This document should give a rough list of topics from the class, and some reminders to help you study.

The most important question is: given an ODE (or system of ODEs), what can you do with it? This usually involves identifying what kind of ODE it is, and knowing which techniques apply.

Approximate List of Topics:

- Recognizing ODEs:
  - ODEs vs. PDEs vs. systems of ODEs vs. other random stuff
  - Order? Autonomous/time-independent/constant coeffs? linear? homogenous? separable?
  - Turning a linear system into matrix-vector form.
- General properties of ODEs and systems:
  - What is an initial value problem (IVP)? Given a general solution, how to solve an IVP.
  - The existence and uniqueness theorem.
  - Where are solutions defined? Difference between the linear and non-linear case.
- What to do with 1ODEs:
  - If separable, separation of variables.
  - If linear, integrating factors.
  - Approximation using the euler method. (Also works for first order systems.)
  - Qualitative analysis using direction fields. Slope lines, isoclines, solution sketching, funnels and anti-funnels, separatrices, maxima and minima.
- General properties of linear ODEs and Linear operators:
  - Superposition. Homogeneous and inhomogeneous solution spaces.
- What to do with nLODEwCC:
  - Finding the characteristic polynomial, finding roots.
  - Finding the general homogeneous solution from the roots. Repeated roots. Complex roots, and basic manipulation of complex numbers.
  - Undetermined coefficients to solve inhomogeneous nLODEwCC with nice forcing functions.
  - The Exponential response formula and sinusoidal response formula, to speed this up.
- What to do with the spring equation (2LODEwCC and positive coeffs):
  - Transforming sinusoidal functions between cos/sin form and amplitude/phase form.

  - Overdamped vs. critically damped vs. underdamped vs. undamped. Behavior of homogeneous solutions.
  - Forced damped motion. Practical resonance.

- What to do with 1LSyswCC, 2 functions:
  - Finding the characteristic polynomial, computing the trace and determinant. (Generally knowing how to multiply matrices and vectors.)
  - Finding the eigenvalues and associated eigenvectors.
  - Getting the general homogeneous solution. (Except the repeated roots case.) Complex eigenvalues.
  - Phase portraits. Nodes, saddles, spirals. Stability.
  - Using trace and determinant to classify behavior.

Recognizing ODEs

- An ODE with order n is something of the form

$$x^{(n)}(t) = F(x^{(n-1)}, \ldots, x'', x', x, t).$$

  That is, the n-th derivative is a function of all the earlier derivatives. Don't mix up the order with other numbers which appear: $x' = x^n$ has order 1.

- A PDE involves a single function of multiple variables, like $x(t_1, t_2)$, and its partial derivatives. Don't mix it up.

- A system of ODEs involves multiple functions of a single variable, like $x_1(t)$ and $x_2(t)$. A solution involves choosing the functions x_i simultaneously.

- An ODE or system is *time-independent* if the formula for the derivative is independent of t. This is called different things in different contexts: *autonomous* for 1ODEs, *constant coefficients* for homogeneous nLODEs or linear systems.

- An ODE is *linear* if it has the form

$$x^{(n)}(t) = p_{n-1}(t)x^{(n-1)} + \ldots + p_2(t)x'' + p_1(t)x' + p_0(t)x + g(t)$$

  and $g(t)$ is called the *forcing function*. It is *homogeneous* if $g(t) = 0$. When $p_i(t)$ are constant functions, it is important whether or not $g(t)$ has a nice form, such as $t^2 e^{5t} \cos(7t)$.

- A 1ODE is *separable* if it has the form $x' = p(x)q(t)$. This is only a notion for first order ODEs.

General properties of ODEs

- An initial value problem for an nODE specifies $x, x', \ldots, x^{(n-1)}$ at a single time t_0. An initial value problem for a first order system with n functions specifies $x_1, x_2, \ldots, x_n$ at a single time t_0.

- If you have the general form of a solution in terms of some parameters (like $c_1, c_2, \ldots, c_n$) then solving an IVP involves computing these parameters.

- So long as the formula for the differential equation is continuous and differentiable in all "variables" at the initial condition, the initial value problem has a unique solution defined in some interval around t_0.

- It is not clear how big the interval of definition is. However, if the differential equation is linear, then the interval is "as big as possible," i.e. it extends until a time where the formula stops being continuous or differentiable.

- In non-linear cases, one can sometimes solve for the solution explicitly, and then figure out the interval of definition from the solution.

First order ODEs, or 1ODEs

- If separable, then integrate both sides of $\frac{dx}{p(x)} = q(t)dt$ to find a general solution. This is *separation of variables*.
- If linear, so $x' = p(t)x + g(t)$, then the general homogeneous solution is $x_h = ce^{\int p(t)}$. Choose any one homogeneous solution x_h. Then the inhomogeneous solution is $x = x_h \int \frac{g}{x_h} + cx_h$. This is called the method of *integrating factors*. The book has a method using $\frac{1}{x_h}$ which they call μ.
- Outside of these cases, we don't know how to solve any 1ODEs. There are some other special cases we can solve, but in general it is too hard.
- For any 1ODE, we can use Euler's method to approximate the value of a solution to an IVP. This involves repeatedly computing tangent lines. The same method can be used for systems, updating each function simultaneously.
- For any 1ODE $x' = F(x, t)$, we can draw a direction field to try to sketch the solution and figure out some qualitative aspects of the solution. This involves marking every point in the (x, t)-plane with a notch of slope $F(x, t)$ indicating what slope a solution through that point would have. Any solution must be tangent to every notch. No two solutions can cross.
- Drawing solutions for separable 1ODEs involves computing the equilibrium solutions, and where the derivative is positive and negative.
- An *isocline for slope* m is the set of all points in the (x, t) plane where a solution has slope m. Any local maximum or minimum must occur on the *nulcline*, which is the isocline for slope 0. To sketch the isocline, one must implicitly solve $F(x, t) = m$. This need not result in the graph of a function, and isoclines should not be confused with solutions.
- On a region where the isocline for slope m is the graph of some function $I(t)$, one can compare $I'(t)$ with m. This will determine whether a solution crosses the isocline from below or from above, or is tangent to it. Sometimes, one can find a pair of isoclines for which this phenomenon implies that any solution which gets between them cannot leave (as time continues). This is called a *funnel*. An *antifunnel* is a pair of isoclines for which solutions which are not between them can not get between them.
- A *separatrix* is a solution which separates other solutions which have drastically different behavior as time continues. Not every differential equation has a separatrix.

Linear operators

- An *operator* $\mathcal{L}$ is something which takes a function and returns another function. An operator is *linear* if for any functions f, g and any number c, one has $\mathcal{L}[cf] = c\mathcal{L}[f]$ and $\mathcal{L}[f + g] = \mathcal{L}[f] + \mathcal{L}[g]$.
- The derivative is a linear operator. Multiplication by a given function is a linear operator. Compositions of linear operators are linear. Any linear nODE has the form $\mathcal{L}[x] = g(t)$ for some linear operator $\mathcal{L}$.

- Solutions to a homogeneous nLODE form a vector space. I.e. if f_1 and f_2 are solutions, then so is $c_1 f_1 + c_2 f_2$ for any numbers c_1, c_2. This is part of *superposition*.
- More generally, if y_1 solves $\mathcal{L}[y] = g_1$ and y_2 solves $\mathcal{L}[y] = g_2$ then $y = c_1 y_1 + c_2 y_2$ solves $\mathcal{L}[y] = c_1 g_1 + c_2 g_2$. This is called *superposition*.
- Superposition implies that, if you find any one particular solution y_p to $\mathcal{L}[y] = g$, then every solution is of the form $y_p + y_h$ for some homogenous solution y_h.

nLODEwCC

- Basically, the only nODEs we know how to solve are nLODEwCC, with nice forcing functions.
- Any nLODEwCC has the form $\mathcal{L}[y] = p(D)y = g(t)$ for some polynomial p, where D is the derivative operator. p is called the *characteristic polynomial*.
- Because $p(D)e^{rt} = p(r)e^{rt}$, it is clear that whenever $p(r) = 0$, e^{rt} is a solution to the homogeneous equation.
- For repeated roots, one gets more solutions by multiplying by t.
- For complex roots of the form $r = a + bi$, the solution e^{rt} is complex-valued. Taking the real and imaginary parts gives two real-valued solutions, which are $e^{at}\cos(bt)$ and $e^{at}\sin(bt)$. The complex number $a - bi$ is also a root, and gives two more real-valued solutions, but these are redundant.
- By superposition, one obtains many homogeneous solutions from the set of all roots. This actually gives you all homogeneous solutions.
- For inhomogeneous equations: when $g(t)$ is nice, you can (correctly!) guess the rough form of a particular solution, and then compute the coefficients. Then all other solutions are obtained by superposition.
- The form of the solution is roughly the form of the forcing function. See p181-2 of the book for a flowchart. Make sure you know how to treat all the possibilities! Common errors:
  - forgetting to multiply by t when part of your guess is already a homogeneous solution;
  - forgetting that every $\cos$ term must be paired with a $\sin$ term;
  - forgetting that the guess for $t^n e^{rt}$ is actually $(A_n t^n + \ldots + A_1 t + A_0)e^{rt}$ when r is not a root, i.e. forgetting the lower terms in a polynomial.
- In certain cases, you can compute the coefficients quickly. One such case is the *exponential response formula*. To solve $p(D)y = Ae^{rt}$ when $p(r) \neq 0$, the answer is $y = \frac{A}{p(r)}e^{rt}$. There are fancier versions, involving derivatives of p, when $p(r) = 0$ is a root.
- Another such case is the *sinusoidal response formula*. Really, it is just the exponential response formula applied to the real parts of a complex-valued ODE. To solve $p(D)y = Ae^{at}\cos(bt - \theta)$, when $p(a + bi) \neq 0$, the answer is $y = \frac{A}{M}\cos(bt - \theta - \varphi)$. Here, $p(a + bi) = Me^{\varphi i}$, so that M is the magnitude of $p(a + bi)$, and φ is the argument.

Spring equation

- The *spring equation* is $my'' + \gamma y' + ky = g(t)$. Here, the *mass* m and the *spring constant* k must be positive numbers. The *damping constant* γ can be zero (*undamped*), or a positive number.
- The behavior of the unforced (i.e. homogeneous) solutions are determined by $\gamma^2 - 4mk$, by the quadratic formula. When this is positive, there are two negative real roots, and it is *overdamped*. When it is negative, there are two imaginary roots with negative real part, and it is *underdamped*. When it is zero, there is a double negative real root, and it is *critically damped*.
- For overdamped ODEs, the solution will be zero at most once, depending on the initial conditions. For critically damped, the solution will be zero exactly once. For underdamped, the solution will be zero infinitely many times.
- All three cases decay to zero as time passes. Thus they are called *transient solutions*. Let us assume the equation is forced, with $g(t)$ sinusoidal. The periodic particular solution is called the *steady-state solution*, and all solutions tend to it as time passes.
- Sinusoidal functions can be written in the form $a\cos(wt) + b\sin(wt)$, or in *amplitude-phase form* $A\cos(wt - \varphi)$. One has $a + bi = Ae^{i\varphi}$.
- By the sinusoidal response formula, if $g(t) = C\cos(wt - \theta)$, then the solution is $y = \frac{C}{M}\cos(wt - \theta - \varphi)$. Then $\frac{1}{M}$ is the *gain*, which tells you how much the input is amplified by the spring equation. M is the magnitude of $p(wi)$, which is some formula involving w. Maximizing the gain is then a calculus problem.
- When there is enough damping, the gain is always maximized when $w = 0$, which is to say that no input force is amplified by the spring. However, when $\gamma^2 - 2mk$ is negative so that the system is very underdamped, then there is a non-zero frequency w which maximizes the gain. This is called *practical resonance*.

First order linear systems with constant coefficients, or 1LSyswCC

- Given a homogeneous 1LSyswCC, you can rewrite it in matrix-vector form as $D\mathbf{x} = A\mathbf{x}$ for some matrix A of real numbers.
- An *eigenvector* $\mathbf{v}$ for a matrix A is a vector satisfying $A\mathbf{v} = \lambda\mathbf{v}$ for some real number λ, which is called its *eigenvalue*. If $\mathbf{v}$ is an eigenvector with eigenvalue λ, then so is $c\mathbf{v}$ for any real number c.
- The *trace* of a matrix is the sum of the diagonal entries. The *determinant* of a matrix is a number which can be computed inductively, involving $n!$ terms for an $n \times n$ matrix. For $n = 2$ and $n = 3$ there are quick rules to compute it.
- The *characteristic polynomial* of a matrix A is $p_A(\lambda) = \det(A - \lambda I)$. For a 2×2 matrix, it is equal to $\lambda^2 - (\mathrm{Tr} A)\lambda + \det A$. Every eigenvector of A is a root of p_A, and vice versa.
- Given an eigenvalue λ, to compute the eigenvector, one should compute $A - \lambda I$ and find a vector $\mathbf{v}$ for which $(A - \lambda I)\mathbf{v} = 0$. For 2×2 matrices, there is a quick trick to do this.

- If $\mathbf{v}$ is an eigenvector for A with eigenvalue λ, then $\mathbf{x} = e^{\lambda t}\mathbf{v}$ is a solution to $D\mathbf{x} = A\mathbf{x}$. This is easy to verify! By superposition, one can obtain other solutions as sums of solutions of this form. If there are no repeated roots of p_A, then all solutions are obtained in this way.

- If there are repeated roots, the solutions can be harder, and involve *generalized eigenvectors*, something we didn't really cover in this course.

- If an eigenvalue $\lambda = a + bi$ is complex, then $e^{\lambda t}\mathbf{v}$ is a complex-valued solution. Taking the real and imaginary parts gives two real-valued solutions, and superposition of these yields all the real-valued solutions corresponding to the eigenvalues $a \pm bi$. (The eigenvalue $a - bi$ also gives two more real-valued solutions, but these are redundant.)

- The *phase portrait* of a 1LSyswCC for two functions is the graph, for a solution $\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$, of x_1 versus x_2 at each time. This is drawn as a trajectory in the x_1x_2-plane, with an arrow indicating the motion of time. Many examples are found in Chapter 7 of the book.

- One can also draw a *phase arrow* at each point $\mathbf{v}$ in the x_1x_2 plane, indicating the trajectory of whichever solution $\mathbf{x}$ passes through that point at some time. The phase arrow at point $\mathbf{v}$ is given by $A\mathbf{v}$, because $\mathbf{x}' = A\mathbf{v}$ for a solution.

- Any trajectory that is a line coming out/in from the origin corresponds to an eigenvector, and a solution of the form $e^{\lambda t}\mathbf{v}$.

- One can classify the behavior of a system with two functions based on the trace and determinant of A. This is because, if r_1 and r_2 are the two eigenvalues, then $r_1 + r_2 = \mathrm{Tr} A$ and $r_1 r_2 = \det A$. This is encoded on the figure on p507 of your book. You should understand this chart, though you need not remember what happens on the boundaries (i.e. proper or improper nodes, or what happens when $\det = 0$). The origin is always an equilibrium solution, and you can tell whether it is stable or unstable based on the real part of the eigenvalues.

- When drawing the phase portrait of a node or saddle, you need to compute the eigenvalues so you can draw the straight line solutions. For a node, you need to know which eigenvalue is bigger. When drawing the phase portrait of a spiral, one can be pretty rough, but one should know whether it is clockwise or counterclockwise. This can be done by evaluating the phase arrow at some point. In general, when drawing any phase portrait, if a point is interesting, you should draw the phase arrow there, and make sure your trajectory is tangent!

- The phase portrait helps for qualitative analysis - telling you what happens to a solution with a given initial value, as time goes forward or backward. Section 9.1 has a good review of this material. In-depth examples are done in Sections 7.5 and 7.6.