Notes for MAT 308, Spring 2025

Joseph Helfer

Last updated: March 10, 2025

Contents

1 Jan 27, Intro and calculus review **3**
1.1 Introductory remarks . . . 3
1.2 Calculus review . . . 3

2 Jan 29, More calculus review and 10.1: Direction fields **7**
2.1 Calculus review continued . . . 7
2.2 10: First-order ODEs . . . 9
2.3 10.1: Direction fields . . . 10

3 Feb 3, 10.2: Separation of variables **12**
3.1 10.1: Direction fields . . . 12
3.2 10.2B Separation of variables . . . 13

4 Feb 5, 10.3: More separation of variables; existence and uniqueness theorem **15**
4.1 More of 10.2B Separation of variables . . . 15
4.2 End of 10.1: existence and uniqueness . . . 15

5 Feb 10, Proof of existence and uniqueness; 10.1: Numerical methods; and 10.3: Linear equations **17**
5.1 End of 10.1: existence and uniqueness . . . 17
5.2 10.1B: Numerical Methods . . . 17
5.3 10.3: Linear equations . . . 18

6 Feb 12, More linear equations, and linear algebra review **19**
6.1 10.3A: Exponential integrating factors . . . 19
6.2 10.3B Applications . . . 19
6.3 Abstract vector spaces . . . 20

7 Feb 17, More linear algebra **22**
7.1 Subspaces . . . 22
7.2 Basis, dimension, etc. . . . 22

8 Feb 19, Linear maps **24**
8.1 Linear maps . . . 24

9 Feb 24, 3.7: Inner products **26**
9.1 3.7A General properties of inner products . . . 26

10 Feb 26, 3.7B: Orthogonal bases; dual spaces **30**
10.1 3.7B Orthogonal bases . . . 30
10.2 3.7C Orthogonal transformations . . . 32
10.3 Dual spaces . . . 33

11 Mar 3, 11.1: differential operators **35**
11.1 11: Second-order equations . . . 35
11.2 11.1: Differential operators . . . 35

12 Mar 5, 11.1B and 11.2A: More second-order equations **37**
12.1 11.1B Factoring operators . . . 37
12.2 11.2A Complex Exponentials . . . 38

13 Mar 10, 11.2A: Complex solutions **40**
13.1 More on complex exponentials . . . 40
13.2 Complex solutions to second-order equations . . . 41

1 Jan 27, Intro and calculus review

1.1 Introductory remarks

- Differential equations are the central formalism in all of the most important physical theories, and are of great importance in many other branches of science and engineering, as well as in pure mathematics.
  - For example: "Newton's Second Law", "Maxwell's Equations" from electromagnetism, "Einstein's Field Equations" from general relativity, "Schrödinger's equation" from quantum mechanics, are all systems of differential equations.
- Roughly speaking a differential equation is an equation containing (ordinary or partial) derivatives.
- Example: Newton's Second Law $F = ma$ is really $F = m\frac{\mathrm{d}^2u}{\mathrm{d}t^2}$, where $u(t)$ is the position of a given object with mass m.
  - F typically depends on u and maybe on $\frac{\mathrm{d}u}{\mathrm{d}t}$.
  - For example, for a massive body in a central gravitation field, we have $F = -G\frac{u}{|u|^3}$ for some constant G, so the equation becomes $-G\frac{u}{|u|^3} = m\frac{\mathrm{d}^2u}{\mathrm{d}t^2}$.
  - The "unknown" in this equation, for which one might try to solve, is u. It appears twice, with different amounts of derivatives. A solution would not be a *number* as in an ordinary equation but a *function* of time: the trajectory of the body in question.
- In Newton's equation, the variable u is actually a *vector*-valued – i.e., $\mathbb{R}^3$-valued – function.
  - Thus, the equation is actually a system of *three* equations.
  - As you learned in linear algebra, a system of linear equations is most fruitfully understood as a single matrix equation $Av = b$.
  - The same kind of thing is true in the theory of differential equations, and moreover there are some interesting and important further ideas in linear algebra which come in.
  - This is why linear algebra is also part of this course.
- Newton's equation only contains *ordinary* derivatives and is thus a so-called **Ordinary Differential Equation** (or **ODE**).
  - By contrast, all of the other equations mentioned above are **Partial Differential Equations** (or *PDE*s).
  - We will be dealing almost exclusively with ODEs in this class.
  - Just to give you an idea, a simple example of a PDE is the *heat equation*, which looks like this: $\frac{\partial u}{\partial t}(t,x) = -\frac{\partial^2 u}{\partial x^2}(t,x)$.

1.2 Calculus review

On engineers and mathematicians

- There are two ways of approaching the subject of differential equations.
  - The "mathematician's way" gives precise definitions of all the objects under consideration, and rigorous proofs on the basis of those definitions.

- The "engineer's (or physicist's) way" rests on a basic intuitive understanding of all of the basic objects (real numbers, continuous functions, limits, etc.) and their properties, and this intuition is based on the physical objects and phenomena that these mathematical objects represent.
- Feynman ("The character of physical law"): *There are two kinds of ways of looking at mathematics, which [...] I [...] call the Babylonian tradition and the Greek tradition. In Babylonian schools in mathematics the student would learn something by doing a large number of examples until he caught on to the general rule. [...] Also he would know a large amount of geometry, a lot of the properties of circles, the theorem of Pythagoras, formulae for the areas of cubes and triangles; in addition, some degree of argument was available to go from one thing to another. [...] But Euclid discovered that there was a way in which all of the theorems of geometry could be ordered from a set of axioms that were particularly simple. The Babylonian attitude – or what I call Babylonian mathematics – is that you know all of the various theorems and many of the connections in between, but you have never fully realized that it could all come up from a bunch of axioms.*

- **Both of these ways are valuable and important**.
  - And in fact, the "ideal mathematician" or "ideal engineer" don't exist; everyone is somewhere in between the two.
  - As this is a *mathematics* course (MAT 308), we must to some extent emphasize the "mathematician's way": in principle, you should know the precise definitions of all the objects we consider, be able to give precise statements of all the theorems we consider, and be able to follow the proofs, and reproduce the simpler ones. (In principle, this also means being familiar with the axioms of set theory – which, though a good thing, we will not insist on.)
  - Thus, we begin with a quick review of the basic elements of set theory and calculus which we will be using, to make sure we are all on the same page.

Sets

- We take for granted all the notions of set theory, and the basic properties of and operations with sets, which you know well by now.
- The most important sets are $\mathbb{N} = \mathbb{Z}_{\geq 0}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$ for the naturals, integers, rationals, reals, and complex numbers, to all of which we shall return shortly.
  - We write $\mathbb{N}_+$ or $\mathbb{Z}_{>0}$ or something like that for the set of positive integers.
- The product of two sets is the set of all ordered pairs $X \times Y = \{(x, y) \mid x \in X, y \in Y\}$.
- We write $A \subset B$ or $A \subseteq B$ for "A is a subset of B".
  - If we want to express that A is a *proper* subset of B, we write $A \subsetneq B$
- A <u>relation</u> R between sets X and Y is a subset of the product $R \subset X \times Y$.
  - We write "xRy" for $(x, y) \in R$.
- A <u>function</u> (or <u>mapping</u> or <u>map</u>) $f\colon X \to Y$ is a relation $f \subset X \times Y$ such that for every $x \in X$, there is a unique $y \in Y$ with $xfy.$, i.e.

$$\forall x \in X, \exists! y \in Y, xfy.$$

  - We write $f(x)$ for the unique $y \in Y$ such that xfy.

- A(n infinite) sequence in a set X is a function $\mathbb{N} \to X$ (or $\mathbb{Z}_{\geq n} \to X$ for any $n \in \mathbb{Z}$).
  - As usual, we write $(x_i)_{i=0}^{\infty}$ for the sequence $\mathbb{N} \to X$ with $i \mapsto x_i$ for $i \in \mathbb{N}$.
- A finite sequence of length $n \in \mathbb{N}$ or n-tuple in a set X is a function $\{1, \ldots, n\} \to X$.
  - We write X^n for the set of finite sequences of length n.
  - As usual, we write $v = (v_1, \ldots, v_n)$ for the tuple $\{1, \ldots, n\} \to X$ given by $i \mapsto v_i$.
  - More generally, given sets $X_1, \ldots, X_n$, we write $X_1 \times \cdots \times X_n$ for the set of finite sequences $(x_1, \ldots, x_n)$ with $x_i \in X_i$ for each i.
- We write $X \cong Y$ if there exists a bijection between X and Y, and we write $f \colon X \xrightarrow{\sim} Y$ to indicate that f is a bijection.
- We write Y^X for the set of functions $X \to Y$ and $\mathcal{P}(X)$ for the power set of X, i.e., the set of subsets of X.
  - We have $2^X \cong \mathcal{P}(X)$ for any set X, where 2 is the set $\{0, 1\}$.

The real numbers

- We recall the main properties of the set of real numbers $\mathbb{R}$ with its operations $+ \colon \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and $\cdot \colon \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and its ordering " $\leq$ " $\subset \mathbb{R} \times \mathbb{R}$:
- **Field axioms**
  - addition and multiplication are commutative and associative, and multiplication distributes over addition
  - 0 and 1 are identity elements for addition and multiplication, respectively (i.e., $0 + a = a$ and $1 \cdot a = a$ for all $a \in \mathbb{R}$).
  - (*Exercise*: 0 and 1 are each uniquely determined by this property!)
  - Every element $a \in \mathbb{R}$ has an additive inverse $(-a)$ (i.e., $a + (-a) = 0$).
  - Every element $a \in \mathbb{R} - \{0\}$ has a multiplicative inverse a^{-1}.
- **Ordering axioms**
  - $\leq$ is reflexive, meaning $a \leq a$ for all $a \in \mathbb{R}$.
  - $\leq$ is antisymmetric, meaning $a \leq b \wedge b \leq a \Rightarrow a = b$ for all $a, b \in \mathbb{R}$.
  - $\leq$ is transitive, meaning $a \leq b \wedge b \leq c \Rightarrow a = c$ for all $a, b, c \in \mathbb{R}$.
  - $\leq$ is total, meaning $a \leq b \vee b \leq a$ for all $a, b \in \mathbb{R}$.
  - If $a \leq b$ then $a + c \leq b + c$ for any $a, b, c \in K$.
  - If $a \leq b$ and $c \geq 0$, then $a \cdot c \leq b \cdot c$ for any $a, b, c \in K$.
- **Completeness axiom**
  - Given a set $S \subset \mathbb{R}$, we say that $a \in \mathbb{R}$ is an upper bound for S if $s \leq a$ for all $s \in \mathbb{R}$, and we say that S is bounded-above if it has some upper bound.
  - Completeness axiom: each bounded-above subset $S \subset \mathbb{R}$ has a least upper bound (or supremum) $\sup S$, meaning that $\sup S$ is an upper bound for S, and that $\sup S \leq a$ for every upper bound a for S.
- In short, one says that $(\mathbb{R}, +, \cdot, \leq)$ is a complete ordered field.

- On the basis of these axioms, one can prove all the familiar properties of the algebraic operations and the ordering relation (an activity you have hopefully done before, and should try your hand at if you haven't).
  - In fact, these axioms **completely determine** the real numbers in the following sense: given any system $(K, \tilde{+}, \tilde{\cdot}, \tilde{\leq})$ consisting of a set K, binary operations $\tilde{+}$ and $\tilde{\cdot}$ on K, and a binary relation $\tilde{\leq}$ on K satisfying the above axioms, there exists a unique bijection $F\colon \mathbb{R} \to K$ satisfying $F(a+b) = F(a)\tilde{+}F(b)$, $F(a \cdot b) = F(a)\tilde{\cdot}F(b)$, and $a \leq b \iff F(a)\tilde{\leq}F(b)$ for all $a, b \in K$.
    - (Such a bijection F is called an "isomorphism of ordered fields".)
    - (The aforementioned theorem establishes the *uniqueness* (up to isomorphism) of the real numbers. One may still reasonably inquire about their *existence*: why must there exist such a set satisfying these axioms at all? We simply take it for granted that it exists; if one wants to *construct* it, this must be on the basis of some more basic axioms: either the Peano axioms of arithmetic, or else the axioms of set theory.)
- The other number systems
  - We define the sets $\mathbb{N}, \mathbb{Z}, \mathbb{Q} \subset \mathbb{R}$ as follows.
  - $\mathbb{N} \subset \mathbb{R}$ is the smallest subset such that $0 \in \mathbb{N}$ and such that $n+1 \in \mathbb{N}$ whenever $n \in \mathbb{N}$.
    - Concretely, this means that $\mathbb{N}$ is the intersection $\bigcap_{S \subset \mathbb{R}} S$ of all subsets of $\mathbb{R}$ satisfying these two properties.
    - It follows immediately from this that $\mathbb{N}$ satisfies the all-important **principle of mathematical induction**: given any $S \subset \mathbb{N}$ such that $0 \in S$ and $n+1 \in S$ whenever $n \in S$, we have $S = \mathbb{N}$.
  - $\mathbb{Z} = \{a - b \mid a, b \in \mathbb{N}\} \subset \mathbb{R}$.
  - $\mathbb{Q} = \{a/b \mid a, b \in \mathbb{Z} \wedge b \neq 0\} \subset \mathbb{R}$.
  - Finally, $\mathbb{C}$ is simply defined to be $\mathbb{R}^2$.
    - Given $(a_1, b_1), (a_2, b_2) \in \mathbb{C}$, we set $(a_1, b_1) + (a_2, b_2) := (a_1 + b_1, a_2 + b_2)$, and $(a_1, b_1) \cdot (a_2, b_2) := (a_1 a_2 - b_1 b_2, a_1 b_2 + a_2 b_1)$.
    - One may check that this makes $\mathbb{C}$ into a *field* (though of course, not an *ordered* field, since there is no natural ordering of points in the plane).
    - Given $a \in \mathbb{R}$, we write a as a shorthand for $(a, 0) \in \mathbb{C}$.
    - We also write i as a shorthand for $(0, 1) \in \mathbb{C}$; we have $i^2 = -1$.
    - Thus, we may write $(a, b) = a + bi$ for every $(a, b) \in \mathbb{C}$, and we then have $(a_1 + b_1 i) + (a_2 + b_2 i) = (a_1 + b_1) + (a_2 + b_2)i$ and $(a_1 + b_1 i) \cdot (a_2 + b_2 i) = a_1 a_2 + b_1 b_2 i^2 + a_1 b_2 i + a_2 b_1 i$.

2 Jan 29, More calculus review and 10.1: Direction fields

2.1 Calculus review continued

Calculus

- We will be reviewing various topics from calculus (and linear algebra) as they come up; we will just recall a couple of the most important ones here
- We consider functions $f\colon I \to \mathbb{R}$ defined on some domain $I \subset \mathbb{R}$.
  - Typically, $I \subset \mathbb{R}$ will be an interval: a set of the form (a, b), $[a, b)$, $(a, b]$, or $[a, b]$, where $a \in \mathbb{R} \cup \{-\infty\}$ and $b \in \mathbb{R} \cup \{\infty\}$ and $a \leq b$.
- Limits:
  - Given a function $f\colon I \to \mathbb{R}$ and $x \in I$, we have the limit $\lim_{x \to a} f(x)$ and the one-sided limits $\lim_{x \nearrow a} f(x)$ and $\lim_{x \searrow a} f(x)$ (each of which may or may not exist).
  - Ideally, you should know the **definition** of the limit: $\lim_{x \to a} f(x) = b$ means $\forall \epsilon > 0, \exists \delta > 0, \forall x \in I, (0 < |x - a| < \delta \Rightarrow |f(x) - b| < \varepsilon)$.
  - But the most important thing is that you know the basic **rules** for limits: $\lim_{x \to a} f(x) + g(x) = \lim_{x \to a} f(x) + \lim_{x \to a} g(x)$ (assuming the right-hand sides exists), and so on.
- A function $f\colon I \to \mathbb{R}$ is continuous at $a \in I$ if $\lim_{x \to a} f(x) = f(a)$, and is continuous on I if it is continuous at each $a \in I$.
  - Continuous functions satisfy the **intermediate value theorem**.
- Derivatives:
  - $f\colon I \to \mathbb{R}$ is differentiable at $a \in I$ if $\lim_{h \to 0} \frac{f(a+h) - f(a)}{h}$ exists, and if so, this limit is called $f'(a)$.
  - f is differentiable on I if it is differentiable at each $a \in I$; in this case, we have the derivative $f'\colon I \to \mathbb{R}$.
    - **Theorem:** Differentiability implies continuity.
  - We sometimes write $y = f(x)$ and then denote the derivative by $f'(x) = \frac{\mathrm{d}y}{\mathrm{d}x}$.
    - Here, we are thinking of x and y as "variables", with y varying "dependently" on x; then $\mathrm{d}x$ represents an "infinitesimal variation" of x, $\mathrm{d}y$ represents the resulting infinitesimal variation of y, and the derivative is their quotient.
    - This is the classical way of thinking about functions, which has been preserved by the physicists and engineers, but among mathematicians has been replaced by the set-theoretic perspective.
    - With some care, this notation can be used in a consistent and rigorous way (in fact, doing so leads to some very interesting mathematics), but in general, it is best to follow the physicists and engineers, and freely use our intuition about infinitesimals without worrying about the formal definition; and then if we want a rigorous proof, we can always revert to the formal, set-theoretic definitions.
    - In any case, we often simply write $\frac{\mathrm{d}f}{\mathrm{d}x}$ for f'.
  - We write f'' or $\frac{\mathrm{d}^2 f}{\mathrm{d}x^2}$ for the derivative of f' (if it exists), and f''' and $f^{(4)}$ and $f^{(5)}$ and so on.
  - Again the most important thing is to know all the **rules** for derivatives: the sum rule, product rule, quotient rule, **chain rule**, the rule for x^n, and the derivatives of everyone's favourite functions: sin, cos, e^x, $\ln x$, and so on.

- We denote by $\mathcal{C}^k(I)$ the set of functions $f\colon I \to \mathbb{R}$ such that the derivatives $f', f'', \ldots, f^{(k)}$ exist and are continuous; $\mathcal{C}^0(I)$ is simply the set of continuous functions on I, and $\mathcal{C}^\infty(I)$ is the set of infinitely differentiable (or smooth) functions on I. All of the most common functions are smooth.

- Partial derivatives
  - We consider functions $f\colon U \to \mathbb{R}$ defined on a domain $U \subset \mathbb{R}^n$.
  - Often, U will be all of $\mathbb{R}^n$, or maybe some product of intervals $U = I_1 \times \cdots \times I_n$.
  - The definition of **limit** and **continuity** for such functions is exactly the same as in the single-variable case,
    - except that now, the absolute values $|x - a|$ and so on appearing in the definition are now considered norms.
    - We recall the norm of $v \in \mathbb{R}^n$ is given as in the Pythagorean theorem: $|v| = \left(\sum_{i=1}^n v_i^2\right)^{1/2}$.
  - Given $f\colon U \to \mathbb{R}$ and $a \in U$, the i-th partial derivative of f at a, if it exists, is the "derivative of f at a in the i-th coordinate direction, holding all other coordinates constant".
  - For example, if $n = 2$, the two partial derivatives at $(a, b) \in U$ are
$$\lim_{h\to 0} \frac{f(a+h, b) - f(a, b)}{h} \qquad \text{and} \qquad \lim_{h\to 0} \frac{f(a, b+h) - f(a, b)}{h}.$$
  - In general, the i-th partial derivative at $a \in U$ is
$$\lim_{h\to 0} \frac{f(a_1, \ldots, a_{i-1}, a_i + h, a_{i+1}, \ldots, a_n) - f(a_1, \ldots, a_n)}{h}.$$
  - There are many notations for the i-th partial derivative; the simplest are $\partial_i f$ or f_i; often we write $y = f(x_1, \ldots, x_n)$, and then write $\frac{\partial y}{\partial x_i}$ or $\frac{\partial f}{\partial x_i}$ or $\partial_{x_i} f$ or f_{x_i}.
  - In practice, one evaluates the partial derivative with respect to x_i by "pretending all the other variables are constant and taking the ordinary derivative with respect to x_i".

- The chain rule
  - The sum, product, quotient rules for partial derivatives are just like those for ordinary derivatives.
  - The **chain rule** is more interesting: if $y_i = f_i(x_1, \ldots, x_m)$ for $i = 1, \ldots, n$ and $z = g(y_1, \ldots, y_n)$, so that
$$z = g\big(f_1(x_1, \ldots, x_m), \ldots, f_n(x_1, \ldots, x_m)\big),$$
    then
$$\frac{\partial z}{\partial x_i} = \sum_{j=1}^n \frac{\partial z}{\partial y_j} \cdot \frac{\partial y_j}{\partial x_i}.$$
  - The chain rule is most elegantly expressed using the **Jacobian** (or **total derivative**) matrix. For this, we must recall a bit about vector-valued functions and matrix multiplication, and we will do this when it comes up.

- Sequences and series
  - The limit $\lim_{n\to\infty} x_n$ of a sequence $(x_n)_{n=0}^\infty$, if it exists, is defined as the unique $a \in \mathbb{R}$ such that $\forall \epsilon > 0, \exists N \in \mathbb{N}, \forall n \geq N, |x_n - a| < \varepsilon$.
    - We also write $x_n \xrightarrow{n\to\infty} a$.

- Again, the most important thing is that you know the main **rules** for limits: the sum rule $\lim_{n\to\infty} a_n + b_n = \lim_{n\to\infty} a_n + \lim_{n\to\infty} b_n$ (assuming the right-hand side exists), the difference and quotient rules, and the rule that if f is a function which is continuous at a and $x_n \xrightarrow{n\to\infty} a$, then $f(x_n) \xrightarrow{n\to\infty} f(a)$.
- The sum of the infinite series $\sum_{n=0}^{\infty} a_n$, if it exists, is the limit $\lim_{N\to\infty} S_n$ of the partial sums $S_N = \sum_{n=0}^{N} a_n$.
- You should know the various **tests** for convergence: the comparison test, alternating series test, absolute convergence test, ratio test, root test, integral test.

- Integration
  - If $I \subset \mathbb{R}$ contains the interval $[a, b]$ and $f\colon I \to \mathbb{R}$ is any function, we can talk about the integral $\int_a^b f(x)\,\mathrm{d}x$, which may or may not exist, but always *does* exist when f is continuous, or even continuous outside of finitely many points.
  - Again, it is best – though perhaps not essential – to know the **definition** of the integral, in terms of Riemann sums: at least for continuous functions, the integral $\int_a^b f(x)\,\mathrm{d}x$ is given by the limit $\lim_{N\to\infty} \frac{b-a}{N} \sum_{i=1}^{N} f(a + i\frac{b-a}{N})$.
    - (In general (for not necessarily continuous f), the definition is a bit more complicated: one has to consider arbitrary partitions $a = x_0 \leq \cdots \leq x_N = b$, and not just the "regular partition" $x_i = a + i\frac{b-a}{N}$.)
    - The notation $\int_a^b f(x)\,\mathrm{d}x$ is meant to represent an "infinite sum of infinitesimals" (the integral sign is an "S" for "sum"). This is the classical notion, which again has been superseded by the modern set-theoretic definition. Again, it can be made rigorous with some effort, and again, we do well to follow the physicists and engineers in freely using this intuition when it is useful.
  - Again, the most important thing to know are the various **rules**, among which the most important is the **fundamental theorem of calculus**:
    - The integral $\int_a^b f(x)\,\mathrm{d}x$ is equal to $F(b) - F(a)$ where $F\colon [a, b] \to \mathbb{R}$ is any anti-derivative of f, meaning a function with $F' = f$ (assuming that an anti-derivative exists!).
    - Moreover, if f is continuous, then an anti-derivative *does* exist.
    - When anti-derivative F of f exists is uniquely determined up to a constant $C \in \mathbb{R}$, and the "indefinite integral" is defined as $\int f(x)\,\mathrm{d}x = F(x) + C$.
  - Other rules you should know include: the sum and difference rule, the substitution rule, and integration by parts, as well as the million little tricks you've learned to apply these in various special cases.
  - We will review the important topics of multi-variable integrals and improper integrals if and when they come up.

2.2 10: First-order ODEs

- We now begin our study of differential equations.
  - As a preliminary definition, a *differential equation of order n* is an equation whose unknown is a function $y(x)$, and which may involve the derivatives of y up to order n, that is $y, y', y'' \ldots, y^{(n)}$, as well as x itself.
  - Examples of order 1 and 2 are

$$y'(x) + y(x) = x \qquad \text{and} \qquad y''(x) + y'(x) = 0.$$

which we usually just write as

$$y' + y = x \qquad \text{and} \qquad y'' + y' = 0.$$

- A solution to this equation on an interval $I \subset \mathbb{R}$ is then a function $y\colon I \to \mathbb{R}$ which satisfies this equation for all $x \in I$ (in particular, y must be n times differentiable).
  - In principle, one can consider differential equations on *any* domain $I \subset \mathbb{R}$, and not just on an interval I. But if I is not connected, say for example, $I = I_1 \cup I_2$, where I_1 and I_2 are disjoint intervals, then a solution to the differential equation on I just amounts to a solution on I_1 and a separate solution on I_2.
- For example, $y(x) = x - 1$ and $y(x) = e^{-x}$ are solutions to the above two equations.
- (As with ordinary equations, we are often not interested in finding *a* solution, but in finding *all* possible solutions.)

- The above definition will be perfectly satisfactory for our purposes, though it is also good to give a more precise, formal definition of what we mean by a "differential equation", which one can do as follows.
  - We define a(n ordinary) differential equation of order n to be an arbitrary function $F\colon \mathbb{R}^{n+1} \to \mathbb{R}$ (or more generally $U \to \mathbb{R}$ for some $U \subset \mathbb{R}^{n+2}$).
    - Intuitively, we think of F as representing the equation $y^{(n)} = F(x, y, y', \dots, y^{(n-1)})$.
    - (One can more generally consider so-called "implicit" equations given by a function $E(x, y, y', \dots, y^{(n)})$, which we think of as representing the equation $0 = F(x, y, y', \dots, y^{(n)})$, but this is rarely done.)
  - A solution to the differential equation F on $I \subset \mathbb{R}$ is then a n-times differentiable function $y\colon I \to \mathbb{R}$ such that $((x, y(x), y'(x), \dots, y^{(n-1)}(x)) \in U$ and) $y^{(n)}(x) = F(x, y(x), y'(x), \dots, y^{(n-1)}(x)) = 0$ for all $x \in I$.
  - For example, in this formalism, the above two equations would be given by the functions $F(x, y_0) = x - y_0$ and $E(x, y_0, y_1) = -y_1$, respectively.
  - An example where we have to take some $U \subsetneq \mathbb{R}^n$ is the equation $y'' = 1/y' + y$, where $F(x, y_0, y_1) = 1/y_1 + y_0$, and so the largest U we can take is $\{(x, y_0, y_1) \in \mathbb{R}^4 \mid y_1 \neq 0\}$.

- To begin with, we will be considering only **first-order** differential equations.

2.3 10.1: Direction fields

- We consider an arbitrary first-order equation

$$y' = F(x, y).$$

  - Any solution $y(x)$ passing through (x_0, y_0) must have slope $F(x_0, y_0)$ at that point.
  - Thus, we think of $F(x, y)$ as assigning a slope to each point (which we can represent as a short line segment through that point with that slope); this is called a **slope field** or **direction field**.
  - This is not to be confused with a **vector field**.
  - We can draw the direction field $F(x, y)$ by sampling several points and drawing line segments, or we can first find several solutions to the differential equation, and then simply plot various tangents along the resulting graphs.

Example 10.1.1

- Consider the equation $y' = -y/x$ for $x \neq 0$.
- We sample the slopes at a few points:

(x, y)	$y' = -y/x$
(1,1)	-1
(1,2)	-2
(2,1)	-1/2
(-1,2)	2
(-2,2)	1

- We obtain a picture as in Figure 10.2 (a)
- Let's solve the equation!
  - Let $I \subset \mathbb{R}$ be one of $(-\infty, 0)$ or $(0, \infty)$.
  - Suppose that $y\colon I \to \mathbb{R}$ is a solution to $y' = -y/x$.
  - Rearrange: $xy' + y = 0$ (equivalent by the assumption $x \neq 0$).
  - Product rule: $(xy)' = 0$.
  - Thus $x \cdot y(x)$ is some constant c.
  - So $y = c/x$.
- Thus, every solution is of the form $y = c/x$.
  - Conversely, since each step above was a **logical equivalence**, it follows that $y = c/x$ **is** always a solution (as one can also check directly).
  - We may say that the general solution to the equation is $y = c/x$ with $x \in \mathbb{R}$.
- We can plot the solutions and see they are tangent to our slope field.

3 Feb 3, 10.2: Separation of variables

3.1 10.1: Direction fields

Example 10.1.2

- If $F(x, y)$ is independent of y, the equation takes the form

$$y' = G(x);$$

  assume G is continuous.

- Then the solutions are simply the antiderivatives of G: $y = \int G(x)\,\mathrm{d}x + C$.
  - All of the solution curves will be parallel to each other, as one can also see from the direction field.
- Example: $y' = \cos x$.
  - Then the general solution is $y = \sin x + c$.
  - Note: when discussing the solutions to a differential equation, one should always first state *on which interval* I one is seeking solutions $y\colon I \to \mathbb{R}$, though this is often left implicit.
  - So here, we should say we are seeking solutions y defined on all of $\mathbb{R}$, and then the general such solution is $y = \sin x + c$.

Example 10.1.3

- We see that differential equations tend to have many solutions.
- We can add extra conditions to a differential equation to single out a particular solution.
- If we demand that y passes through a given point (x_0, y_0), i.e., that $y(x_0) = y_0$, this is called an **initial condition**. (More generally, we can demand that $y^{(i)}(x_i) = y_i$.)
- The problem of satisfying a differential equation with a given initial condition is called an **initial value problem** (IVP).
- The name comes from thinking of x as the "time parameter" and x_0 as the "initial time".
- We return to $y' = -y/x$, $x \neq 0$ with solutions $y = c/x$.
  - Now consider the initial value problem $(x_0, y_0) = (1/2, 2)$.
  - We thus have $2 = c/(1/2)$ and hence $c = 1$.
  - The (unique!) solution to the IVP is thus $y = 1/x$, $x > 0$.
- We can also solve the general IVP for a given (x_0, y_0).
  - We have $y_0 = c/x_0$, so $c = x_0 y_0$, and the solution is $y = x_0 y_0/x$, with $x \in (0, \pm\infty)$ depending on the sign of x_0.

3.2 10.2B Separation of variables

- Among the first-order equations $y' = F(x, y)$, we have seen that the simplest case is when $F(x, y)$ doesn't depend on y. We can then solve it simply by integration.
- We now study the more general case in which in which $F(x, y)$ is of the form $f(x) \cdot h(y)$.
  - Supposing h is non-zero and setting $g(y) = h^{-1}(x)$, these are thus the equations of the form $g(y) \cdot y' = f(x)$.
  - These can often also be solved simply by integration, using a method called "separation of variables".
  - Writing $g(y)\frac{\mathrm{d}y}{\mathrm{d}x} = f(x)$, we obtain "$g(y)\,\mathrm{d}y = f(x)\,\mathrm{d}x$" and hence
$$\int g(y)\,\mathrm{d}y = \int f(x)\,\mathrm{d}x.$$
  - Computing antiderivatives F and G of f and g, we thus have $G(y) = F(x) + C$, and can thus find y as long as we can solve for y in this (non-differential-)equation.
  - (We can do the above computation without using "differentials": we simply have $\frac{\mathrm{d}G(y(x))}{\mathrm{d}x} = y'(x) \cdot g(y(x)) = f(x)$ by the chain rule, and hence $G(y(x))$ is an antiderivative of $f(x)$.)

Example 10.2.3

- We consider the differential equation $\frac{\mathrm{d}P}{\mathrm{d}t} = kP$, where $k > 0$ is a constant.
  - This describes the rate of growth of a population of bacteria of size $P(t)$ (say in grams), in which each bacterium reproduces at a rate of k bacteria per second.
  - As always in mathematical modelling, we are making simplifying assumptions in this equation, so we can only expect the solution of the equation to give an approximation to the population growth of the bacteria.
- We know from experience that the solution is $P(t) = Ke^{kt}$ for some constant K.
  - But let us see how we could have arrived at that using separation of variables.
  - Assuming P is always positive, we have $\int \frac{1}{P}\,\mathrm{d}P = \int k\,\mathrm{d}t$, hence $\ln P = kt + C$, hence $P(t) = e^{kt+C} = e^C \cdot e^{kt}$.
  - The solution to the IVP $P' = kP; P(0) = P_0$ is thus $P(t) = P_0 \cdot e^{kt}$.
  - (Assuming instead that P is always *negative* would given $\ln(-P) = kt + C$ and hence $P(t) = -e^C \cdot e^{kt}$ and thus again $P(t) = P_0 e^{kt}$; assuming just that P is non-zero would also lead to the same conclusion, since being differentiable and hence continuous, it is then always positive or negative.)
  - **Be careful**: in this computation, we had to make the assumption that P is non-zero, so that we could divide by it. If this assumption wasn't valid, our solution may be incorrect. However, we can just check directly that our solution *is* correct.
- This obviously isn't a realistic solution for large t; one problem is that our model doesn't include the amount of *food* available to the bacteria.
- If we look carefully at what we did, we see that we *almost* showed that this is the unique solution to the IVP; the only problem is we had to assume P was non-zero; the possibility remains that there are other solutions which are zero somewhere (besides the obvious one $P(t) = 0$).

- The trick to show uniqueness without this assumption is to consider the product $P(t) \cdot e^{-kt}$.
- We then have $\frac{\mathrm{d}}{\mathrm{d}t}(Pe^{-kt}) = kPe^{-kt} - kPe^{-kt} = 0$. Hence $e^{-kt} \cdot P(t)$ is equal to a constant K and $P = Ke^{-t}$, as desired.
- We will return to this trick of "exponential multipliers" soon.
- (We also have uniqueness by the general uniqueness theorem we stated, but that is overkill.)

Example 10.2.6

- We consider a tank of a chemical solution, in which a particular chemical is flowing in and out at a particular rate.
- We let $S = S(t)$ be the amount of chemical at time t. Then

$$\frac{\mathrm{d}S}{\mathrm{d}t} = (\text{rate of inflow}) - (\text{rate of outflow}).$$

- Example: a 100-gallon tank contains 150 pounds of salt of salt in solution at time $t = 0$.
  - A salt solution with 2 pounds of salt per gallon is being added at a rate of 2 gallons per minute.
  - The solution, which we take to be homogeneous, is flowing out of the tank at a rate of 2 gallons per minute.
  - Thus salt is flowing in at 4 pounds per minute, and out at $2S(t)/100$ pounds per minute at time t.
- Thus S satisfies the IVP $\frac{\mathrm{d}S}{\mathrm{d}t} = 4 - \frac{2S}{100}; S(0) = 150$.
- We solve it using separation of variables
  - We have $\int \frac{\mathrm{d}S}{S-200} = -\int \frac{\mathrm{d}t}{50}$; we see that this will only be valid if $S - 200$ remains non-zero for all time. Since $S(0) - 200 = -50$, we suppose it remains negative for all t.
  - Hence $\ln\big(-(S-200)\big) = -t/50 + C$.
  - Hence $S = 200 - e^C e^{-t/50}$.
  - The initial condition gives $150 = S(0) = 200 - e^C$ and hence $e^C = 50$.
  - Hence $S = 200 - 50e^{-t/50}$; this is plotted in Figure 10.9; it has a positive slope and an asymptote $\lim_{t\to\infty} S = 200$.
  - Checking directly, we see that this is indeed a solution to the equation.
- (Again, we had to assume $S \neq 200$ here, so we haven't proven that this is the unique solution; again, we can circumvent this using an exponential multiplier.)
  - Note however that the uniqueness *is* important, from a scientific perspective: to draw the conclusion that the amount of chemical in the tank evolves in the way that we found, on the basis of the fact that it solves the given IVP, we have to know that it is the **only** solution to the IVP.

4 Feb 5, 10.3: More separation of variables; existence and uniqueness theorem

4.1 More of 10.2B Separation of variables

Example 10.2.4

- Consider $y' = y/x$, for $x \in \mathbb{R} - \{0\}$.
  - The direction field is shown in Figure 10.7.
- By separation of variables we have $y^{-1}\,\mathrm{d}y = x^{-1}\,\mathrm{d}x$, hence $\ln|y| = \ln|x| + C$, hence $|y| = e^C|x|$, hence $y = \pm e^C x$.
  - Thus, the solutions are $y = kx$ for $k \in \mathbb{R}$.
  - (We should really be more careful here and say that we are looking for solutions on $(-\infty, 0)$ or on $(0, \infty)$, since otherwise we could also have solutions like
$$y(x) = \begin{cases} x & x < 0 \\ 2x & x > 0. \end{cases}$$
  - (The same is true in general when we say $\int x^{-1}\,\mathrm{d}x = \ln|x| + C$: if we take the function x^{-1} to be defined on all of $\mathbb{R} - \{0\}$, then the anti-derivative is only determined up to *two* constants: one on $(-\infty, 0)$ and one on $(0, \infty)$.)
- To prove uniqueness without needing y to be non-zero, we can use a similar trick to last time: we have $0 = y' - y/x$ and hence $0 = y'/x - y/x^2 = \frac{\mathrm{d}(y/x)}{\mathrm{d}x}$. Hence y/x is a constant.
  - (This is secretly again a case of the method of "exponential multipliers".)

4.2 End of 10.1: existence and uniqueness

Example 10.1.4

- In the previous examples, there was a unique solution to each IVP $y' = F(x, y); y(x_0) = y_0$, and this solution extended over the whole given domain; both of these features can fail.
- Consider
$$y' = \begin{cases} \sqrt{y}, & y \geq 0 \\ 0, & y < 0. \end{cases}$$
  - The direction field for this is shown in Figure 10.3 (a).
  - It has two solutions $y\colon \mathbb{R} \to \mathbb{R}$ passing through $(x_0, y_0) = (0, 0)$: $y(x) = 0$ and
$$y(x) = \begin{cases} 0, & x < 0 \\ x^2/4, & x \geq 0. \end{cases}$$
    - To conclude that y is differentiable and compute $y'(0)$, we use the **Theorem** from calculus that if a function f has a "left" derivative $\lim_{h \nearrow 0}(f(a+h)-f(a))/h$ and a "right" derivative $\lim_{h \searrow 0}(f(a+h)-f(a))/h$ at a given point a, and they agree, then f is differentiable at a.
  - In fact, there are infinitely many solutions to this IVP.
- As we will see, there is a simple condition on $F(x, y)$ which guarantees that this cannot not happen.

Example 10.1.5

- Now consider $y' = 1 + y^2$.
  - The direction field is shown in Figure 10.3 (b).
  - It has the solution $y = \tan x$ passing through $(x_0, y_0) = (0, 0)$.
  - This can only be extended to $x \in (-\pi/2, \pi/2)$ since it tends to $\pm\infty$, despite $F(x, y)$ being well-behaved everywhere.
  - In particular, had we stated the problem in the form "find a solution to this IVP which is defined on all of $\mathbb{R}$", the conclusion would have been that *this problem does not have a solution.*
- Again, this can be avoided by putting a simple condition on F.

Existence and uniqueness theorem

- We now state a general theorem which guarantees that a unique solution to a given IVP exists, if we assume that F satisfies certain conditions which in particular rule out the above pathologies.
  - The existence result should be cared to analogous phenomena for polynomial equations.
  - For degree 2 polynomial equations, we have an explicit method of solution, and even an explicit formula; this is analogous to the explicit methods we have introduced and will introduce to solve certain classes of differential equations.
  - On the other hand, we also know the general result (whose proof is easy): *every odd degree polynomial equation has a solution.* In this case, we are guaranteed on general grounds that a given equation has a solution, but we are not given any means to find it. The following existence theorem is of a similar nature.
- **Theorem**: Suppose $F\colon U \to \mathbb{R}$ is continuous, where $U = I_1 \times I_2$ for some intervals $I_1, I_2 \subset \mathbb{R}$, and that $F_y\colon U \to \mathbb{R}$ exists and is continuous.
  - Then for any $(x_0, y_0) \in U$, the IVP $y' = F(x, y); y(x_0) = y_0$ has a solution defined on some interval $I \subset I_1$, and this solution is unique in the sense that for any two solutions $y_1\colon I \to \mathbb{R}$ and $y_2\colon I' \to \mathbb{R}$, $y_1(x) = y_2(x)$ for $x \in I \cap I'$.
  - If moreover $I_2 = \mathbb{R}$ and there exists $B \in \mathbb{R}$ with $F_y(x, y) < B$ for all $(x, y) \in U$, then the solution will exist on the entire interval I_1.
- The first condition fails for Example 10.1.4, since $y \mapsto \sqrt{y}$ is not differentiable at $y = 0$.
  - The second condition fails for Example 10.1.5, since $1 + y^2$ is not bounded as $y \to \pm\infty$.

5 Feb 10, Proof of existence and uniqueness; 10.1: Numerical methods; and 10.3: Linear equations

5.1 End of 10.1: existence and uniqueness

- We will sketch the proof of the existence and uniqueness statement from last time.
- The proof of the second part is fairly easy: the point is that if the slope of $y(x)$ is bounded, then it can only increase by a finite amount in finite time by the **mean value theorem**.
  - The general fact is that if a function $y\colon I \to \mathbb{R}$ satisfies $|y'(x)| < B$ for all $x \in I$, then for any $a < b$ in I, we have $|y(b) - y(a)| < B(b-a)$.
  - Indeed, if we had $|y(b) - y(a)|/(b-a) \geq B$, then the mean value theorem would given $y'(\xi) = |y(b) - y(a)|/(b-a) \geq B$ for some $\xi \in (a,b)$, contradicting the assumption that $y'(x) < B$ for all $x \in I$.
  - One then has to show that the only way that a solution to the IVP can *fail* to extend over the entire interval I is if it diverges to $\pm\infty$.
- The proof of the first part uses an ingenious trick called **Picard iteration**.
  - One first converts the equation to the equivalent **integral equation** $y(x) = y_0 + \int_{x_0}^{x} F(x, y(x))\,\mathrm{d}x$; this integral exists since F is continuous.
  - One then finds a sequence of better and better **approximations** to the solution:
  - The first is simply $f_1(x) = y_0$.
  - Then we inductively define $f_{n+1}(x) = y_0 + \int_{x_0}^{x} F(x, f_n(x))\,\mathrm{d}x$.
  - Finally, we define $y(x) = \lim_{n\to\infty} f_n(x)$.
  - We then have

    $$y(x) = \lim_{n\to\infty} f_n(x) = \lim_{n\to\infty} f_{n+1}(x) = \lim_{n\to\infty} y_0 + \int_{x_0}^{x} F(x, f_n(x))\,\mathrm{d}x = y_0 + \int_{x_0}^{x} F(x, y(x))\,\mathrm{d}x,$$

    as desired.
  - The tricky part is to show that the limit defining y **actually converges** (and that the exchanging of limit and integral in the last equation is legitimate); this is where the assumption is used that F_y exists and is continuous.
- (This is an instance of the general technique of *finding fixed points using iteration*: given a domain U and a continuous function $G\colon U \to U$, if we want to find a *fixed point* of G, i.e., a point $x \in U$ with $G(x) = x$, we can choose some arbitrary $x_0 \in U$, and iteratively define $x_{n+1} = G(x_n)$, and set $x = \lim_{n\to\infty} x_n$ *if this limit exists.* Using the continuity of G, we then have $G(x) = G(\lim_{n\to\infty} x_n) = \lim_{n\to\infty} G(x_n) = \lim_{n\to\infty} x_{n+1} = x$, as desired.)

5.2 10.1B: Numerical Methods

- In scientific applications, it is often important not to explicitly *solve* a given IVP $y' = F(x,y); y(x_0) = y_0$ (which often cannot be done anyway), but to compute numerically an approximation to its solution.
  - This amounts to running a simulation of a quantity y whose rate of change at each time x is given by $F(x,y)$.
  - Geometrically, it amounts to tracing out a curve tangent to a given direction field.

- A straightforward way to do this **Euler's method**:
  - Fix some step size $h > 0$.
  - We are given x_0 and y_0.
  - Now define $x_n = x_0 + n \cdot h$ for $n > 0$, and recursively define $y_{n+1} = y_n + F(x_n, y_n) \cdot h$.
  - It is easy to implement this in any programming language; an example is given in the book.
- There are many ways to improve this algorithm, and there is a whole field dedicated to studying such things.
  - A central problem that arises is the accumulation of **rounding errors**, for example if one tries to make the above step-size parameter h too small.
  - One such improved algorithm is discussed in the book.

5.3 10.3: Linear equations

- A first-order equation $y' = F(x, y)$ is called linear if F is of the form $F(x, y) = -g(x)y + f(x)$.
  - The equation can then be written as $y' + g(x)y = f(x)$; this is called its **normalized form**.
  - The name *linear* corresponds to the fact that $L(y) = y' + g(x)y$ is a *linear* function of y: we have $L(y_1 + y_2) = L(y_1) + L(y_2)$ and $L(a \cdot y) = a \cdot L(y)$ for $a \in \mathbb{R}$ (we will return to this when we do some linear algebra review).
  - Thus, a linear equation has the form $L(y) = f(x)$ for some linear operator L.
  - Just like in linear algebra, we call the equation homogeneous if $f(x) = 0$ and inhomogeneous otherwise.

Example 10.3.1

- In the homogeneous case, if we assume y is never 0, we obtain $\frac{y'}{y} = -g(x)$.
- Integrating, this gives $\ln y = -G(X) + C$ for some $C \in \mathbb{R}$ if $y > 0$ and $\ln(-y) = -G(x)\,\mathrm{d}x + C$ if $y < 0$, where G is some anti-derivative of g, i.e., $\int g(x)\,\mathrm{d}x = G(x) + C$.
- Hence $y = \pm e^C e^{G(x)}$ and so in either case, $y = Ke^{-G(x)} = Ke^{-\int g(x)\,\mathrm{d}x}$ for some $K \neq 0$.
  - As usual, we can now check directly that this is indeed a solution.
  - (And we notice that $K = 0$ also yields a solution.)

6 Feb 12, More linear equations, and linear algebra review

6.1 10.3A: Exponential integrating factors

- In order to show that the solution $y = Ke^{-G(x)}$ is the *general* solution (which we don't quite know since we had to assume $y \neq 0$), we use a trick we have used before, and divide the original equation by the known solution $y = e^{-G(x)}$.
  - We obtain $0 = e^{G(x)}y' + e^{G(X)}g(x) = \frac{\mathrm{d}}{\mathrm{d}x}(e^{G(x)}y)$.
  - Note that this is equivalent to the original equation, since we have multiplied by both sides by a non-vanishing quantity.
  - Thus, we see that for any solution y to this equation, $e^{G(x)y}$ must be constant, and so $y = Ke^{-G(x)}$ for some $K \in \mathbb{R}$, as desired.
- This is the trick of "exponential integrating factors" and it also allows us to solve linear equations in the *inhomogeneous case*.
- For the equation $y' + g(x)y = f(x)$ in normalized form, the exponential integrating factor is $M(x) = e^{\int g(x)\,\mathrm{d}x}$.
- We multiply the equation by M and obtain

$$e^{\int g(x)\,\mathrm{d}x}y' + g(x)e^{\int g(x)\,\mathrm{d}x} = f(x)e^{\int g(x)\,\mathrm{d}x}.$$

- Using the product rule, the left-hand side is $\frac{\mathrm{d}}{\mathrm{d}x}(ye^{\int g(x)\,\mathrm{d}x})$, and the equation becomes

$$\frac{\mathrm{d}}{\mathrm{d}x}(ye^{\int g(x)\,\mathrm{d}x}) = f(x)e^{\int g(x)\,\mathrm{d}x}.$$

- Now we can just integrate the right-hand side and solve for y.
  - Again, the original equation is equivalent to the equation after multiplying by M since M is nowhere zero.

Example 10.3.2

- Let's solve $y' = xy + x$.
- We rewrite it in normalized form $y' - xy = x$.
- Multiply by the integrating factor $y'e^{-x^2/2} - xye^{-x^2/2} = xe^{-x^2/2}$.
- Now integrate $ye^{-x^2/2} = -e^{-x^2/2} + C$.
- Thus $y = -1 + Ce^{-x^2/2}$.
- Since each step was an equivalence, this is the *general* solution to the equation (with domain $\mathbb{R}$).

6.2 10.3B Applications

Example 10.3.5

- Newton's law of cooling says the surface temperature $u(t)$ of an objects changes at a rate proportional to the difference to the ambient temperature $f(t)$ (which we are assuming might also vary with time).
  - Thus $u' = k(f - u)$ for some $k > 0$.

  - k should be positive so that $u' < 0$ precisely when $f < u$.

- It is supposedly quite accurate in certain situations, but as with all empirical laws, it has its limitations.
  - Ideally, in addition to testing the law with experiments, one should develop a *mechanism* (for example, in terms of the jiggling of molecules) which explain why it should hold; this makes it easier to understand under which circumstances it should and shouldn't apply.

- Writing this as $u' + ku = kf(t)$, this is just a special case of a linear first-order equation in which the coefficient of u is constant.

- We obtain $u'e^{kt} + ue^{kt} = kf(x)e^{kt}$ and hence $ue^{kt} = \int kf(x)e^{kt}\,\mathrm{d}t + C$ and hence $u = ke^{-kt}\int f(t)e^{kt}\,\mathrm{d}t + Ce^{-kt}$.
  - We fix the lower-bound of the integral to $t = 0$, giving $C = u(0)$ and hence $u = ke^{-kt}\int_0^t f(s)e^{ks}\,\mathrm{d}s + u(0)e^{-kt}$.

- If the ambient temperature is constant, $f(t) = f_0$, we obtain $u = kf_0(1 - e^{-kt}) + u(0)e^{-kt} = kf_0 + (u(0) - kf_0)e^{-kt}$.
  - Thus $u(t) \xrightarrow{t\to\infty} f_0$, as we would expect.

6.3 Abstract vector spaces

- We review the concept of an abstract vector space.

- There are both real and complex vector spaces. In order to handle both at the same time, in what follows, we let K stand for one of $\mathbb{R}$ or $\mathbb{C}$.
  - (Actually, the notion of vector space makes sense when K is any *field*.)

- The motivating examples of vector spaces are the familiar ones K^n; the abstract definition results from isolating the most important features and properties from these. One should keep this example in mind when considering the general definition, but as we'll see, there are lots of other interesting examples as well.

- **Definition**: a vector space over K or K-vector space is a triple $(V, +, \times)$ where V is an arbitrary *set* (whose elements we call *vectors*), $+\colon V \times V \to V$ is a binary operation on V (called "*addition*") taking any two elements $\mathbf{u}, \mathbf{v} \in V$ to an element $\mathbf{u} + \mathbf{v} \in V$, and $\times\colon K \times V \to V$ is an operation (called "*scalar multiplication*") taking an element $r \in K$ and an element $\mathbf{v} \in V$ to an element $r \times \mathbf{v} \in V$ (which we also write as $r \cdot \mathbf{v}$ or just $r\mathbf{v} \in V$).
  - These are required to satisfy the following *axioms*:
  - Addition is associative and commutative, and has an identity element $\mathbf{0} \in V$ (i.e., $\mathbf{v} + \mathbf{0} = \mathbf{v}$ for all $\mathbf{v} \in V$). (*Exercise*: it follows that there is a *unique* such element, which we call the "*zero vector*" of V.)
  - Each $\mathbf{v} \in V$ has an additive inverse $(-\mathbf{v})$ (i.e., $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$). (*Exercise*: it follows that for each $\mathbf{v}$, there is a *unique* such element $-\mathbf{v}$.)
  - Scalar multiplication distributes over vector addition and scalar addition, i.e., $r \cdot (\mathbf{u} + \mathbf{v}) = r \cdot \mathbf{u} + r \cdot \mathbf{v}$ and $(r + s) \cdot \mathbf{v} = r \cdot \mathbf{v} + s \cdot \mathbf{v}$ for all $r, s \in K$ and $\mathbf{u}, \mathbf{v} \in V$.
  - $1 \cdot \mathbf{v} = \mathbf{v}$ for all $\mathbf{v} \in V$.
  - $(r \cdot s) \cdot \mathbf{v} = r \cdot (s \cdot \mathbf{v})$ for all $r, s \in K$ and $\mathbf{v} \in V$.

- As is typical, we will often *abuse notation* and simply write V in place of $(V, +, \times)$.
  - For example, we may speak of "the vector space $\mathbb{R}^3$", whereas the vector space is really the triple $(\mathbb{R}^3, +, \times)$.
- Examples:
  - The original and most important example of a K-vector space is the set K^n of n-tuples of elements of K. Addition and multiplication are given component-wise, i.e., $(u_1, \ldots, u_n) + (v_1, \ldots, v_n) = (u_1 + v_1, \ldots, u_n + v_n)$, and $r \cdot (v_1, \ldots, v_n) = (rv_1, \ldots, rv_n)$ for $r \in K$.
    - (One should check that these operations indeed satisfy the axioms!)
    - Of course, we usually picture $\mathbb{R}^2$ as the set of points in a plane with a fixed pair of coordinate axes and chosen units of length, and likewise $\mathbb{R}^3$ is pictured as 3-dimensional space (and $\mathbb{R}^1$ as a line).
    - (However, when considering them as vector spaces, we should really regard $\mathbb{R}^2$, for instance, as the set of *arrows* in the given plane, where we identify two arrows if one can be brought on top of the other by a translation; the reason is that it doesn't make sense to add or scale *points*, but we do know how to add *arrows* (in the usual head-to-tail manner) and scale them. The identification between arrows and points of course follows from placing the tail an arrow at the origin and passing to the point at its head.)
  - A related example is the set $K^{m \times n}$ of $m \times n$-matrices. Formally, this is the set of functions

$$\{(i, j) \mid 1 \leq i \leq m; 1 \leq j \leq n\} \to K.$$

    - As with tuples, given a matrix A, we write A_{ij} in place of $A(i, j)$, and we represent a matrix in the familiar way as a box of numbers:

$$A = \begin{bmatrix} A_{11} & \cdots & A_{1n} \\ \vdots & & \vdots \\ A_{m1} & \cdots & A_{mn} \end{bmatrix}.$$

    - We also write $A = (A_{ij})_{1 \leq i \leq m, j \leq 1 \leq n}$ or just $A = (A_{ij})$.
    - Addition and scalar multiplication are again given component-wise: $(A + B)_{ij} = A_{ij} + B_{ij}$ and $(r \cdot A)_{ij} = r \cdot A_{ij}$ for $r \in K$.
    - Of special interest are the vector spaces $K^{n,1}$ of column vectors and $K^{1,n}$ of row vectors of length n. Of course, there are bijections $K^n \cong K^{n,1} \cong K^{1,n}$, and we may sometimes *abuse notation* and simply write K^n in place of $K^{n,1}$.
  - The above two examples generalize: given any set S, the set K^S of all functions $S \to K$ is a vector space with operations $(f + g)(s) = f(s) + g(s)$ and $(r \cdot f)(s) = r \cdot f(s)$.
    - In particular, the set $\mathbb{R}^{\mathbb{R}}$ of all functions $\mathbb{R} \to \mathbb{R}$ becomes a vector space in this way.

7 Feb 17, More linear algebra

7.1 Subspaces

- A linear subspace (or just subspace) of a vector space V is a subset $W \subset V$ which is closed under addition and scalar multiplication:
  - This means that $\mathbf{u}, \mathbf{v} \in W$ implies $\mathbf{u} + \mathbf{v} \in W$ and $r\mathbf{v} \in W$ for all $r \in K$.
  - In this case, W is again a vector space, with the same operations $+$ and $\times$.
- The most familiar examples are the subspaces of $\mathbb{R}^3$, which are the planes and lines passing through the origin, as well as $\mathbb{R}^3$ itself and the singleton set $\{\mathbf{0}\} \subset \mathbb{R}^3$.
- Other interesting examples of vector spaces arise as subspaces of $\mathbb{R}^{\mathbb{R}}$ (or more generally of $\mathbb{R}^I$ for any domain $I \subset \mathbb{R}$):
  - The set of *continuous* functions $\mathcal{C}^0(I)$ and the set of *smooth* functions $\mathcal{C}^\infty(I)$ are each subspaces of $\mathbb{R}^{\mathbb{R}}$ (as is $\mathcal{C}^k(I)$ for any k).
  - The set of *polynomial functions* $p(x) = \sum_{i=0}^n a_i x^i$ (with $n \in \mathbb{N}$ and $a_i \in \mathbb{R}$) is a subspace of $\mathbb{R}^{\mathbb{R}}$, and is usually denoted $\mathbb{R}[x]$.
  - There is a further subspace $\mathbb{R}_{\leq d}[x]$ consisting of polynomials of degree at most d.
  - Similarly, one can consider polynomials $\mathbb{C}[x]$ (and $\mathbb{C}_{\leq d}[x]$) over $\mathbb{C}$.
  - One should contemplate why these are all in fact subspaces!

7.2 Basis, dimension, etc.

- Given vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ in a vector space V, a linear combination of these vectors is any vector of the form $\sum_{i=1}^n a_i \mathbf{v}_i = a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n$ with $a_1, \ldots, a_n \in K$.
- The span $\mathrm{Span}(S)$ of a subset $S \subset V$ is the set of all linear combinations $\sum_{i=1}^n a_i \mathbf{v}_i$ with $\mathbf{v}_1, \ldots, \mathbf{v}_n \in S$.
- We say that S spans V or is a spanning set if $\mathrm{Span}(S) = V$, i.e., if every element of V is a linear combination of elements in S.
- We say that S is linearly independent if, whenever $\mathbf{v}_1, \ldots, \mathbf{v}_n \in S$ are distinct elements and $a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n = \mathbf{0}$, then $a_1 = \cdots = a_n = 0$
  - Equivalently, each element of V can be represented *in at most one way* as a linear combination of elements of S.
  - Equivalently, no element of S is a linear combination of the other elements.
- A basis for V is a subset $\mathcal{B}$ which is both spanning and linearly independent (or equivalently, such that every vector in V can be represented in *exactly one way* as a linear combination of elements of S).
  - K^n has the standard basis $\mathbf{e}_1, \ldots, \mathbf{e}_n$, where $\mathbf{e}_i$ is the vector $\mathbf{e}_i = (0, \ldots, 1, \ldots, 0)$ with a single 1 in the i-th place and 0 elsewhere.
  - It is often useful to regard a basis not as a *set* $\mathcal{B} \subset V$, but as a *tuple* $B \in V^n$; in other words, we remember the *order* of the elements in B, and in this case refer to B as an ordered basis.
- If V has a basis $\mathcal{B}$ which is *finite*, then it is called finite-dimensional; otherwise, it is called infinite-dimensional.

- **Theorem**: if V is finite-dimensional, then any two bases of V have the *same* number of elements, and this number is called the *dimension* $\dim(V)$ of V; moreover, any *linearly independent* set of vectors in V has at most $\dim V$ elements.
  - Roughly speaking, the proof proceeds by taking an arbitrary linearly independent set $S = \{\mathbf{u}_1, \ldots, \mathbf{v}_k\} \subset V$ and an arbitrary spanning set $T = \{\mathbf{v}_1, \ldots, \mathbf{v}_l\} \subset V$, and swapping out vectors in T for vectors in S one by one until a new spanning set is produced which includes all the vectors in S. This proves that $k \leq l$ for any such sets S and T (and in particular that $k = l$ if S and T are both bases).
- Examples of finite-dimensional vector spaces:

V	K^n	$K^{m\times n}$	$K_{\leq d}[x]$
$\dim(V)$	n	$m \cdot n$	$d+1$

  - Examples of infinite-dimensional vector spaces: $\mathbb{R}^{\mathbb{R}}, \mathcal{C}^0(\mathbb{R}), \mathcal{C}^\infty(\mathbb{R}), \mathbb{R}[x]$.
- **Lemma**: given any set $\mathbf{v}_1, \ldots, \mathbf{v}_k$ of linearly independent vectors in a vector space V, and a vector $\mathbf{v}_{k+1} \notin \mathrm{Span}\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$, the set $\{\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}\}$ is still linearly independent.
- **Corollary** (of the theorem and the lemma): if V is finite-dimensional, then any subspace $W \subset V$ is finite-dimensional, and if $\dim W = \dim V$, then $W = V$.
  - Indeed, if W were infinite-dimensional, we would have more than $\dim V$ linearly independent vectors in V.
  - And if $\dim W = \dim V$ and $W \subsetneq V$, then by the Lemma, we could extend a basis of W to a larger linearly independent set of vectors, which would thus have more than $\dim V$ elements.
- Another **Corollary**: any linearly independent set of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k \in V$ in a finite-dimensional space can be extended to a basis $\mathbf{v}_1 \ldots, \mathbf{v}_n$ of V.
  - Indeed, as long as $k < n$, we must have $W = \mathrm{Span}\{\mathbf{v}_1, \ldots, \mathbf{v}_k\} \subsetneq V$, hence we can find some v_{k+1} in $V - W$, and the set $\{\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}\}$ is still linearly independent; eventually, we will have $k = n$, and then $\mathbf{v}_1, \ldots, \mathbf{v}_n$ will have to constitute a basis.

8 Feb 19, Linear maps

8.1 Linear maps

- A function $f\colon V \to W$ between two vector spaces is linear if it preserves addition and scalar multiplication, i.e., $f(\mathbf{v}+\mathbf{w}) = f(\mathbf{v}) + f(\mathbf{w})$ and $f(r\mathbf{w}) = r \cdot f(\mathbf{w})$ for all $r \in K$.
  - It follows (by induction) that f preserves arbitrary linear combinations, i.e., that $f(a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n) = a_1(\mathbf{v}_1) + \cdots + a_n f(\mathbf{v}_n)$ for any $\mathbf{v}_1, \ldots, \mathbf{v}_n \in V$ and $a_1, \ldots, a_n \in K$.
  - A linear map $f\colon V \to V$ from a vector space to itself is sometimes called a linear operator.
- Given a basis $\mathbf{b}_1, \ldots, \mathbf{b}_n$ of V, any linear map $f\colon V \to W$ is determined by its values $f(\mathbf{b}_i)$ on the basis elements.
  - And conversely, given any vectors $\mathbf{w}_1, \ldots, \mathbf{w}_n$, there is a unique linear map $f\colon V \to W$ with $f(\mathbf{b}_i) = \mathbf{w}_i$ for all i; it is given by $f\left(\sum_{i=1}^n v_i\mathbf{b}_i\right) = \sum_{i=1}^n v_i\mathbf{w}_i$.
- *Exercise*: the linear maps $f\colon K^m \to K^n$ are exactly those of the form

$$f\begin{bmatrix} v_1 \\ \vdots \\ v_m \end{bmatrix} = \begin{bmatrix} a_{11}v_1 + \cdots + a_{1n}v_n \\ \vdots \\ a_{m1}v_1 + \cdots + a_{mn}v_n \end{bmatrix}$$

  for some elements $a_{ij} \in K$.
  - In other words, we have $f(\mathbf{v}) = A \cdot \mathbf{v}$, where $A = (a_{ij})$.
  - Here, we are using matrix multiplication, which we recall is the operation $K^{l\times m} \times K^{m\times n} \to K^{l\times n}$ given by $(A \cdot B)_{ij} = \sum_{k=1}^m A_{ik} \cdot B_{kj}$.
- For any two vector spaces V and W, we write $\mathcal{L}(V, W)$ for the set of all linear maps $V \to W$; this is a subspace of the vector space W^V.
  - In the case $V = W$, we simply write $\mathcal{L}(V)$ for the space $\mathcal{L}(V, V)$ of *linear operators* on V.
- Another important example of a linear map is the derivative operator $D = \frac{\mathrm{d}}{\mathrm{d}x}\colon \mathcal{C}^\infty(I) \to \mathcal{C}^\infty(I)$; this can also be considered as a linear map $\mathbb{R}_{\leq d}[x] \to \mathbb{R}_{\leq d-1}[x]$ or for any $d > 0$.
  - Similarly, we have seen the first-order *linear differential operators* $L\colon \mathcal{C}^\infty(I) \to \mathcal{C}^\infty(I)$ given by $L(y) = Dy + g(x)y = y' + g(x)y$ for some $g\colon I \to \mathbb{R}$.
  - There is also the integral map $\int_0^x\colon \mathcal{C}^0(\mathbb{R}) \to \mathcal{C}^0(\mathbb{R})$.
- The kernel or nullspace of a linear map $f\colon V \to W$ is the subspace $\ker(f) \subset V$ defined by $\ker(f) = \{\mathbf{v} \in V \mid f(\mathbf{v}) = \mathbf{0}\}$.
  - *Exercise*: a linear map is injective if and only if its nullspace is $\{\mathbf{0}\}$.
- The image $\operatorname{im}(f)$ of a linear map $f\colon V \to W$ (or in fact of any function between two sets) is the set $\operatorname{im}(f) = \{f(\mathbf{v}) \mid \mathbf{v} \in V\}$.
- **Theorem** (the *rank-nullity theorem* or *dim sum theorem*): if V and W are finite-dimensional and $f\colon V \to W$ is linear, then $\dim(V) = \dim\ker(f) + \dim\operatorname{im}(f)$.
  - In particular, if $\dim V = \dim W$, then f is injective if and only if it is surjective.
- An isomorphism between vector spaces V and W is a linear map $V \to W$ which is also a bijection.

  - In this case, the inverse map $W \to V$ is also linear (hence also an isomorphism).
  - If there exists an isomorphism between V and W, we say that V and W are isomorphic, and write $V \cong W$.

- Note that isomorphism is an *equivalence relation*: it is reflexive ($V \cong V$), symmetric ($V \cong W \Rightarrow W \cong V$), and transitive $U \cong V \wedge V \cong W \Rightarrow U \cong W$.

- **Theorem**: V is of finite-dimension n if and only if $V \cong K^n$.

  - Indeed, if V has a basis $\mathbf{b}_1, \dots, \mathbf{b}_n$, then each vector $\mathbf{v} \in V$ has *coordinates* $v_1, \dots, v_n$ with respect to this basis, determined by $\mathbf{v} = v_1\mathbf{b}_1 + \cdots v_n\mathbf{b}_n$; the map $V \to K^n$ taking $\mathbf{v}$ to $(v_1, \dots, v_n)$ is then the desired isomorphism.
  - An important consequence of this is that any linear map $f\colon V \to W$ between finite-dimensional vector spaces (say, with $\dim V = m$ and $\dim W = n$) can be represented by a $n \times m$ matrix, because $V \cong K^m$ and $W \cong K^n$.
    * The matrix representation of f will depend on a choice of isomorphisms $V \xrightarrow{\sim} K^m$ and $W \xrightarrow{\sim} K^n$, which is to say on a choice of ordered bases of V and W.
    * Given such ordered bases $\mathcal{B} = (\mathbf{v}_1, \dots, \mathbf{v}_m)$ and $\mathcal{C} = (\mathbf{w}_1, \dots, \mathbf{w}_n)$, the resulting matrix $A \in K^{n\times m}$ can be described directly by the formula $f(\mathbf{v}_i) = \sum_{j=1}^{n} A_{ji}\mathbf{w}_j$ (i.e., the i-th column of A consists of the coordinates of $f(\mathbf{v}_i)$ with respect to $\mathcal{C}$).
    * Another way to say this is: if we write $\mathbf{v}_\mathcal{B} \in K^{n\times 1}$ for the coordinate vector of $\mathbf{v} \in V$ with respect to the ordered basis $\mathcal{B}$, and similarly write $\mathbf{w}_\mathcal{C} \in K^{m\times 1}$ for the coordinates of $\mathbf{w} \in W$ with respect to $\mathcal{C}$, then we have $\big(f(\mathbf{v})\big)_\mathcal{C} = A \cdot \mathbf{v}_\mathcal{B} \in K^{m\times 1}$ for all $\mathbf{v} \in V$.

- Other examples of isomorphisms:

  - If $\dim V = m$ and $\dim W = n$, then $\mathcal{L}(V, W) \cong K^{n\times m}$ and in particular $\mathcal{L}(V) \cong K^{m\times m}$.
  - The existence of these isomorphisms follows from the above theorem just by comparing dimensions, but one can also establish the isomorphisms more directly by sending each linear map to the matrix representing it.
    * But, again, it is important to remember that this isomorphism will depend on choosing bases for V and W.
  - *Important exercise:* The above isomorphism takes *composition of linear maps* to *matrix multiplication*: given finite-dimensional spaces U, V, W with chosen ordered bases, and linear maps $f\colon U \to V$ and $g\colon V \to W$ represented by matrices $A \in K^{m\times l}$ and $B \in K^{n\times m}$, the composite map $g \circ f\colon U \to W$ is represented by the product $B \cdot A$.

9 Feb 24, 3.7: Inner products

9.1 3.7A General properties of inner products

- We introduce the notion of an *inner product* on a general vector space, which is a scalar-valued binary function $V \times V \to \mathbb{R}$ satisfying certain axioms.
  - While the concept of *abstract vector space* allows us to reproduce many of the geometric features of the motivating example $\mathbb{R}^n$, it does not allow us to define the notions of *length* or *angle*, which are of course crucial to studying geometry in $\mathbb{R}^2$ and $\mathbb{R}^3$; it is inner products that allows us to reproduce these notions in an arbitrary abstract vector space.
  - Just as the very notion of vector space is meant to capture the main features of the archetypical case $\mathbb{R}^n$, so the notion of inner product is meant to capture the main features of the dot product $\mathbf{x} \cdot \mathbf{y} = \sum_{i=1}^n x_i y_i$ on $\mathbb{R}^n$; again, this motivating example should be kept in mind while considering the definition, but it's also important to remember that there are other interesting examples as well.
- If U, V, W are K-vector spaces, a function $f\colon U \times V \to W$ is bilinear if it is linear in each argument, when the other argument is held fixed, i.e., $f(a\mathbf{u} + \mathbf{u}', \mathbf{v}) = af(\mathbf{u}, \mathbf{v}) + f(\mathbf{u}', \mathbf{v})$ and $f(\mathbf{u}, a\mathbf{v} + \mathbf{v}') = af(\mathbf{u}, \mathbf{v}) + f(\mathbf{u}, \mathbf{v}')$ for all $\mathbf{u} \in U$, $\mathbf{v} \in V$, and $a \in \mathbb{R}$..
  - (Examples: the cross-product $\times\colon \mathbb{R}^3 \times \mathbb{R}^3 \to \mathbb{R}^3$, matrix multiplication $K^{a,b} \times K^{b,c} \to K^{a,c}$, composition of linear maps $\mathcal{L}(U, V) \times \mathcal{L}(V, W) \to \mathcal{L}(U, W)$.)
- A bilinear form on a K-vector space V is a bilinear function $\beta\colon V \times V \to K$.
- An inner product on a real vector space V is a bilinear form $\beta\colon V \times V \to \mathbb{R}$ which is moreover
  - symmetric, meaning $\langle \mathbf{u}, \mathbf{v} \rangle = \langle \mathbf{v}, \mathbf{u} \rangle$ for all $\mathbf{u}, \mathbf{v} \in V$
  - positive-definite, meaning that $\langle \mathbf{u}, \mathbf{u} \rangle > 0$ for all $\mathbf{u} \neq \mathbf{0}$
- An inner product space is a pair (V, β), where V is a real vector space, and $\beta\colon V \times V \to \mathbb{R}$ is an inner product on V.
  - (Again, we often abuse notation, and just write say that "V is a inner product space").
  - (In particular, by "the inner product space $\mathbb{R}^n$", we mean $\mathbb{R}^n$ equipped with the standard inner product.)
- **Notation**: for a given inner product $\beta\colon V \times V \to \mathbb{R}$, we will usually prefer to write $\langle \mathbf{u}, \mathbf{v} \rangle$ in place of $\beta(\mathbf{u}, \mathbf{v})$.
  - Accordingly, we may forego the additional name β, and instead, inserting placeholders "$-$" for $\mathbf{u}$ and $\mathbf{v}$ in the expression "$\langle \mathbf{u}, \mathbf{v} \rangle$", just refer to the inner product as $\langle -, - \rangle\colon V \times V \to \mathbb{R}$.
- Remarks:
  - In light of the symmetry property, the bilinearity is equivalent to just being linear in one argument.
  - It follows from bilinearity that $\langle \mathbf{0}, \mathbf{0} \rangle = 0$, and hence (by positive-definiteness) that $\langle \mathbf{v}, \mathbf{v} \rangle \geq 0$ for all $\mathbf{v}$.
- Example: the bilinear forms on $\mathbb{R}^n$ are exactly the functions of the form $\langle \mathbf{u}, \mathbf{v} \rangle = \sum_{i,j=1}^n a_{ij} u_i v_j$ for some real numbers a_{ij}.

- In other words, $\langle u, v\rangle = \mathbf{u}^\top A\mathbf{v}$ where $A = (a_{ij})_{i,j=1}^n$, and where we are identifying $\mathbb{R}^{1\times 1}$ with $\mathbb{R}$.
- Recall here that for a matrix $B \in K^{l\times m}$, we write $B^\top \in K^{m\times l}$ for its transpose defined by $(B^\top)_{ij} = B_{ji}$.
- In particular, if we take A to be the identity matrix I_n (defined by $(\mathrm{I}_n)_{ij} = \delta_{ij}$), we obtain the dot product $\mathbf{u}\cdot\mathbf{v}$, also known as the standard inner product on $\mathbb{R}^n$.
  * The symbol δ_{ij}, called the Kronecker delta symbol, is defined to be 1 if $i = j$ and 0 otherwise
  * Often we will just write I instead of I_n when n is clear from context.

- It is easy to see that such a bilinear form $\mathbf{u}^\top A\mathbf{v}$ is symmetric precisely when the *matrix A is* symmetric, meaning $A^\top = A$.

- Example 3.7.1
  - It is more difficult to determine whether the bilinear form arising from a matrix A is *positive-definite*.
  - But here is an example: if A is a diagonal matrix (meaning $A_{ij} = 0$ for $i \neq j$) with *positive* diagonal entries, then $\mathbf{u}^\top A\mathbf{v}$ is positive definite (and hence an inner product, since A is clearly symmetric).
  - In this case $\langle \mathbf{u}, \mathbf{v}\rangle$ is given by the simple formula $a_{11}u_1v_1 + \cdots + a_{nn}u_nv_n$.
  - For a completely specific example, $\langle \mathbf{u}, \mathbf{v}\rangle = u_1u_2 + 2v_1v_2$ is an inner product on $\mathbb{R}^2$.

Complex inner products

- There is also a notion of inner products over the complex numbers, the standard example being the function $\mathbb{C}^n \times \mathbb{C}^n \to \mathbb{C}$ given by $(\mathbf{w}, \mathbf{z}) \mapsto \overline{\mathbf{w}}\cdot\mathbf{z}$.
  - Here, we recall that for a complex number $z = a + bi$, its complex conjugate is defined by $\bar{z} = a - bi$ (reflection across the real axis).
  - This has the fundamental property $z\cdot\bar{z} = |z|^2$, where we recall that the absolute value (or *modulus*) of a complex number is defined by $|z| = \sqrt{a^2+b^2}$ (distance from the origin in the complex plane).
  - Another fundamental property of the complex conjugate is that it is compatible with addition and multiplication: $\overline{z+w} = \bar{z} + \bar{w}$ and $\overline{z\cdot w} = \bar{z}\cdot\bar{w}$.
  - For a vector $\mathbf{w} \in \mathbb{C}^n$, the complex conjugate is defined component-wise: $\overline{\mathbf{w}} = (\overline{w}_1, \ldots, \overline{w}_n)$.
  - We recall the important formulas $\operatorname{Re} z = \frac{1}{2}(z + \bar{z})$ and $\operatorname{Im} z = \frac{1}{2}(z - \bar{z})$, where Re and Im are the real and imaginary parts, defined by $\operatorname{Re}(a+bi) = a$ and $\operatorname{Im}(a+bi) = b$.

- In general, on a complex vector space V, we define a function $\langle -, -\rangle \colon V \times V \to \mathbb{C}$ to be:
  - sesquilinear if it is linear in its second argument and *antilinear* in its first argument, meaning $\langle a\mathbf{u} + \mathbf{v}, \mathbf{w}\rangle = \bar{a}\langle \mathbf{u}, \mathbf{w}\rangle + \langle \mathbf{v}, \mathbf{w}\rangle$ and $\langle \mathbf{w}, a\mathbf{u} + \mathbf{v}\rangle = a\langle \mathbf{w}, \mathbf{u}\rangle + \langle \mathbf{w}, \mathbf{v}\rangle$ for all $\mathbf{u}, \mathbf{v}, \mathbf{w} \in V$ and $a \in \mathbb{C}$.
  - conjugate symmetric if $\langle \mathbf{u}, \mathbf{v}\rangle = \overline{\langle \mathbf{v}, \mathbf{u}\rangle}$ for all $\mathbf{u}, \mathbf{v} \in V$.
  - an inner product if it is sesquilinear, conjugate symmetric, and positive definite (which as before means $\langle \mathbf{v}, \mathbf{v}\rangle > 0$ for all $\mathbf{v} \neq \mathbf{0}$).
  - Again, a complex inner product space is a complex vector space equipped with an inner product.

- Note that if V is a *real* vector space, and $\langle -, - \rangle \colon V \times V \to \mathbb{R}$ is a *real*-valued function, then the definitions of sesquilinear and conjugate symmetric reduce to those of "bilinear" and "symmetric" (since $\bar{a} = a$ for all $a \in \mathbb{R}$).
  - Thus, for conciseness, we can define an inner product on a real *or* complex vector V all at once by saying it is a sesquilinear, conjugate symmetric, positive definite function $V \times V \to K$.

Norms

- Now let V be a real or complex inner product space.
- We define the norm or length of a vector $\|\mathbf{v}\|$ to be the non-negative real number $\|\mathbf{v}\| = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}$.
  - Note that $\|\mathbf{v}\| = 0 \iff \mathbf{v} = \mathbf{0}$ by positive-definiteness of $\langle -, - \rangle$.
  - Some people just write $|\mathbf{v}|$ for the norm instead of $\|\mathbf{v}\|$.
- The norm has the following properties:
  - ("Positivity") $\|\mathbf{v}\| > 0$ for all $\mathbf{v} \neq \mathbf{0}$
  - ("Linearity") $\|a\mathbf{v}\| = |a| \cdot \|\mathbf{v}\|$ for all $a \in K$ and $\mathbf{v} \in V$.
  - ("Triangle inequality") $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$
  - (In general, one calls a function $V \to \mathbb{R}$ having these properties a norm on V; there are examples of norms that do not arise from inner products, an example being $\|-\|_3 \colon \mathbb{R}^n \to \mathbb{R}$ given by $\|\mathbf{v}\|_3 = (\sum_{i=1}^n v_i^3)^{1/3}$.)
- We have the important formula (in the real case): $\|u + v\|^2 = \|u\| + \|v\| + 2\langle u, v \rangle$.
  - In the complex case, the formula is $\|u + v\|^2 = \|u\| + \|v\| + 2\,\mathrm{Re}\langle u, v \rangle$
  - These formulas show that the inner product can be recovered from the norm, namely (in the real case) as $\langle u, v \rangle = \frac{1}{2}(\|u + v\| - \|u\| - \|v\|)$.
- **Theorem (the Cauchy-Schwartz inequality)**: $|\langle \mathbf{u}, \mathbf{v} \rangle| \leq \|\mathbf{u}\| \cdot \|\mathbf{v}\|$ for all $\mathbf{u}, \mathbf{v} \in V$, where equality holds if and only if $\mathbf{u}$ and $\mathbf{v}$ are colinear (i.e., linearly dependent).
- **Proof**:
  - We give the proof in the real case; the complex case is very similar but a little bit more complicated.
  - The case in which $\mathbf{u} = \mathbf{0}$ or $\mathbf{v} = \mathbf{0}$ is obvious, since then both sides are equal to 0 (and $\mathbf{u}, \mathbf{v}$ are colinear in this case).
  - Next, assume $\mathbf{u}, \mathbf{v} \neq 0$. Using bilinearity of the inner product, and the linearity of the norm, we may divide both sides by $\|\mathbf{u}\| \cdot \|\mathbf{v}\|$ to obtain the equivalent inequality $|\langle \mathbf{a}, \mathbf{b} \rangle| \leq \|\mathbf{a}\| \cdot \|\mathbf{b}\|$, where $\mathbf{a} = \mathbf{u}/\|\mathbf{u}\|$ and $\mathbf{b} = \mathbf{v}/\|\mathbf{v}\|$.
  - We observe that $\|\mathbf{a}\| = \|\mathbf{b}\| = 1$; moreover, $\mathbf{a}, \mathbf{b}$ are colinear if and only if $\mathbf{u}, \mathbf{v}$ are. In other words, we have reduced to the case of proving the inequality for unit vectors.
  - In this case, we have $0 \leq \|\mathbf{a} \pm \mathbf{b}\|^2 = \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 \pm 2\langle \mathbf{a}, \mathbf{b} \rangle = 2 \pm 2\langle \mathbf{a}, \mathbf{b} \rangle$, and hence $\pm\langle \mathbf{a}, \mathbf{b} \rangle \leq 1 = \|\mathbf{a}\| \cdot \|\mathbf{b}\|$, as desired.
  - Note that the inequality above is an equality if and only if $\mathbf{a} \pm \mathbf{b} = 0$.
- Example 3.7.1 again

- Note that for the inner product $\langle \mathbf{u}, \mathbf{v} \rangle = u_1 v_1 + 2u_2 v_2$ on $\mathbb{R}^2$, the resulting geometry in the plane is "distorted". For example, $(1, 0)$ is still a unit vector, but $(0, 1)$ now has length $\sqrt{0^2 + 2 \cdot 1^2} = \sqrt{2}$.
- The "unit circle" (the set of all vectors of length 1) is now an ellipse.

Angles

- We continue to fix a real or complex inner product space V.
- For $\mathbf{u}, \mathbf{v} \neq \mathbf{0}$, the Cauchy-Schwartz inequality can be rephrased as $-1 \leq \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\| \|\mathbf{v}\|} \leq 1$ for any $\mathbf{u}, \mathbf{v} \in V$.
- It follows that there is a unique $\theta \in [0, \pi]$ with $\cos \theta = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\| \|\mathbf{v}\|}$.
- We define the angle $\angle(\mathbf{u}, \mathbf{v})$ between $\mathbf{u}$ and $\mathbf{v}$ to be this number θ.
  - This recovers the usual notion of angle in $\mathbb{R}^2$ and $\mathbb{R}^3$ because of the well-known formula $\mathbf{u} \cdot \mathbf{v} = \|\mathbf{u}\| \|\mathbf{v}\| \cos \theta$.
  - $\mathbf{u}$ and $\mathbf{v}$ are called orthogonal, written $\mathbf{u} \perp \mathbf{v}$, if $\theta = \pi/2$, or in other words if $\mathbf{u} \cdot \mathbf{v} = 0$.
  - This definition still makes sense when $\mathbf{u} = \mathbf{0}$ or $\mathbf{v} = \mathbf{0}$.
- We have the following version of the **Pythagorean theorem**:
  - If $\mathbf{u} \perp \mathbf{v}$, then $\|\mathbf{u} + \mathbf{v}\|^2 = \|\mathbf{u}\|^2 + \|\mathbf{v}\|^2$.
- Example 3.7.1 again
  - Note again that the non-standard inner product $\langle \mathbf{u}, \mathbf{v} \rangle = u_1 v_1 + 2u_2 v_2$ on $\mathbb{R}^2$ "distorts" the notion of angle.
  - For example, $(1, 0)$ and $(0, 1)$ are still orthogonal, but $(1, 1)$ and $(1, -1)$ are no longer orthogonal, since $\langle (1, 1), (1, -1) \rangle = 1 - 2 \neq 0$.

Example 3.7.2

- On the vector space $\mathcal{C}^0([0, 2\pi])$ of continuous real-valued functions on the interval $[0, 2\pi]$, we can define an inner product by the formula

$$\langle f, g \rangle = \int_0^{2\pi} f(x) g(x) \, \mathrm{d}x.$$

  - One should check that this is indeed an inner product.
- This inner product is of central importance in *Fourier analysis*, the mathematical subject underpinning of signal processing.
  - In this subject, it is also important to consider continuous complex-valued functions; in this case, the inner product can be defined by $\langle f, g \rangle = \int_0^{2\pi} \overline{f(x)} g(x) \, \mathrm{d}x$.

10 Feb 26, 3.7B: Orthogonal bases; dual spaces

10.1 3.7B Orthogonal bases

- A set $S \subset V$ of vectors in a real or complex inner product space V is *orthonormal* if $\langle \mathbf{u}, \mathbf{u} \rangle = 1$ and $\langle \mathbf{u}, \mathbf{v} \rangle = 0$ for all $\mathbf{u}, \mathbf{v} \in V$ with $\mathbf{u} \neq \mathbf{v}$; i.e., all the vectors are unit vectors, and they are mutually orthogonal.
  - Any orthonormal set is linearly independent; indeed, if $a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n = \mathbf{0}$ where $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \delta_{ij}$ for all i, j, then using linearity of $\langle -, - \rangle$ and orthonormality, $0 = \langle \mathbf{v}_i, \mathbf{0} \rangle = \langle \mathbf{v}_i, a_1\mathbf{v}_1 + \cdots + a_n\mathbf{v}_n \rangle = a_i$ for all i.
- An orthonormal basis is a basis which is also an orthonormal set.
- The standard example of an orthonormal basis is the standard basis of $\mathbb{R}^n$ or $\mathbb{C}^n$ (equipped with the standard inner product).
- In general, a basis $\mathbf{b}_1, \ldots, \mathbf{b}_n$ of $\mathbb{R}^n$ is orthonormal if and only if the matrix B with columns $\mathbf{b}_1, \ldots, \mathbf{b}_n$ is an orthogonal matrix, meaning that $B^\top \cdot B = \mathrm{I}$.

We have the following nice formula for the representation of any vector with respect to an orthonormal basis.

Theorem 7.4

- If $\mathbf{b}_1, \ldots, \mathbf{b}_n$ is an orthonormal basis for an inner product space V, then $\mathbf{v} = \langle \mathbf{b}_1, \mathbf{v} \rangle \mathbf{b}_1 + \cdots + \langle \mathbf{b}_n, \mathbf{v} \rangle \mathbf{b}_n$ for any $\mathbf{v} \in V$.
- Proof:
  - Since $\mathbf{b}_1, \ldots, \mathbf{b}_n$ is a basis, we know that $\mathbf{v} = v_1\mathbf{b}_1 + \cdots + v_n\mathbf{b}_n$ for some uniquely determined $v_1, \ldots, v_n \in K$.
  - Using linearity and orthonormality (as in the above proof that orthonormal sets are linearly independent), we conclude that $\langle \mathbf{b}_i, \mathbf{v} \rangle = v_i$ for all i.

Similarly, we have a nice formula for the inner product of any two vectors represented in terms of an orthonormal basis.

Theorem 7.6

- If $\mathbf{b}_1, \ldots, \mathbf{b}_n \in V$ is an orthonormal basis for an inner product space V, then for any two vectors $\mathbf{x} = x_1\mathbf{b}_1 + \cdots + x_n\mathbf{b}_n$ and $\mathbf{y} = y_1\mathbf{b}_1 + \cdots + y_n\mathbf{b}_n$, their inner product is given by $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n \bar{x}_i y_i$.
  - This follows immediately from (sesqui-)linearity and orthonormality.

Projection

- The above Theorem 7.4 can be understood geometrically as "any vector is the sum of its projections onto the vectors of any orthonormal basis"; to understand this, we need the notion of projection.
- If $\mathbf{u} \in V$ is a vector in an inner product space V, the (orthogonal) projection onto $\mathbf{u}$ is the linear map $\Pi_{\mathbf{u}} \colon V \to \operatorname{Span} \mathbf{u}$ given by $\Pi_{\mathbf{u}}\mathbf{v} = \frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\langle \mathbf{u}, \mathbf{u} \rangle}\mathbf{u}$.
  - When $\mathbf{u}$ is a unit vector, this simplifies to $\Pi_{\mathbf{u}}\mathbf{v} = \langle \mathbf{u}, \mathbf{v} \rangle \mathbf{u}$.

- In this case, the *length* of the projection of $\mathbf{v}$ onto $\mathbf{u}$ is simply given by $|\langle \mathbf{u}, \mathbf{v} \rangle|$.
- To understand where this definition comes from, one should consider the case of $\mathbb{R}^2$ or $\mathbb{R}^3$, and draw a triangle and do some trigonometry. (The definition is also justified by Theorem 7.5 below.)

- More generally, let $\mathbf{b}_1, \ldots, \mathbf{b}_m \in V$ be any orthonormal set of vectors, and let $W = \mathrm{Span}(\mathbf{b}_1, \ldots, \mathbf{b}_m)$ (so that $\mathbf{b}_1, \ldots, \mathbf{b}_m$ is an orthonormal basis of W).
  - We define the (orthogonal) projection onto W to be the linear map $\Pi_W \colon V \to W$ given by $\Pi_W \mathbf{v} = \langle \mathbf{b}_1, \mathbf{v} \rangle \mathbf{b}_1 + \cdots \langle \mathbf{b}_m, \mathbf{v} \rangle \mathbf{b}_m$ (this should be compare to Theorem 7.4 above).

Theorem 7.5

- With $V, W, \mathbf{b}_1, \ldots, \mathbf{b}_m$ as above, for any $\mathbf{v} \in V$, the orthogonal projection $\Pi_W \mathbf{v} \in W$ is the unique closest vector to $\mathbf{v}$ in W, i.e., we have $\|\mathbf{v} - \Pi_W \mathbf{v}\| \leq \|\mathbf{v} - \mathbf{w}\|$ for any $\mathbf{w} \in W$, and if this is an equality, then $\mathbf{w} = \Pi_w \mathbf{v}$.

- In particular, this proves that the orthogonal projection map Π_W depends only on the finite-dimensional subspace W and not on the given choice of orthonormal basis.

- The proof rests on the following observation, which is of independent interest: that $\mathbf{v} - \Pi_W \mathbf{v} \perp$ W \) ("$\mathbf{v} - \Pi_W \mathbf{v}$ is orthogonal to W"), by which we mean $(\mathbf{v} - \Pi_W \mathbf{v}) \perp \mathbf{w}$ for every $\mathbf{w} \in W$.
  - This follows by linearity as soon as we show that $(\mathbf{v} - \Pi_W \mathbf{v}) \perp \mathbf{b}_i$ for every $\mathbf{b}_i$, and this follows immediately from sesquilinearity and orthonormality.

- Proof of the theorem:
  - Using the Pythagorean theorem and $(\mathbf{v} - \Pi_W \mathbf{v}) \perp W$, we have, for any $\mathbf{w} \in W$, that

$$\|\mathbf{v} - \mathbf{w}\|^2 = \|(\mathbf{v} - \Pi_W \mathbf{v}) + (\Pi_W \mathbf{v} - \mathbf{w})\|^2 = \|\mathbf{v} - \Pi_W \mathbf{v}\|^2 + \|\Pi_W \mathbf{v} - \mathbf{w}\|^2 \geq \|\mathbf{v} - \Pi_W \mathbf{v}\|^2.$$

  - Moreover, the above inequality is an equality if and only if $\|\Pi_W \mathbf{v} - \mathbf{w}\| = 0$, i.e., if and only if $\Pi_W \mathbf{v} = \mathbf{w}$; this proves uniqueness.

The Gram-Schmidt process

- What we showed above is that there exists an orthogonal projection operator Π_W onto any finite-dimensional subspace W *assuming W has an orthonormal basis.*

- We now recall the Gram-Schmidt process, which shows that **any finite-dimensional inner product space *does* have an orthonormal basis**.
  - More precisely, for any finite set of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m \in V$ is an inner product space, the Gram-Schmidt process gives explicit formulas for a set of orthonormal vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k$ having the same span.

- The process is to first produce a set of *orthogonal* vectors $\mathbf{w}_i$, by subtracting from each $\mathbf{v}_i$ the orthogonal projection onto the span of the previous vectors.
  - That is, we define $\mathbf{w}_1 = \mathbf{v}_1$, and then inductively define

$$\mathbf{w}_{i+1} = \mathbf{v}_i - \Pi_{\mathrm{Span}\{\mathbf{w}_1, \ldots, \mathbf{w}_i\}} \mathbf{v}_i = \mathbf{v}_i - \frac{\langle \mathbf{w}_1, \mathbf{v} \rangle}{\|\mathbf{w}_1\|^2} \mathbf{w}_1 - \cdots - \frac{\langle \mathbf{w}_i, \mathbf{v} \rangle}{\|\mathbf{w}_i\|^2} \mathbf{w}_i.$$

- If along the way, any $\mathbf{v}_i$ was dependent on the previous $\mathbf{v}_j$'s, then we get $\mathbf{w}_i = 0$, and in this case, we discard it (that is, we replace our original sequence $\mathbf{v}_1, \ldots, \mathbf{v}_m$ with the sequence $\mathbf{v}_1, \ldots, \mathbf{v}_{i-1}, \mathbf{v}_{i+1}, \ldots, \mathbf{v}_m$).
- We then define $\mathbf{b}_i = \frac{\mathbf{w}_i}{\|\mathbf{w}_i\|}$ to obtain an orthonormal set.

- (One can normalize the $\mathbf{w}_i$ to obtain the $\mathbf{b}_i$ "along the way" while performing the algorithm, instead of doing it at the end, but doing so usually makes the computations more complicated.)

10.2 3.7C Orthogonal transformations

- A linear map $f\colon V \to W$ between inner product spaces is orthogonal if it preserves the inner product, i.e., $\langle f\mathbf{u}, f\mathbf{v}\rangle = \langle \mathbf{u}, \mathbf{v}\rangle$ for all $\mathbf{u}, \mathbf{v} \in V$.
  - Equivalently, $\|f\mathbf{v}\| = \|\mathbf{v}\|$ for all $\mathbf{v} \in V$ (this follows from the formula which recovers the inner product from the norm).
  - Equivalently, $\|f\mathbf{u} - f\mathbf{v}\| = \|\mathbf{u} - \mathbf{v}\|$ for all $\mathbf{u}, \mathbf{v} \in V$, i.e., f preserves the distances between points.

- If $\mathbf{b}_1, \ldots, \mathbf{b}_n$ is an orthonormal basis of V, then f is orthogonal if and only if $f(\mathbf{b}_1), \ldots, f(\mathbf{b}_n)$ is still an orthonormal set.
  - The $\Rightarrow$ direction is immediate.
  - The other direction follows because if $\mathbf{u} = \sum_i u_i\mathbf{b}_i$ and $\mathbf{v} = \sum_i v_i\mathbf{b}_i \in V$, then $\langle \mathbf{u}, \mathbf{v}\rangle = \sum_i \bar{u}_i v_i = \sum_{i,j} \bar{u}_i v_j \langle \mathbf{b}_i, \mathbf{b}_j\rangle$, and $\langle f\mathbf{u}, f\mathbf{v}\rangle = \sum_{i,j} \bar{u}_i v_j \langle f(\mathbf{b}_i), f(\mathbf{b}_j)\rangle$ using the linearity of f and sesquilinearity of the inner product.

- If V and W are real vector spaces, and we fix orthogonal bases of V *and* W, then $f\colon V \to W$ is orthogonal if and only if the matrix $A \in \mathbb{R}^{m\times n}$ representing f with respect to these bases satisfies $A^\top \cdot A = \mathrm{I}$.

- One can show that the orthogonal transformations $\mathbb{R}^2 \to \mathbb{R}^2$ and $\mathbb{R}^3 \to \mathbb{R}^3$ (equipped with the standard inner product) are exactly the usual rotations and reflections.

- An isomorphism of inner product spaces (or linear isometry) is a bijective orthogonal linear map.
  - We say that two inner product spaces $(V, \langle -, -\rangle_V)$ and $(W, \langle -, -\rangle_W)$ are isomorphic, denoted $(V, \langle -, -\rangle_V) \cong (W, \langle -, -\rangle_W)$ if there exists an isomorphism of inner product spaces $V \to W$.

- **Theorem**: any n-dimensional inner product space V over K is isomorphic to K^n (with the standard inner product).
  - Indeed, we have already seen that any choice $\mathbf{b}_1, \ldots, \mathbf{b}_n$ of basis for V determines a linear isomorphism $f\colon V \to K^n$ with $f(\mathbf{b}_i) = \mathbf{e}_i$, the i-th standard basis vector.
  - If this is moreover an *orthonormal basis*, then f is an orthogonal map, and hence an isomorphism of inner product spaces.

- Example 3.7.1 again
  - The theorem gives a new perspective on the "distorted" inner product $\langle \mathbf{u}, \mathbf{v}\rangle = u_1v_1 + 2u_2v_2$: we have an isomorphism $f\colon \mathbb{R}^2 \to \mathbb{R}^2$ taking the distorted inner product to the standard inner product, namely $f(v_1, v_2) = (v_1, \sqrt{2} \cdot v_2)$.
  - Hence, we can understand the distorted inner product $\langle \mathbf{u}, \mathbf{v}\rangle$ as being described by the procedure "first apply f, then take the ordinary inner product".

10.3 Dual spaces

- The dual space V^* of a vector space V over K is the vector space $\mathcal{L}(V, K)$ of all linear maps $V \to K$; elements of V^* are sometimes referred to as *covectors*, *dual vectors*, *functionals*, or *linear forms*.
- Each linear map $f\colon V \to W$ between vector spaces induces a dual map $f^*\colon W^* \to V^*$ by composition: for $\varphi \in W^*$, we set $f^*(\varphi) = \varphi \circ f \in V^*$.
  - The dual map f^* is linear.
  - Moreover, the operation $f \mapsto f^*$ determines a linear map $\mathcal{L}(V, W) \to \mathcal{L}(W^*, V^*)$.
  - Given linear maps $U \xrightarrow{f} V \xrightarrow{g} W$, we have $(g \circ f)^* = f^* \circ g^*\colon W^* \to U^*$.
- Any basis $\mathbf{b}_1, \ldots, \mathbf{b}_n$ of V gives rise to a dual basis $\varphi_1, \ldots, \varphi_n$ of V^*, determined by the conditions $\varphi_i(\mathbf{b}_j) = \delta_{ij}$ for all i, j.
  - That these conditions determine functionals φ_i follows from the fact that a linear map is completely specified by giving its values on any given basis.
  - To see that the φ_i form a basis: linear independence follows from the fact that if $\varphi = \sum_{i=1}^n a_i\varphi_i$, then $a_i = \varphi(\mathbf{b}_i)$ for all i; and the spanning property follows from the fact that for any $\varphi \in V^*$, we have $\varphi = \sum_{i=1}^n \varphi(\mathbf{b}_i)\varphi_i$.
    - Both these facts are established by evaluating both sides of the equation on all of the basis vectors $\mathbf{b}_i$.
- Dual maps are related to transposes of matrices in the following way:
  - Let V and W be vector spaces over K of dimensions $m = \dim V$ and $n = \dim W$, and fix bases $\mathcal{B}_V$ of V and $\mathcal{B}_W$ of W, and hence dual bases $\mathcal{B}_V^*$ and $\mathcal{B}_W^*$ of V^* and W^*. Given a linear map $f\colon V \to W$ represented by a matrix $A \in K^{n\times m}$ with respect to $\mathcal{B}_V$ and $\mathcal{B}_W$, the dual map $f^*\colon W^* \to V^*$ is represented by the matrix $A^\top$, with respect to the dual bases of $\mathcal{B}_W^*$ and $\mathcal{B}_V^*$.
  - To prove this, write $\mathcal{B}_V = (\mathbf{v}_1, \ldots, \mathbf{v}_m)$ and $\mathbf{w}_1, \ldots, \mathbf{w}_n$, and $\mathcal{B}_V^* = (\varphi_1, \ldots, \varphi_m)$ and $\mathcal{B}_W^* = (\psi_1, \ldots, \psi_n)$.
    - We then have by definition that $f(\mathbf{v}_i) = \sum_{j=1}^n A_{ji}\mathbf{w}_j$, and hence that $\psi_j(f(\mathbf{v}_i)) = A_{ji}$.
    - Similarly, if $B \in K^{m\times n}$ is the matrix representing f^* with respect to $\mathcal{B}_W^*$ and $\mathcal{B}_V^*$, then $f^*(\psi_j) = \sum_{i=1}^m B_{ij}\varphi_i$, and hence $(f^*(\psi_j))(\mathbf{v}_i) = B_{ij}$.
    - Hence
      $$B_{ij} = (f^*(\psi_j))(\mathbf{v}_i) = \psi_j(f(\mathbf{v}_i)) = A_{ji},$$
      so $B = A^\top$, as desired.
- **Lemma**: for any non-zero vector $\mathbf{v} \in V$ in a finite-dimensional vector space, there is some covector $\varphi \in V^*$ with $\varphi(\mathbf{v}) \neq 0$.
  - Indeed, we may extend $\{\mathbf{v}\}$ to a basis, and take φ to be the dual basis vector corresponding to $\mathbf{v}$, so that $\varphi(\mathbf{v}) = 1$.
  - (In fact, this Lemma also holds for infinite-dimensional spaces.)
- There is a natural linear map $\alpha\colon V \to V^{**}$ from a vector space to the dual of its dual space, given by $\alpha(v) = \mathrm{ev}_{\mathbf{v}}$, where $\mathrm{ev}_{\mathbf{v}}\colon V^* \to K$ is defined by $\mathrm{ev}_{\mathbf{v}}(\varphi) = \varphi(v)$.
  - By the Lemma, α is always injective, since if $\mathrm{ev}_{\mathbf{v}} = 0$, then $\varphi(\mathbf{v}) = \mathrm{ev}_{\mathbf{v}}(\varphi) = 0$ for every $\varphi \in V^*$.

  - If V is finite-dimensional, it follows that α is an isomorphism, since V and V^{**} have the same dimension.
  - That is, **the map α provides a canonical isomorphism between any finite-dimensional vector space and its double dual space**.

- Note that if V is finite-dimensional, we not only have $V \cong V^{**}$, but also $V \cong V^*$, simply for dimension reasons: $V \cong K^{\dim V} \cong V^*$.
  - However, unlike with V^{**}, there is no *canonical* such isomorphism; one has to *choose* bases of V and V^* in order to produce an isomorphism $V \to V^*$.
  - (Actually, it suffices to just choose a basis of V, since this induces a dual basis of V^*, and hence an isomorphism $V \xrightarrow{\sim} V^*$ taking each basis vector in V to the corresponding dual basis vector.)

- There is another important source of isomorphisms $V \xrightarrow{\sim} V^*$: on any real inner product space V there is a canonical (i.e., only depending on the inner product, but on no further choices) isomorphism $\rho\colon V \xrightarrow{\sim} V^*$, which takes a vector $\mathbf{v} \in V$ to the functional $\rho(\mathbf{v})\colon V \to \mathbb{R}$ defined by $\rho(\mathbf{v})(\mathbf{w}) = \langle \mathbf{v}, \mathbf{w} \rangle$.
  - ρ is injective by positive-definiteness: if $\rho(\mathbf{v}) = 0$, then $\langle \mathbf{v}, \mathbf{v} \rangle = \rho(\mathbf{v})(\mathbf{v}) = 0$ and hence $\mathbf{v} = 0$.
  - It then follows that ρ is an isomorphism since $\dim V = \dim V^*$.

- If $V = \mathbb{R}^{n,1}$ is the space of column vectors, then V^*, the space of linear maps $\mathbb{R}^{n,1} \to \mathbb{R} \cong \mathbb{R}^{1,1}$, is naturally identified with the space $\mathbb{R}^{1,n}$ of $(1 \times n)$-matrices, i.e., row vectors.
  - Under these identifications, the isomorphism $V \xrightarrow{\sim} V^*$ induced by the standard inner product on $\mathbb{R}^{n,1}$ is then simply the transposition map $\mathbb{R}^{n,1} \xrightarrow{\sim} \mathbb{R}^{1,n}$.
  - This happens to be the same as the isomorphism induced by taking the standard basis of $\mathbb{R}^{n,1}$ to its dual basis in $\mathbb{R}^{1,n}$.

11 Mar 3, 11.1: differential operators

11.1 11: Second-order equations

- We'll now start to study second order equations $y'' = F(x, y, y')$.
- Mostly we'll look at *linear* equations $y'' + f(x)y' + g(x)y = h(x)$, and more specifically, linear equations with *constant coefficients* $y'' + ay' + by = h(x)$ with $a, b \in \mathbb{R}$.
  - Some examples:

$$y'' - 2y = x$$
$$y'' - 3y = e^x$$
$$y'' - 3y' + 2y = f(x)$$

- Later, we'll also look at some *nonlinear* equations such as $y'' + (y')^2 = 0$.

11.2 11.1: Differential operators

- To solve a first-order linear equation $y' + g(x)y = f(x)$, we used the exponential multiplier $e^{\int g(x)\,\mathrm{d}x}$.
  - If the equation has constant coefficients $y' + ry = f(x)$ (as in Newton's law of cooling), this simply becomes e^{rx}.
  - We recall that if y is any solution, we have $e^{rx}(y' + ry) = e^{rx}f(x)$ and hence $\frac{\mathrm{d}}{\mathrm{d}x}(e^{rx}y) = e^{rx}f(x)$ and $y = Ce^{-rx} + e^{-rx}\int e^{rx}f(x)\,\mathrm{d}x$ for some $C \in \mathbb{R}$.
  - In what follows, we will be making further use of such exponential multipliers.
- Before proceeding, we introduce some notation
- For a given interval I, recall that we have the vector space $\mathcal{C}^\infty(I)$ of smooth functions on I.
  - We denote by $D\colon \mathcal{C}^\infty(I) \to \mathcal{C}^\infty(I)$ the *differentiation operator* given by $Dy = y'$.
  - We recall that D can alternatively be considered as having different domains: for example, it can regarded as an operator $D\colon \mathbb{R}[x] \to \mathbb{R}[x]$ acting only on *polynomial* functions; or alternatively as a linear map $D\colon \mathcal{C}^{k+1}(I) \to \mathcal{C}^k(I)$ which lowers the degree of differentiability by one.
  - It is good to keep these in mind and to be flexible in one's interpretation of D.
  - (There is a modern answer to the question of what is the "correct" "largest" space of functions on which D can be said to act: this is the space of so-called Schwartz distributions; these are generalizations of functions, which include all of the continuous functions (even the non-differentiable ones!), and include (useful!) exotic entities like the *Dirac delta distribution*, which is the mass-distribution function of a point particle with non-zero mass. We may return to this later.)
- We recall that the set $\mathcal{L}(V)$ of operators on any vector space V itself forms a vector space.
  - Thus we can form new *differential operators* by addition and scalar multiplication. For example $(2D + 3\mathrm{I})y = 2Dy + 3\mathrm{I}y = 2y' + 3y$, where $\mathrm{I} \in \mathcal{L}(V)$ is the *identity operator*; normally we just omit I and write $2D + 3$ in place of $2D + 3\mathrm{I}$.
  - Moreover, we have the bilinear composition map $\circ\colon \mathcal{L}(V) \times \mathcal{L}(V) \to \mathcal{L}(V)$; given $f, g \in \mathcal{L}(V)$, we may just write $g \cdot f$ or gf for $g \circ f$ and write f^2 for $f \circ f$.

- We can thus obtain further examples by composition: for example $(D^2-1)y = D^2y - 1\cdot y$ and $(D+s)(D+t)y = (D+s)(y'+ty) = D(y'+ty) + s(y'+ty) = y'' + (s+t)y' + sty$. Note that using bilinearity of composition, we can say directly that $(D+s)(D+t) = D^2 + (s+t)D + st$.

- We may thus say that a second-order linear differential equation with constant coefficients is one of the form $(D^2 + aD + b)y = f(x)$, the entity on the left being a *second-order linear differential operator*.

Characteristic equations

- Let us now consider a *homogeneous* second-order linear differential equation $y'' + ay' + b = 0$.

- Motivated by the first-order homogeneous linear equation, let us look for solutions of the form $y = e^{rx}$.

  - (Such a guess about the general form of a solution is often referred to using the German word *Ansatz*, meaning "approach" or "starting point"; so we can say "we make the *Ansatz* $y = e^{rx}$".)

  - We find that this y is a solution if and only if $(r^2 + ar + b)e^{rx} = 0$.
  - The resulting equation $r^2 + ar + b = 0$ is called the characteristic equation associated to the given differential equation.

Example 11.1.3

- $y'' - 3y' + 2y = 0$ has characteristic equation $r^2 - 3r + 2 = 0$.

- Factoring, we find this has roots 1 and 2.

- Thus $y_1(x) = e^x$ and $y_2(x) = e^{2x}$ are both solutions.

- Since the operator $L = D^2 - 3D + 2$ (for which our equation is $Ly = 0$) is *linear*, we have $L(c_1y_1 + c_2y_2) = 0$ for any $c_1, c_2 \in \mathbb{R}$.

- Thus, $y(x) = c_1e^x + c_2e^{2x}$ is a solution for any $c_1, c_2 \in \mathbb{R}$..

- In fact, this is the general solution, as we will learn how to prove next.

12 Mar 5, 11.1B and 11.2A: More second-order equations

12.1 11.1B Factoring operators

- To find the general solution to a homogeneous second-order linear differential equation $Ly = 0$, we factor L and reduce to solving first-order linear equations.

Example 11.1.4

- We seek all solutions $y\colon \mathbb{R} \to \mathbb{R}$ to $y'' + 5y' + 6 = 0$.
- We have $D^2 + 5D + 6 = (D + 3)(D + 2)$.
- Thus, we are seeking solutions of $(D + 3)(D + 2)y = 0$.
- We see that that the function $u = (D + 2y)$ must solve the first-order linear ODE $(D + 3)u = 0$; but we know the solutions to such an ODE: we must have $(D + 2)y = u = c_1e^{-3x}$ for some $c_1 \in \mathbb{R}$.
- We have reduced to the (non-homogeneous) first-order ODE $y' + 2y = c_1e^{-3x}$; we know how to solve this with exponential multipliers.
- We have $e^{2x}(y' + 2y) = c_1e^{-x}$, hence $\frac{\mathrm{d}}{\mathrm{d}x}(ye^{2x}) = c_1e^{-x}$, hence $ye^{2x} = -c_1e^{-x} + c_2$, hence $y = -c_1e^{-3x} + c_2e^{-2x}$.
- We have thus found the general solution; as expected, it is the same as the one resulting from our Ansatz $y = e^{rx}$.

Theorem 11.1.1

- The above procedure leads to the following general theorem:
  - Given a differential equation $y'' + ay' + b = 0$, if the corresponding characteristic equation $r^2 + ar + b = 0$ has two distinct roots $r_1, r_2 \in \mathbb{R}$, then the general solution $y\colon \mathbb{R} \to \mathbb{R}$ to the equation is $y = c_1e^{r_1x} + c_2e^{r_2x}$.
  - If the characteristic equation has a double root $r_1 = r_2 \in \mathbb{R}$, then the general solution is $y = c_1e^{r_1x} + c_2xe^{r_1x}$.
  - In both cases, for any $x_0, y_0, z_0 \in \mathbb{R}$, the initial value problem

$$\begin{cases} y'' + ay' + b = 0 \\ y(x_0) = y_0 \\ y'(x_0) = z_0 \end{cases}$$

    has a unique solution (i.e., these initial conditions uniquely determine the coefficients c_1, c_2 in the general solution).

- The proof:
  - As above, we can write the equation as $(D - r_1)(D - r_2)y = 0$.
  - We thus have that $y\colon \mathbb{R} \to \mathbb{R}$ is a solution if and only if $(D-r_2)y = c_1e^{r_1x}$, i.e., $y'-r_2y = c_1e^{r_1x}$ for some $c_1 \in \mathbb{R}$.
  - This is equivalent to $\frac{\mathrm{d}}{\mathrm{d}x}(ye^{-r_2x}) = c_1e^{(r_1-r_2)x}$.
  - If $r_1 \neq r_2$, this is equivalent to $ye^{-r_2x} = (r_1 - r_2)^{-1}c_1e^{(r_1-r_2)x} + c_2$ for some $c_2 \in \mathbb{R}$, and hence to $y = c_1e^{r_1x} + c_2e^{r_2x}$ for some $c_1, c_2 \in \mathbb{R}$.

- If $r_1 = r_2$, so that $e^{(r_1-r_2)x} = 1$, the above is equivalent to $ye^{-r_2x} = c_1x + c_2$ for some $c_2 \in \mathbb{R}$, and hence to $y = c_1e^{r_1x} + c_2xe^{r_1x}$ for some $c_1, c_2 \in \mathbb{R}$.
- In the first case, the initial conditions $y(x_0) = y_0$ and $y'(x_0) = z_0$ give

$$\begin{cases} y_0 = c_1e^{r_1x_0} + c_2e^{r_2x_0} \\ z_0 = r_1c_1e^{r_1x_0} + r_2c_2e^{r_2x_0} \end{cases}$$

- In the second case, they give

$$\begin{cases} y_0 = c_1e^{r_1x_0} + c_2x_0e^{r_2x_0} \\ z_0 = r_1c_1e^{r_1x_0} + c_2(1 + r_1x_0)e^{r_2x_0} \end{cases}$$

- In both cases, this is a system of two *linear equations* in c_1 and c_2, and we claim in both cases that this has a unique solution $(c_1, c_2) \in \mathbb{R}^2$.

- In fact, it is not hard to solve these systems of linear equations directly, but it is better to do a little more linear algebra review and recall some general principles.
  - Namely, recall that an inverse of a square matrix $A \in K^{n\times n}$, is a matrix $A^{-1} \in K^{n\times n}$ with the property that $A \cdot A^{-1} = A^{-1} \cdot A = \mathrm{I}_n$. If A has an inverse (or as we say, if it is invertible), it has a *unique* inverse.
  - If A is invertible, then for any $b \in K^n$, the system of equations $A\mathbf{x} = \mathbf{b}$ has the unique solution $\mathbf{x} = A^{-1}\mathbf{b}$.
  - A criterion for invertibility is that A is invertible if and only if its *determinant* $\det A$ is non-zero.
  - We will recall more about determinants later; for now, we recall that for a 2×2-matrix $A \in K^{2\times 2}$, the determinant is given by $\det A = A_{11}A_{22} - A_{12}A_{21}$.

- Returning to the above case, the determinants of the matrices appearing in those systems of linear equations are in the first case $e^{(r_1+r_2)x_0}(r_1 - r_2)$, which is non-zero since we are assuming $r_1 \neq r_2$, and in the second case $e^{(r_1+r_2)x_0}(1 + r_1x_0 - r_1x_0) \neq 0$.

12.2 11.2A Complex Exponentials

- We have solved the homogeneous second-order linear equation $y'' + ay' + b = 0$ in the case that the characteristic equation $r^2 + ar + b = 0$ has two *real* roots.

- As we know, a general quadratic polynomial may have complex roots; to deal with the corresponding differential equation in this case, we must introduce the *complex exponential function.*

- For this purpose, let us first ask ourselves: what *is* the exponential function, and also: what are the cosine and sine functions?

- As we have seen, one way to characterize the exponential function is as the unique solution $y = e^x$ to the IVP $y' = y;\ y(0) = 1$.
  - We saw how to prove uniqueness *assuming* that we already have the exponential function satisfying $\frac{\mathrm{d}}{\mathrm{d}x}e^x = e^x$ and $e^0 = 1$; but how do we know that such a function exists?
  - One way is to use power series.

- Remember that any power series function $f(x) = \sum_{n=0}^{\infty} a_nx^n$ has a *radius of convergence* $R \in [0, \infty]$, such that the series converges whenever $|x| < R$, and diverges when $|x| > R$ (and this is true even for $x \in \mathbb{C}$!).

- One can see (for example using the comparison test against a geometric series) that the radius of convergence of the series $\exp(x) = \sum_{n=0}^{\infty} \frac{x^n}{n!}$ is ∞. (This series should look familiar: it is the Taylor series of the exponential function; but now, we are using this series to *define* the exponential function!).
- Moreover, one can show that any power series function $f(x) = \sum_{n=0}^{\infty} a_n x^n$ is differentiable within its radius of convergence and its derivative is given term-by-term: $f'(x) = \sum_{n=0}^{\infty} n a_n x^{n-1}$.
- It follows that $\exp'(x) = \exp(x)$, and we clearly have $\exp(0) = 1$, so this proves the existence of a solution to the IVP defining e^x.

- Note that the addition law $e^{a+b} = e^a \cdot e^b$ follows from the characterization of e^x as the unique solution to $y' = y;\ y(0) = 1$.
  - Indeed, if we set $y(x) = e^{a+x}/e^a$, then $y' = y$ (by the chain rule) and $y(0) = e^a/e^a = 1$, and hence $y(x) = e^x$, i.e., $e^{a+x} = e^a \cdot e^x$.
  - In particular, this also gives $1 = e^{a-a} = e^a \cdot e^{-a}$ and hence $e^{-a} = 1/e^a$.

13 Mar 10, 11.2A: Complex solutions

13.1 More on complex exponentials

- Last time, we defined the exponential function as $e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!}$.
- A consequence of this is that we can immediately make sense of e^z for *complex* values $z \in \mathbb{C}$, simply by plugging z into the power series.
  - In particular, for $x \in \mathbb{R}$, we find that

$$\begin{aligned} e^{ix} &= 1 + ix - \frac{x^2}{2} - i\frac{x^3}{3!} + \frac{x^4}{4!} + \cdots \\ &= \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n}}{2n!} + i \sum_{n=0}^{\infty} (-1)^n \frac{x^{2n+1}}{(2n+1)!} \end{aligned}$$

  - We recognize the two terms appearing as the Taylor series for cos and sin.
  - As with the exponential function, let us *define* cos and sin to be given by these two power series.
- It follows immediately from this definition that $\sin' = \cos$ and $\cos' = -\sin$.
  - Soon, we will see that there are in fact *unique* functions satisfying the initial value problems $y'' = -y; y(0) = 0; y'(0) = 1$ and $y'' = -y; y(0) = 1; y'(0) = 0$, respectively. Hence, as with the exponential function, we get a nice condition characterizing cos and sin uniquely, and the above series prove the existence of functions satisfying these conditions.
  - It is also not too hard to show from these definitions that $\cos^2 + \sin^2 = 1$ and that they agree with the geometric definition of the trigonometric functions, i.e., that $\mathbf{p} = (\cos(\theta), \sin(\theta)) \in \mathbb{R}^2$ is the point on the unit circle such that the arc-length from $(1, 0)$ is $\mathbf{p}$, measured counter-clockwise, is θ.
- This way of defining e^x, cos, and sin immediately gives rise to the famous *Euler's formula*:

$$e^{ix} = \cos x + i \sin x.$$

  - This gives a concise way to express points in the plane using polar coordinates: the point with radius r and angle θ is $re^{i\theta}$.
    - As usual with radial coordinates, the angle is not uniquely determined: we always have $re^{i\theta} = re^{i(\theta+2\pi)}$ (and conversely, if $r_1 e^{i\theta_1} = r_2 e^{i\theta_2}$, then $r_1 = r_2$ and $\theta_1 - \theta_2 \in 2\pi\mathbb{Z}$ – the one exception being that $0e^{i\theta} = 0$ for *any* θ)

Some properties of the complex exponential

- A variant (using a bit of complex analysis) of the argument given above to deduce the exponential law $e^{a+b} = e^a \cdot e^b$ for $a, b \in \mathbb{R}$ proves that this holds as well for $a, b \in \mathbb{C}$.
  - From this, we can deduce the addition laws for sin and cos:

$$\begin{aligned} &\cos(a+b) + i\sin(a+b) = e^{i(a+b)} = e^{ia}e^{ib} \\ &= (\cos a + i \sin a)(\cos b + i \sin b) \\ &= (\cos a \cos b - \sin a \sin b) + i(\sin a \cos b + \cos a \sin b) \end{aligned}$$

    hence by comparing real and imaginary parts, we get $\cos(a+b) = \cos a \cos b - \sin a \sin b$ and $\sin(a+b) = \sin a \cos b + \cos a \sin b$.

- This also makes complex numbers easy to multiply when written in polar coordinates: $(r_1 e^{i\theta_1})(r_2 e^{i\theta_2}) = (r_1 r_2) e^{i(\theta_1+\theta_2)}$.
- In particular, this allows us to easily find square roots (and more generally n-th roots): if $z = re^{i\theta}$, then its square roots – i.e., the numbers $w \in \mathbb{C}$ such that $w^2 = z$ – are just $w = \pm\sqrt{r}e^{i\theta/2}$.
  - Indeed, we see that these *are* square roots, and given any other square root w, we have $w^2 - z = (w - \sqrt{r}e^{i\theta/2})(w + \sqrt{r}e^{i\theta/2})$, and hence $w = \pm\sqrt{r}e^{i\theta/2}$.
  - (Regarding the ambiguity of θ: had we written $z = re^{i\theta+2pi}$, we would have gotten the same square roots $w = \pm\sqrt{r}e^{i\theta/2+\pi} = \mp\sqrt{r}e^{i\theta/2}$.
- In particular, if z is a negative real number, then $z = re^{i\pi}$, and we have the familiar imaginary square root $\sqrt{z} = \pm\sqrt{r}e^{i\pi/2} = \pm i\sqrt{r}$.

- Next, recall that the derivative of a function $f\colon \mathbb{R} \to \mathbb{C} = \mathbb{R}^2$ is defined component-wise: if $f(x) = u(x) + iv(x)$, then $f'(x) = u'(x) + iv'(x)$.
  - It follows that
$$\tfrac{\mathrm{d}}{\mathrm{d}x} e^{ix} = \cos' x + i \sin' x = -\sin x + i \cos x = ie^{ix}$$
  - Hence, the *anti-derivative* $\int e^{ix}\,\mathrm{d}x$ of e^{ix} (the unique-up-to-a-constant function $f\colon \mathbb{R} \to \mathbb{C}$ whose derivative is e^{ix}) is $\frac{1}{i}e^{ix} + C = -ie^{ix} + C$, where $C \in \mathbb{C}$ is an arbitrary *complex* constant.
  - More generally, using that $e^{a+ib} = e^a \cdot e^{ib}$, we have that $\frac{\mathrm{d}}{\mathrm{d}x} e^{(a+ib)x} + (a+ib)e^{(a+ib)x}$ and $\int e^{(a+ib)x}\,\mathrm{d}x = \frac{1}{a+ib}e^{(a+ib)x} + C$.

13.2 Complex solutions to second-order equations

Example 11.2.1

- Now consider a general homogeneous second order linear equation $Ly = (D^2 + aD + b)y = 0$, with its characteristic equation $r^2 + ar + b = 0$.
  - We are perhaps still mainly interested in solutions $y\colon \mathbb{R} \to \mathbb{R}$, but now we can also try to find all solutions $y\colon \mathbb{R} \to \mathbb{C}$; and in this case, we can also consider equations with coefficients $a, b \in \mathbb{C}$.
- We now factor this polynomial as $(r - r_1)(r - r_2)$, with roots $-a \pm \frac{1}{2}\sqrt{a^2 - 4b}$ (where this square root is now possibly complex).
- We can thus factor the differential operator as $L = (D - r_1)(D - r_2)$.
  - We will see that the above exponential multiplier method still works because $D(e^{rx}y) = e^{rx}(D + r)y$ as before, even for $r \in \mathbb{C}$.
- As a first example, consider $y'' + y = 0$, i.e., $(D^2 + 1)y = 0$, i.e., $(D - i)(D + i)y = 0$.
  - We would like to conclude from this that $(D + i)y = c_1 e^{ix}$ for some $c_1 \in \mathbb{C}$.
  - That is, we would like to say that $u = Ce^{ix}$ is the *general* solution $u\colon \mathbb{R} \to \mathbb{C}$ to the differential equation $u' = iu$.
  - Indeed, the usual method works: given any solution u, we have $\frac{\mathrm{d}}{\mathrm{d}x}(ue^{-ix}) = u'e^{-ix} - iue^{-ix} = 0$, hence $ue^{-ix} = C$ for some constant $C \in \mathbb{C}$, and hence $u = Ce^{ix}$.
- It now remains to solve $y' + iy = c_1 e^{ix}$, and for this, we again use exponential multipliers.

– This equation is equivalent to $e^{ix}(y' + iy) = c_1 e^{2ix}$ (because we can multiply e^{-ix} to go back!), or in other words $\frac{\mathrm{d}}{\mathrm{d}x}(ye^{ix}) = c_1 e^{2ix}$.

  * Note that this use of the product rule is legitimate: it is just the ordinary product rule applied to each of the two components of the function $\mathbb{R} \to \mathbb{C}$ given by $x \mapsto y(x)e^{ix}$.

– This is equivalent to $ye^{ix} = \frac{1}{2}c_1 e^{2ix} + c_2$ for some $c_2 \in \mathbb{C}$ and hence to

$$y = c_1 e^{ix} + c_2 e^{-ix}$$

for some $c_1, c_2 \in \mathbb{C}$; thus, this is the general solution.

- We can rewrite this as

$$\begin{aligned} y &= c_1(\cos x + i \sin x) + c_2(\cos x - i \sin x) \\ &= (c_1 + c_2)\cos x + (c_1 - c_2)i \sin x \\ &= d_1 \cos x + d_2 \sin x, \end{aligned}$$

where we have set $d_1 = c_1 + c_2$ and $d_2 = i(c_1 - c_2)$.

– Since we can solve for c_1 and c_2 in terms of d_1 and d_2, this formula also expresses the general solution.

– It is the general solution we might have expected for $y'' = -y$, except that now d_1 and d_2 may be *complex*.

– The general *real* solution is obtained by restricting to the case $d_1, d_2 \in \mathbb{R}$.

Theorem 11.2.3

- The above example shows how the proof of Theorem 11.1.1 above can be adapted to the complex context, to yield the following statement:

- Given any $a, b \in \mathbb{C}$, the differential equation $y'' + ay' + by = 0$ has the general solution $y \colon \mathbb{R} \to \mathbb{C}$ given by

$$\begin{cases} y = c_1 e^{r_1 x} + c_2 e^{r_2 x}, & r_1 \neq r_2 \\ y = c_1 x e^{r_1 x} + c_2 e^{r_2 x}, & r_1 = r_2, \end{cases}$$

with $c_1, c_2 \in C$, where $r_1, r_2 \in \mathbb{C}$ are the roots of $r^2 + ar + v = 0$.

- In the case where $a, b \in \mathbb{R}$ and $a^2 - 4b < 0$, so that $r_1 = \alpha + i\beta$ and $r_2 = \alpha - i\beta$ for some $\alpha, \beta \in \mathbb{R}$, the general solution can also be written

$$y = c_1 e^{\alpha x} \cos \beta x + c_2 e^{\alpha x} \sin \beta x,$$

where $c_1, c_2 \in \mathbb{C}$. In this case the *real* solutions $y \colon \mathbb{R} \to \mathbb{R}$ are precisely those with $c_1, c_2 \in \mathbb{R}$.

- Moreover, the initial conditions $y(x_0) = y_0$ and $y'(x_0) = z_0$, for any $x_0 \in \mathbb{R}$ and $y_0, z_0 \in \mathbb{C}$ can always be satisfied by a unique choice of c_1 and c_2.

Example 11.2.2

- A generalization of the equation $y'' + y = 0$ from Example 11.2.1 is the equation $y'' + \omega^2 y = 0$, where $\omega \in \mathbb{R} - \{0\}$; this is called the *harmonic oscillator equation* with *angular frequency* ω.

- This has characteristic equation $r^2 + \omega^2 = 0$ with roots $r = \pm i\omega$, and hence general solution

$$y = c_1 e^{i\omega x} + c_2 e^{-i\omega x} = d_1 \cos \omega x + d_2 \sin \omega x.$$

- Such functions are called *harmonic oscillators*, and arise frequently in physics.

Example 11.2.3

- We summarize the form of the general solution to $y'' + ay' + by = 0$ with $a, b \in \mathbb{R}$.
- This has characteristic equation $r^2 + ar + b =$ with roots $r_1, r_2 = -a/2 \pm \sqrt{a^2 - 4b}/2$.
- We now consider three cases, based on the sign of the discriminant $a^2 - 4b$:
- If $a^2 - 4b > 0$, then r_1, r_2 are real and distinct, and the general solution is $y = c_1 e^{r_1 x} + c_2 e^{r_2 x}$ with $c_1, c_2 \in \mathbb{R}$.
- If $a^2 - 4b < 0$, then r_1, r_2 are imaginary and distinct, and the general solution is $y = c_1 e^{\alpha x} \cos \beta x + c_2 e^{\alpha x} \sin \beta x$ with $c_1, c_2 \in \mathbb{R}$, where $\alpha = -a/2$, and $\beta = \sqrt{a^2 - 4b}/2$.
- If $a^2 - 4b = 0$ represents the dividing line between the above oscillatory and non-oscillatory behaviour; the general solution is $y = c_1 x e^{r_1 x} + c_2 e^{r_2 x}$.