the positive direction is “down”, we have the initial value problem

$$mv' = mg - kv \quad v(0) = v_0$$

with solution

$$v(t) = \frac{mg}{k} + \left(v_0 - \frac{mg}{k}\right)e^{-\frac{kt}{m}}$$

As $t \rightarrow \infty$, $v(t) \rightarrow \frac{mg}{k}$ = the *terminal velocity*.

The Approximation Method of Euler

Most first order equations cannot be solved by the methods we will present here; the function $f(y, t)$ is too complicated. For such equations, the best we can often do is to *approximate* the solutions, using numerical techniques. One method is the *tangent line method*, also known as *Euler's method*. The idea is that our differential equation $y' = f(t, y)$ tells us the slope of the tangent line at every point of our solution, and the tangent line can be used to approximate the graph of a function, at least close to the point of tangency. In other words, for a solution to our differential equation,

$$y(t) \approx y(t_0) + y'(t_0)(t - t_0) = y_0 + f(t_0, y_0)(t - t_0)$$

for $t - t_0$ small. If we wish to approximate $y(t)$ for a value of t far away from our initial value t_0 , we use the above idea in several steps. We cut up the interval into n pieces of length h (called the *stepsize*), and then set

$$\begin{aligned} y_1 &= y_0 + f(t_0, y_0)h & , & & t_1 &= t_0 + h \\ y_2 &= y_1 + f(t_1, y_1)h & , & & t_2 &= t_1 + h \\ y_3 &= y_2 + f(t_2, y_2)h & , & & t_3 &= t_2 + h \end{aligned}$$

and continue until we reach y_n , which will be our approximation to $y(t) = y(t_n)$. Each step can be thought of as a mid-course correction, using information about the direction field at each stage to determine which way the solution is tending.

Calculus teaches us that at each stage the error introduced is approximately proportional to the square of h . So with a stepsize half as large, we will require twice as many steps, but each introduces an error only about one-fourth as large, so overall we get an error only half as large. This leads us to conclude that as the stepsize goes to 0, the error between our approximate solution y_n and $y(t_n)$ goes to 0.