

Contents

1	Vectors and matrices	1
1.1	Basic matrix results	1
1.2	Determinants	1
1.3	Other results	2
1.4	Cross product properties	2
2	Calculus	5
2.1	Polar coordinate integration	5
2.2	Power series	6
2.3	Differential equations	8
2.4	Other results	12
3	Conic sections	12
4	Recurrence relations	14

1 Vectors and matrices

1.1 Basic matrix results

Proposition. Matrix multiplication is associative (ie it satisfies $(AB)C = A(BC)$)

Proof. Let A be $n * m$, B be $m * k$, and C be $k * l$.

Then $(AB)_{ij}$ denotes the entry of AB in row i and column j . By matrix multiplication this is $\sum_{r=1}^m A_{ir}B_{rj}$. Also, $((AB)C)_{ij} = \sum_{r=1}^k (AB)_{ir}C_{rj}$. Using the sum for $(AB)_{ir}$ we get $((AB)C)_{ij} = \sum_{r=1}^k (AB)_{ir}C_{rj} = \sum_{r=1}^k \sum_{s=1}^m A_{is}B_{sr}C_{rj}$. If you do the same for $(A(BC))_{ij}$ you will get the same sum. I'm too lazy to actually do it but I think you probably can. If the elements all agree the matrices are equal.

□

We'll be extra rigorous here and give an obvious proposition.

Proposition. Matrices satisfy $(A+B)C=AC+BC$

Proof. This is kind of obvious since matrix addition just sums the elements, then it is immediate that:

$$\sum_{r=1}^m A_{ir}C_{rj} + \sum_{r=1}^m B_{ir}C_{rj} = \sum_{r=1}^m (A + B)_{ir}C_{rj}$$

Again, the elements agree so the matrices are equal.

□

1.2 Determinants

Theorem. The determinant of a general $n * n$ matrix as given by the formula stated in level 5 equals the volume scale factor of a shape transformed by the matrix.

Proof. This has been proven in a video using some visual arguments about vectors. Note that in this proof we show that the sign of the determinant changes when we swap the columns. The fact that whether we have an even or odd number of swaps is well defined (ie for some matrix you can tell which situation we are in) is something which we will prove in Level 7 since the proof is not hard but it is not needed for any of this, and the video basically shows geometrically that it has to be well defined because we know intuitively that the concept of a determinant makes sense. For the three dimensional

case you can think of whether the transformed axes are left or right handed (right handed meaning if you take the x-y plane to be parallel to the page with the positive y direction less than 180 degrees counterclockwise from the positive x direction then the z axis will go out of the page).

□

Theorem. The matrix inverse formula we stated in level 5 actually gives the inverse

Proof. This has been proven in the same video as the determinant proof video using the determinant formula and the fact that it is volume.

□

Note that a left inverse is also a right inverse, in the sense that if $BA = I$ then $AB = I$ because A and B are invertible and $BAB = B = BI$ and by associativity $BAB = BI$ and B has a left inverse C, meaning we can say $CB = I$ so $CBAB = CBI$ so $AB = I$ as required.

1.3 Other results

Proposition. The closest distance between a point (x_1, y_1, z_1) and a plane $Ax + By + Cz + D = 0$ is given by the formula $\frac{|Ax_1 + By_1 + Cz_1 + D|}{\sqrt{A^2 + B^2 + C^2}}$

Proof. This is because if a plane is defined by $ax + by + cz + d = 0$ then we know that any point's distance from the plane is given by the length of the line segment normal to the plane, from the plane to the point. We can also rewrite the equation of the plane as $(a, b, c) \cdot (x, y, z) = -d$. We know that the normal has direction (a, b, c) since for any points P, Q on the plane, $(P - Q) \cdot (a, b, c) = -d - (-d) = 0$ and also it is visually obvious that P-x is perpendicular to the normal of the plane. So the distance from any point R to the plane is the length of the component $R - P$ in the direction of n, where n is normal to the plane and P is a point on the plane. This means it is given by $\frac{|n \cdot (R - P)|}{|n|}$ but $n \cdot P = -d$ so it is given by $\frac{|n \cdot R - d|}{|n|}$, and the result follows.

□

Proposition. Matrices satisfy $(AB)^T = B^T A^T$

Proof. We want to show that $(AB)^T = B^T A^T$ for matrices A and B. This is true because the row i, column j element of AB is given by doing (row i of A) \cdot (column j of B) and so the row j, column i element of $(AB)^T$ is given by that same formula. The row j, column i element of the right hand side of the equation above is equal to (row j of B^T) \cdot (column i of A^T) which by the definition of the transpose is indeed equal to (row i of A) \cdot (column j of B). The matrices are equal element-wise so they are equal.

□

1.4 Cross product properties

We want to show that:

1. $A \times B$ is perpendicular to A and B and satisfies the right hand rule (ie right \times back = up)
2. $A \times B = -(B \times A)$ (anticommutativity)
3. $A \times (B + C) = A \times B + A \times C$ (distributivity)
4. $A \times B$ = area of parallelogram spanned by A and B

It will be helpful to think of the cross product as the determinant of the following matrix, since if you actually compute this determinant you can see that it is equal to the formula for the cross product that we assume is true and want to prove these properties from. (Note: i, j, k are the standard unit vectors, and the subscript denotes the coordinate)

$$A \times B = \det \begin{bmatrix} i & A_x & B_x \\ j & A_y & B_y \\ k & A_z & B_z \end{bmatrix}$$

This proves $A \times B = -(B \times A)$ as swapping two columns changes the sign of the determinant, and that $A \times (B + C) = A \times B + A \times C$ since the determinants are linear in columns as shown earlier. Now we can consider the determinant of the following matrix:

$$\begin{bmatrix} A_x & A_x & B_x \\ A_y & A_y & B_y \\ A_z & A_z & B_z \end{bmatrix}$$

Clearly it is zero because it has two columns the same, but another way to see this is to consider expanding it out by the first column, ie the determinant is

$$A_x(A_y B_z - B_y A_z) + A_y(B_x A_z - A_x B_z) + A_z(A_x B_y - B_x A_y)$$

But notice, this is actually equal to $A \cdot (A \times B)$, and so we have that $A \cdot (A \times B) = 0$. Applying the same argument to B gives that $A \times B$ is perpendicular to both A and B.

To prove the thing about area, we notice that the sum of the squares of the components of the cross product is the square of the cross product by pythagoras, and that value can also be shown (by a direct computation) to be given by this formula:

$$\det \begin{bmatrix} A_y B_z - B_y A_z & A_x & B_x \\ B_x A_z - A_x B_z & A_y & B_y \\ B_x A_z - A_x B_z & A_z & B_z \end{bmatrix}$$

So now picture this geometrically: We have a vector $A \times B$ perpendicular to vectors A and B and we want to find the length of $A \times B$. We know that the volume of the parallelepiped spanned by A, B and $A \times B$ (By the volume property of the determinant proven earlier and the determinant of the matrix above) equals the square of the magnitude of the cross product, but the volume of this parallelepiped also equals the magnitude of $A \times B$ times the area of the parallelogram spanned by A and B. Cancelling a factor of $A \times B$ on both sides gives the desired result.

Proposition. Let P be an eigenvector matrix of A and D a diagonal matrix with the eigenvalues, then $A = PDP^{-1}$. Quite frankly it's stupid that A level textbooks don't prove this.

Proof. The reason $A = PDP^{-1}$ is essentially that $(PDP^{-1})v$ takes v as input, and if we take v to be separated into eigenvector components, it first moves v such that those are its normal components (by the P^{-1}), scales everything by the eigenvalue factors (by the D), then moves it back (by the P).

□

Proposition. Eigenvectors with different eigenvalues are actually linearly independent (ie they don't collapse onto a lower dimensional subspace) so the matrix P in the decomposition is actually invertible

Proof. Suppose we have r different eigenvalues and we pick r eigenvectors, one for each eigenvalue. We want to show that they are linearly independent, ie none is a linear combination (sum of constant multiples) of the others.

Suppose that they are linearly independent, meaning there exists constants a_1, a_2, \dots, a_r not all 0 such that $a_1 v_1 + a_2 v_2 + \dots + a_r v_r = 0$ where v is our eigenvectors.

Take the p that is the least p such that p of the a's can be different from 0 with this satisfied.

So now suppose $a_1 v_1 + a_2 v_2 + \dots + a_p v_p = 0$ with all coefficients non-zero. Then

$(A - \lambda_1 I)(a_1 v_1 + a_2 v_2 + \dots + a_p v_p) = 0$ by assumption, but we know that $(A - \lambda_1 I)(a_1 v_1) = 0$ by definition, so $(A - \lambda_1 I)(a_2 v_2 + a_3 v_3 + \dots + a_p v_p) = 0$. Now this means that

$$0 = (A - \lambda_1 I)(a_2 v_2 + a_3 v_3 + \dots + a_p v_p)$$

$$= A(a_2v_2 + a_3v_3 + \dots + a_pv_p) - \lambda_1(a_2v_2 + a_3v_3 + \dots + a_pv_p)$$

Therefore, since these are eigenvectors,

$$(a_2(\lambda_1 - \lambda_2)v_2 + a_3(\lambda_1 - \lambda_3)v_3 + \dots + a_p(\lambda_1 - \lambda_p)v_p) = 0$$

This contradicts that p is the least number of non zero coefficients we can have.

□

Theorem. Symmetric matrices decompose into $P^TAP = D$ and this is the same as $P^{-1}AP = D$

Proof. It is sufficient to show that the eigenvectors with distinct eigenvalues are perpendicular, as then we see that if the eigenvalue matrix is A then $A^T A = AA^T = I$ by a dot product argument, so the transpose of A is the inverse of A . Here we assume that the symmetric matrix has real values, otherwise the claim about the transpose being the inverse, the more general result that works for complex numbers is about Hermitian matrices, where if you take the complex conjugate of the transpose of A you get the inverse of A , and for real numbers taking the complex conjugate simply does nothing, and the proof for that result goes exactly like the proof shown below.

If S is our symmetric matrix and v is an eigenvector with eigenvalue λ then by the definition of eigenvectors we have that $Sv = \lambda v$.

So, suppose we have eigenvalues λ and μ (which are not equal to each other, this is another assumption used in the mathematical statement we are proving) with eigenvectors v and u respectively, then we have

Suppose $Sv = \lambda v$ (1) and $Su = \mu u$ (2). We can multiply both sides of (1) on the left by the row vector u^T and both sides of (2) on the left by the row vector v^T . This will give $u^T Sv = \lambda u^T v$ and $v^T Su = \mu v^T u$, where the constants can be moved out to the front. Now we will take the transpose of both sides of this last equation so we get $(v^T Su)^T = \mu \mu^T (v^T u)$ and then use the relation about matrix transposes from earlier to simplify this to $u^T S^T v = \mu u^T v$. It is here that we use the assumption that S is symmetric, as by definition we have that $S = S^T$ so $u^T Sv = \mu u^T v$.

Since we have $u^T Sv = \lambda u^T v$ from earlier as well, it means $\mu u^T v = \lambda u^T v$ as if they are both equal to $u^T Sv$ they must be equal to each other. Therefore, we have 2 options: Either $\lambda = \mu$ (so the eigenvalues would not be distinct as per our assumption), or $u^T v = 0$. Because of how matrix multiplication works, finding $u^T v$ is actually the same as taking the dot product $u \cdot v$, so we are done and we have proven the perpendicularity as required since dot product is 0 implies perpendicular.

□

Theorem. (Cayley hamilton theorem) Every matrix satisfies its own characteristic equation.

Proof. We define $Adj(M)$ as the transpose of the cofactor matrix (ie, we are in the process of finding the inverse before we divide by the determinant, meaning this always exists as we never run into a division by 0 problem). We have $det(M) M^{-1} = Adj(M)$ if M is invertible, and we always have that $det(M) I = M Adj(M)$.

We define $B = Adj(tI - A)$ where A is an $n * n$ matrix with a characteristic equation in t that is $det(tI - A) = 0$. Since the elements of an adjugate matrix are built by taking the original matrix and deleting a row and column and finding the determinant (which in the case of $tI - A$ will be polynomials in t with a power of t not exceeding $n - 1$), we know that the elements of B must be polynomials in t of degree up to $n - 1$, so we can write B as

$\sum_{k=0}^{n-1} t^k B_k$ for matrices B_k with coefficients not depending on t .

We know that $det(tI - A)I = (tI - A)Adj(tI - A)$ so $det(tI - A) I = (tI - A) \sum_{k=0}^{n-1} t^k B_k$

On the left hand side, we have $t^n I + c_{n-1}t^{n-1}I + c_{n-2}t^{n-2}I + \dots + c_1 tI + c_0 I$ where the c 's are the coefficients of the characteristic polynomial of A and on the right hand side we have, by expanding,

$$t^n B_{n-1} + t^{n-1} (B_{n-2} - AB_{n-1}) + t^{n-2} (B_{n-3} - AB_{n-2}) + \dots + t(B_0 - AB_1) - AB_0.$$

$$= t^n B_{n-1} + \sum_{k=1}^{n-1} (t^k (B_{k-1} - AB_k)) - AB_0$$

Now we equate the coefficients of powers of t from our two equations. We get that

$$I = B_{n-1}$$

And that

$$c_k I = B_{k-1} - AB_k$$

For k from 1 to n-1, and

$$c_0 I = -AB_0$$

Now by multiplying both sides of each equation by A^n , A^k and nothing respectively we get

$$A^n = A^n B_{n-1}$$

$$c_k A^k = A^k B_{k-1} - A^{k+1} B_k$$

$$c_0 I = -AB_0$$

So we have that

$$A^n + c_{n-1} A^{n-1} + \dots + c_1 A^1 + c_0 I = A^n B_{n-1} + A^{n-1} B_{n-2} - A^n B_{n-1} + A^{n-2} B_{n-3} - A^{n-1} B_{n-2} + \dots + AB_0 - A^2 B_1 - AB_0$$

The latter expression is 0 because clearly every negative term cancels every positive term, and this is equal to the characteristic equation of A of A, where A is an arbitrary matrix. This completes the proof. □

2 Calculus

2.1 Polar coordinate integration

Figure 1 shows the area of an arc of a polar curve approximated by thin right angled triangles




Figure 1

We see from the diagram above that we are summing the area of right angled triangles which have a base of r , a height of $r \sin(d\theta)$ which is therefore approximately $rd\theta$ and approaches $rd\theta$ as $d\theta$ goes to 0. So the area of these triangles is $\frac{1}{2}r^2 d\theta$. Summing these areas and taking a limit as $d\theta$ goes to 0 is kind of what an integral is, so we get the formula $\int \frac{1}{2}r^2 d\theta$.

2.2 Power series

We will have to do some real analysis preamble to make sense of the results in this section.

Definition. Least upper bound / Supremum essentially means the smallest number greater than or equal to a set of numbers. This always exists, but it is not necessarily always attained, for example the sequence 0.9, 0.99, 0.999 has a least upper bound of 1 but 1 is not attained in this sequence. It is obvious that any bounded set has a least upper bound, so therefore any increasing sequence which is bounded above converges as it converges to its least upper bound, similarly for decreasing sequences bounded below.

The infimum is the opposite, it is the highest lower bound.

If I have a sequence, then the lim inf of the sequence is defined as the limit of the infima of the sequence as I remove more terms at the start. Take some time to think about this definition and the following corollaries if you need to, since I'm not sure how confusing this will be for someone who's never done analysis.

As an example, if my sequence is 0, 2, 0.9, 2, 0.99, 2, 0.999, etc then if I remove terms from the start of this sequence the infimum of this tail sequence will approach 1 so the lim inf is 1. The limit of the sequence does not exist since the terms do not eventually get as close as you want a value. The lim sup, which is the opposite of the lim inf (Essentially the limit of suprema) of this sequence is 2. Since the infimum of a sequence does not decrease when you remove terms, the lim inf can be thought of the supremum of the infima of the tails of the sequence (ie the sequences you get after removing terms). In more precise terms, $\liminf_{n \rightarrow \infty} a_n = \sup_n \inf_{k \geq n} a_k$. In fact, the lim inf of a sequence always exists, it just might be infinity if the sequence is unbounded, since the infima of the tails are increasing, and if they are increasing and bounded they will converge to something (as if something is bounded it has a least upper bound), otherwise they will go to infinity.

We will now prove some properties of power series.

Lemma. Suppose that $\sum a_n z^n$ converges and $|w| < |z|$, then $\sum a_n w^n$ converges absolutely.

Proof. We know that $|a_n w^n| = |a_n z^n| \left| \frac{w}{z} \right|^n$. Since $\sum a_n z^n$ converges the terms $|a_n z^n|$ are bounded so we can bound the absolute value by a geometric series which converges, specifically the geometric series $C \left| \frac{w}{z} \right|^n$ for some C which is an upper bound for $|a_n z^n|$. □

Theorem. All power series, written as $a_0 + a_1 z + a_2 z^2 + \dots$ with complex coefficients have a unique radius of convergence R with the property that if $|z| < R$ then the series converges absolutely, and if $|z| > R$ then the series diverges. Note that R could be 0 or infinity or anything in between.

Proof. Let $R = \sup\{|z| : a_n z^n \text{ converges}\}$. Then by the lemma, this R has the property we want from the theorem. Specifically, for any $|w| < R$ there is a k such that $a_n t^n$ and $|w| < |t| = k < |z|$ since by definition of supremum we can find something in the $\{|z| : a_n z^n \text{ converges}\}$ set as close as we want to $|z|$, then $|w| < k$ and the lemma implies $a_n w^n$ converges absolutely. Furthermore, for any $|w| > R$, $a_n w^n$ must diverge since R is an upper bound for the set of points where it converges. □

Definition. R as defined above is the radius of convergence of a power series.

Theorem. In fact, $R = \frac{1}{\limsup \sqrt[n]{|a_n|}}$

Proof. Suppose $|z| < \frac{1}{\limsup \sqrt[n]{|a_n|}}$. Then $|z| \limsup \sqrt[n]{|a_n|} < 1$. Therefore there exists N and $\varepsilon > 0$ such that

$$\sup_{n \geq N} |z| \sqrt[n]{|a_n|} \leq 1 - \varepsilon$$

by the definition of limsup. Therefore,

$$|a_n z^n| \leq (1 - \varepsilon)^n$$

for every $n \geq N$ which implies by comparison with geometric series that the original power series $\sum a_n z^n$ converges absolutely.

On the other hand, if $|z| \limsup \sqrt[n]{|a_n|} > 1$ then $|a_n z^n| \geq 1$ for infinitely many n so the series diverges. □

Now differentiating changes $|a_n|$ to $|n a_n|$ but the n 'th root of this is $\sqrt[n]{|a_n|} \sqrt[n]{n}$. Note that $\sqrt[n]{n}$ is just $e^{\frac{\ln(n)}{n}}$ which goes to 1 by L'hospital's rule on the exponent. So the derivative of a power series has the same radius of convergence in the general case.

These results will be used when we prove a condition for when series solutions for differential equations are valid.

We will now prove that sums and products of power series converge inside the common radius of convergence of the power series in question, and that compositions of power series converge on a slightly different radius. We write x as if this were not centered at 0 it would just be $x - x_0$.

Proposition. Sums of power series converge inside their common radius of convergence

Proof. We know from earlier that $\sum a_n x^n$ and $\sum b_n x^n$ are absolutely convergent if $x < R$. Therefore $\sum |a_n + b_n| x^n \leq \sum |a_n| x^n + \sum |b_n| x^n < \infty$ (by the triangle inequality, recall the idea that $|a + b| \leq |a| + |b|$). □

Proposition. Products of power series converge inside their common radius of convergence

Proof. We note that if we have two power series A and B, then the x^n term of AB has contributions from the x^0 term of A with the x^n term of B, the x^1 term of A with the x^{n-1} term of B, the x^2 term of A with the x^{n-2} term of B, and so on. Therefore, the coefficient of the x^n term can be given by $\sum_{k=0}^n a_k c_{n-k}$. So AB can be written as $\sum_{n=0}^{\infty} (\sum_{k=0}^n a_k c_{n-k}) x^n$. We need absolute convergence to justify rearranging the terms like this, but luckily we have that, since if $|x| < R$ then $(\sum |a_n x^n|) (\sum |b_n x^n|) < \infty$ since each sum converges in the common radius of convergence of the power series a and b. □

Proposition. If $f(x)$ and $g(x)$ are power series convergent in some interval then $f(g(x))$ is a power series convergent in some interval, but we don't claim anything other than that $R > 0$.

Proof. Suppose we want to find a power series for $f(g(x))$ where f is $\sum a_n x^n$, g is $\sum b_n x^n$. Set the radius of convergence of f and g to be R_f and R_g respectively. Suppose that $g(x_0) = 0$, because if $g(x_0) = k$ then we can rename f to $f(x - k)$ then g will still be 0 at x_0 . Let $g^+(x) = \sum |b_n| x^n$. Then we have a closed disc where $|g^+| < R_f$ (possible since $g^+(0) = 0$ and power series are continuous inside their radius of convergence since they are differentiable by level 4 so g^+ can be as small as we want if we make the disc small enough, also making sure the disc does not have a radius larger than R_g). Then inside this disc, the sum $\sum a_n g(x)^n$ is absolutely convergent, since the absolute values are given by $\sum |a_n| |g^+(x)|^n < \sum |a_n| R_f^n < \infty$. Therefore, we can safely rearrange or regroup the terms as we like, so composite power series converge inside a finite radius. In the case where f is the exponential function, the power series in question converges in the radius of convergence of g , because g is always within R_f so we just need the other condition that x is within R_g . We will use this fact for the differential equations series solutions proof as well. □

2.3 Differential equations

Note that integrating factors are unique despite +c because a constant of integration added to the integrating factor can be moved to the constant out front of the exponential so it does not matter which antiderivative we use in this context.

We will now discuss how to derive the form of second order differential equations solutions.

Technically this is Second order ordinary linear differential equations with constant coefficients.

To explain why the solutions are of the form they are and why those are the general solutions, I will start by imitating how I think textbooks should introduce it. Currently it feels like you are just following a bunch of rules for the general solution, such as adding x factors for repeated terms, but now I will show where these rules come from. Pretend like you've only seen how to solve first order differential equations using A level techniques, and that you are given the following problem:

Use the substitution $u = \frac{dy}{dx} - 2y$ to solve the differential equation $\frac{d^2y}{dx^2} - 5\frac{dy}{dx} + 6y = 0$.

Lets see how we might do this. We observe that $\frac{du}{dx} = \frac{d^2y}{dx^2} - 2\frac{dy}{dx}$ by differentiating both sides of the given substitution, which is a natural step as we need to somehow get a $\frac{d^2y}{dx^2}$ term. Now we can try subtracting $3u$ so that that we have the right $\frac{dy}{dx}$ coefficient of -5, specifically we get $\frac{du}{dx} - 3u = \frac{d^2y}{dx^2} - 5\frac{dy}{dx} + 6y$. How lucky indeed, we get the original differential equation! To see what is going on, consider this:

$$\begin{aligned} & \frac{d^2y}{dx^2} - 5\frac{dy}{dx} + 6y \\ &= \frac{d}{dx} \left(\frac{dy}{dx} - 2y \right) - 3 \left(\frac{dy}{dx} - 2y \right) \end{aligned}$$

I'm going to do a notational trick to illustrate the point:

$$= \left(\frac{d}{dx} - 3 \right) \left(\frac{dy}{dx} - 2y \right)$$

Of course, it doesn't make much sense to say " $\frac{d}{dx} - 3$ ", but by "expanding" this to the form before, it hopefully makes sense that what we are effectively doing with the substitution is "factoring" the differential equation, ie we use the fact that

$$x^2 - 5x + 6 = (x - 2)(x - 3).$$

Now, let's actually solve it. Recall that we got $\frac{du}{dx} - 3u = \frac{d^2y}{dx^2} - 5\frac{dy}{dx} + 6y$ so by the original equation we have that $\frac{du}{dx} - 3u = 0$. The general solution to this is $u = Ae^{3x}$. We now reverse the substitution to get $\frac{dy}{dx} - 2y = Ae^{3x}$. We can solve this using the integrating factor method and we will end up getting that $ye^{-2x} = \int Ae^x dx = Ae^x + B$

(Note, in general, if it was not e^x in the integral but something like e^{2x} the A may become a constant multiple of A, but an arbitrary constant multiplied by something is still an arbitrary constant, so we can just write A. It may be a good idea to make a note of this when you use it, but I don't think this is generally necessary for marks on exams.)

Anyway we have that $y = Ae^{3x} + Be^{2x}$. This approach should shed some light on where the form of the solutions to these equations comes from.

Now, we tackle a problem involving repeated roots. We will use the substitution $u = \frac{dy}{dx} + 4y$ to solve $\frac{d^2y}{dx^2} + 8\frac{dy}{dx} + 16y = 0$. We eventually get that $u = Ae^{-4x}$ and therefore $\frac{dy}{dx} + 4y = Ae^{-4x}$. Now the reason that the repeated roots are difference is rooted in what happens when you now try to do the integrating factor method, which gives $\frac{d}{dx} (ye^{4x}) = A$. What happens is that the exponential terms on the right hand side are cancelled. We get that $ye^{4x} = \int A dx = Ax + B$, giving the general solution as $y = (Ax + B)e^{-4x}$. The same thing happens when a term on the right hand side of the differential equation is of the same form as a term in the general solution, but hopefully now it feels less like a rule to follow but something with a clear reason.

Now we look at the case of complex roots. The equation $\frac{d^2y}{dx^2} + y = 0$ gives $y = Ae^{ix} + Be^{-ix}$ for reasons the same as above. We will show that this can be converted into a linear combination of sin and cos, and is in fact equivalent to that. We now change our constants to be complex numbers so that

$y = (A + Bi)e^{ix} + (C + Di)e^{-ix}$ where A, B, C, D are arbitrary real numbers. This, by Euler's identity, is equivalent to $y = (A + Bi)(\cos(x) + i\sin(x)) + (C + Di)(\cos(-x) + i\sin(-x))$

$$\begin{aligned} &= (A + Bi)(\cos(x) + i\sin(x)) + (C + Di)(\cos(x) - i\sin(x)) \\ &= (A + C + Bi + Di)(\cos(x)) + (-B + D + Ai - Ci)(\sin(x)) \end{aligned}$$

And we see that these are still arbitrary complex constants. To get $A + C + Bi + Di = X + Yi$ for arbitrary X and Y and $-B + D + Ai - Ci = U + Vi$ for arbitrary U and V we can set

$$A = \frac{X + V}{2}, B = \frac{Y - U}{2}, C = \frac{X - V}{2}, D = \frac{Y + U}{2}$$

So the general solution is all and only all of the solutions of the form $Asin(x) + Bcos(x)$ for arbitrary complex numbers A and B. The point is we can convert arbitrary coefficients times conjugate complex exponentials into arbitrary coefficients times sines and cosines, usually without going through the calculation above.

This is now enough theory to understand everything going on with these differential equations.

In fact, there is another result, which says that if you find any particular solution to a differential equation, which you can use the ideas above to make an educated guess for without blindly following rules, then that solution plus the complementary function (the solution to the DE with 0 on the right hand side) is the general solution. To see this, if we have a particular solution y_p and a general solution y consider $y - y_p$. If the coefficients in the differential equation are constant, or more generally, arbitrary functions of x, so the differential equation is like this

$$A \frac{d^2 y}{dx^2} + B \frac{dy}{dx} + Cy = f(x)$$

Where A, B, C can be constants or functions of x, but we know that y_p satisfies the differential equation so

$$A \frac{d^2 y_p}{dx^2} + B \frac{dy_p}{dx} + Cy_p = f(x)$$

But now consider this:

$$A \frac{d^2 (y - y_p)}{dx^2} + B \frac{d(y - y_p)}{dx} + C(y - y_p) = \left(A \frac{d^2 y}{dx^2} + B \frac{dy}{dx} + Cy \right) - \left(A \frac{d^2 y_p}{dx^2} + B \frac{dy_p}{dx} + Cy_p \right) = f(x) - f(x) = 0$$

Therefore $y - y_p$ is the general complementary function, so we can simply do complementary function plus particular solution.

As for series solutions, we don't need to prove anything to be able to find series solutions, so the unproven assertion here is the fact that these solutions are valid on an interval. It would be rather silly if they were not. We use the fact that we can differentiate power series inside the radius of convergence.

We will prove it for the case that differential equations are linear, ie either

$$y' = a(x)y + b(x)$$

or

$$y'' = a_1(x)y' + a_0(x)y + b(x)$$

This result will be useful in the level 7 differential equations course.

For the first order case, suppose $y' = a(x)y + b(x)$. This is very easy, as we can find an explicit solution using integrating factors that is therefore given only in terms of integrals of products and compositions of the exponential, a, and b, which all converge within the common radius of convergence of a(x) and b(x) (by the stuff on power series from earlier).

For the second order case, suppose $y'' = a_1(x)y' + a_0(x)y + b(x)$

Note: If the equation is not of this form, then any A level question on it just requires you to find what the power series solution would be if it was valid in some interval, and no information is relevant as to if that is actually the case.

Theorem. A second order differential equation as above, if the coefficients are power series, has a convergent power series solution in at least the common radius of convergence of the power series coefficients.

Proof. We will say $a_1(x) = \sum_{n=0}^{\infty} a_n^{(1)} x^n$, $a_0(x) = \sum_{n=0}^{\infty} a_n^{(0)} x^n$, $b(x) = \sum_{n=0}^{\infty} b_n x^n$. We also assume $y = \sum_{n=0}^{\infty} c_n x^n$, and we will treat the derivative of y as some series which we do not yet know is the true derivative (ie, $y' := \sum_{n=0}^{\infty} (n+1)(c_{n+1} x^n)$). After we prove that the power series for y converges when the actual derivative is replaced with this series derivative, we will know that in fact, our power series y satisfies the original differential equation with the true derivative.

We therefore get the following relation by equating the x^n coefficient of both sides of our differential equation.

$$(n+2)(n+1)c_{n+2} = \sum_{p=0}^n (n-p+1)a_p^{(1)}c_{n-p+1} + \sum_{p=0}^n a_p^{(1)}c_{n-p} + b_n$$

Now set $A_r^{(1)} := \sum_{p=0}^{\infty} |a_p^{(1)}| |x|^p$, $A_r^{(0)} := \sum_{p=0}^{\infty} |a_p^{(0)}| |x|^p$, $B_r := \sup_n |B_n| |x|^n$. Note that since $|x| < R$, where R is the common radius of convergence of $a_1(x)$, $a_0(x)$, $b(x)$, these are all finite. We will assume that $|x| < R$ for the rest of this proof. Set $M_n := \max_{0 \leq k \leq n} |c_k| |x|^k$. Using the triangle inequality on the equation above that relates the coefficients, we have that

$$(n+2)(n+1)|c_{n+2}| \leq \sum_{p=0}^n (n-p+1) |a_p^{(1)}| |c_{n-p+1}| + \sum_{p=0}^n |a_p^{(1)}| |c_{n-p}| + |b_n|$$

Multiplying both sides by $|x|^{n+2}$ gives

$$(n+2)(n+1)|c_{n+2}| |x|^{n+2} \leq \sum_{p=0}^n (n-p+1) |a_p^{(1)}| |c_{n-p+1}| |x|^{n+2} + \sum_{p=0}^n |a_p^{(1)}| |c_{n-p}| |x|^{n+2} + |b_n| |x|^{n+2}$$

$$(n+2)(n+1)|c_{n+2}| |x|^{n+2} \leq \sum_{p=0}^n (n-p+1) |a_p^{(1)}| |c_{n-p+1}| |x|^{n-p+1} |x|^{p+1} + \sum_{p=0}^n |a_p^{(1)}| |c_{n-p}| |x|^{n-p} |x|^{p+2} + |x|^2 B_r$$

By the definition of B_r so this inequality still holds.

$$(n+2)(n+1)|c_{n+2}| |x|^{n+2} \leq (n+1) \sum_{p=0}^n |a_p^{(1)}| |c_{n-p+1}| |x|^{n-p+1} |x|^{p+1} + \sum_{p=0}^n |a_p^{(1)}| |c_{n-p}| |x|^{n-p} |x|^{p+2} + |x|^2 B_r$$

By the definition of M and A,

$$(n+2)(n+1)|c_{n+2}| |x|^{n+2} \leq (n+1)M_{n+1} \sum_{p=0}^n |a_p^{(1)}| |x|^{p+1} + M_n \sum_{p=0}^n |a_p^{(1)}| |x|^{p+2} + |x|^2 B_r$$

$$(n+2)(n+1)|c_{n+2}| |x|^{n+2} \leq (n+1)M_{n+1}A_r^{(1)} + M_n A_r^{(0)} + |x|^2 B_r$$

$$|c_{n+2}| |x|^{n+2} \leq \frac{M_{n+1}A_r^{(1)}}{n+2} + \frac{M_n A_r^{(0)}}{(n+2)(n+1)} + \frac{|x|^2 B_r}{(n+2)(n+1)}$$

But by the definitions, $M_{n+2} - M_{n+1} \leq |c_{n+2}| |x|^{n+2}$ since the amount M_{n+2} can increase by from M_{n+1} is no more than the new term in our list of terms we are finding a maximum from, since they are all positive. Therefore,

$$M_{n+2} \leq M_{n+1} + \frac{M_{n+1}A_r^{(1)}}{n+2} + \frac{M_n A_r^{(0)}}{(n+2)(n+1)} + \frac{|x|^2 B_r}{(n+2)(n+1)}$$

Now set $S_n := M_n + C_n$ where $C_n := \frac{|x|^2 B_r}{A_r^{(1)}(n+1) + A_r^{(0)}}$, so we get

$$S_{n+2} \leq S_{n+1} + \frac{S_{n+1}A_r^{(1)}}{n+2} + \frac{S_n A_r^{(0)}}{(n+2)(n+1)} + \frac{|x|^2 B_r}{(n+2)(n+1)} - \frac{C_n A_r^{(1)}}{n+2} - \frac{C_n A_r^{(0)}}{(n+1)(n+2)}$$

Therefore we have

$$S_{n+2} \leq S_{n+1} \left(1 + \frac{A_r^{(1)}}{n+2} \right) + \frac{S_n A_r^{(0)}}{(n+2)(n+1)}$$

Since the three terms on the right of the previous equation cancel. Since M, and therefore S, is non decreasing, we get

$$S_{n+2} \leq S_{n+1} \left(1 + \frac{A_r^{(1)}}{n+2} + \frac{A_r^{(0)}}{(n+2)(n+1)} \right) \leq S_{n+1} \left(1 + \frac{A_r^{(1)} + A_r^{(0)}}{n+2} \right).$$

Call $A_r^{(1)} + A_r^{(0)}$ a , then $S_{n+2} \leq S_1 \prod_{k=0}^n \left(1 + \frac{a}{k+2}\right)$, where this symbol is a product like how sigma is a sum. Although these inequalities assumed n was “large enough”, we just have to have multiplication by a constant to deal with the terms that are not large enough, so it will not affect the radius of convergence argument. We can turn this product into a sum by taking logs on both sides to obtain

$$\ln(S_{n+2}) \leq \ln(S_1) + \sum_{k=0}^n \ln\left(1 + \frac{a}{k+2}\right)$$

Now $\ln(1+x) \leq x$ for all $x > -1$ because $\ln(1+x) - x$ has a derivative of $\frac{1}{1+x} - 1$ which is positive when x is between -1 and 0 and negative when x is > 0 . $\ln(1+x) - x$ has a stationary point at $(0, 0)$, and because it is increasing before that and decreasing after that, it is never positive. Therefore,

$$\ln(S_{n+2}) \leq \ln(S_1) + \sum_{k=0}^n \frac{a}{k+2} < \ln(S_1) + \sum_{k=0}^n \frac{a}{k+1} = \ln(S_1) + \sum_{k=1}^{n+1} \frac{a}{k}$$

H_n is commonly used as shorthand for the quantity $\sum_{r=1}^n \frac{1}{r}$. We have that $\ln(S_{n+2}) \leq \ln(S_1) + a(H_{n+1}) \leq \ln(S_1) + a(H_n) + a$

$H_n \leq 1 + \ln(n)$. Why? Suppose, for example, $n = 5$. Then $1 + \ln(n)$ is given by this blue area with the graph being $y = \frac{1}{x}$ (easily shown by integration)

Figure 2 shows an area equal to $1 + \ln(n)$ for $n = 5$ as an example using the area bounded by $y = 0, y = 1, y = \frac{1}{x}, x = 0$ and $x = 5$.




Figure 2

And H_n is this area which is clearly smaller:

Figure 3 shows an area equal to $H(n)$ for $n = 5$ as an example by using the area bounded by $x = 0, x = 5, y = 0, y = \frac{1}{\text{smallest integer } > x}$, visually demonstrating that $H(n) < 1 + \ln(n)$




Figure 3

So $\ln(S_{n+2}) \leq c + a(H_n) \leq c + a(1 + \ln(n)) \leq c + a(\ln(n))$

Where in the last step I have renamed c : What matters is it's still a constant. Un-logging both sides gives

$$S_{n+2} \leq cn^a$$

Where again, c is still some constant.

We now have that $S_n \leq c(n-2)^a \leq cn^a$

Why do all this? Because the radius of convergence of our power series is given by

$$\frac{1}{\limsup_{n \rightarrow \infty} |c_n|^{1/n}}.$$

Now using the definitions, this result, and the fact that $S_n \geq M_n \geq 0$, we have that

$$|c_n|^{1/n} \leq \left(\frac{M_n}{|x|^n}\right)^{\frac{1}{n}} \leq \left(\frac{S_n}{|x|^n}\right)^{\frac{1}{n}} \leq \left(\frac{cn^a}{|x|^n}\right)^{\frac{1}{n}} = \frac{c^{1/n}n^{a/n}}{|x|}$$

So, since $c^{1/n}n^{a/n} \rightarrow 1$ (You can take logs and use L'hopital, proved in the video, to check this), the lim sup of this is no larger than $\frac{1}{|x|}$. Since x could be anything less than R (where R is the common radius of convergence that we've assumed $|x|$ is less than the entire time), the lim sup is no larger than $\frac{1}{R}$, so the reciprocal is no smaller than R, so done (at last). The series solution does in fact converge within the common radius of convergence of the coefficients of a 2^{nd} order linear differential equation.

□

2.4 Other results

Theorem. L'hopital's rule works (see the technical results document for a more detailed proof)

Proof. I have a video on this which shows a visual argument for the $\frac{0}{0}$ case and then at the end shows images of both the technical justification for that case and the proof of the other cases from that case. Note that the rule assumes that the fraction with the derivatives exists – The original limit may exist but that one may not.

□

A tricky integral problem

We can derive using partial fractions a formula that says that $\int \frac{1}{x^2-a^2} dx = \frac{1}{2a} \ln\left(\frac{x-a}{x+a}\right) + c$. So, putting in $a = i$ we know that $\int \frac{1}{x^2+1} dx = \frac{1}{2i} \ln\left(\frac{x-i}{x+i}\right) + c$. But this integral is also equal to $\arctan(x) + c$. So we have two different answers. However, it turns out these are secretly equivalent, and we will see how this comes from the +c. A simpler example of this is how the integral of $\frac{1}{2x}$ gives both $\frac{\ln(x)}{2} + c$ and $\frac{\ln(2x)}{2} + c$ depending on the method used to solve it. The resolution to this second example is that the two answers actually differ by $\frac{\ln(2)}{2}$ by logarithm properties, which is a constant. We can do the same for the first example: Since $\left|\frac{x-i}{x+i}\right| = \frac{x^2+1}{x^2+1} = 1$ if x is real, $\ln\left(\frac{x-i}{x+i}\right) = i * \arg\left(\frac{x-i}{x+i}\right) = i * \arg\left(\frac{ix+1}{ix-1}\right) = i * \arg(ix+1) - i * \arg(ix-1)$

$$= i * \arctan(x) - i * \arctan(-x) = 2i \arctan(x)$$

So $\frac{1}{2i} \ln\left(\frac{x-i}{x+i}\right) = \arctan(x)$.

So, two different answers could be secretly equivalent: This often happens when trig identities involved. Note that if $|x| = 1$ then x can be written as $e^{i \arg(x)}$, and $\arg(ix+1) = \arctan(x)$ because this can be shown geometrically.

3 Conic sections

They are called conic sections because you get them from slicing a cone, but typically at A level this is not proven. If a cone has equation $c_1 z^2 = x^2 + y^2$ and a plane has equation $x + c_2 z = 1$ (where we have rotated and scaled the figure so that the plane intersection with $z = 0$ closest to the origin has been standardized to the point (1,0) in the xy plane). Then we can make a substitution (specifically we rearrange the second equation for z and substitute that into the first

equation): $A(1 - x)^2 = x^2 + y^2$ where A is some constant. If $a = 1$ this cancels to give $y^2 = 1 - 2x$ which is a parabola. This means a parabola is what you would see if you were to look at the plane from the z axis (ie, looking down from a point like $(0, 0, 10000)$). If you were to look at the plane from its perpendicular axis, it would stretch out, below is an illustration of what I mean, we see that if we ignore perspective issues and act like we are simply seeing a projection, the plane looks like a contracted version of itself when viewed from an axis not perpendicular to it.

Figure 4 shows a plane viewed from above and from the side, with its gridlines, so you can see my point.




Figure 4

Stretching a parabola still gives a parabola, if the x or y is scaled it's still a parabola.

Similarly, if $A < 1$, we get an ellipse equation, if $A > 1$ we get a hyperbola equation. In both cases the stretching argument still applies. Therefore we know that all three conic sections are gotten by slicing a cone.

Also, you can consider a parabola to be like a "limit" of an ellipse as it gets more and more stretched out, by slowly tilting the plane intersecting the cone so A is 0.99, 0.999, 0.9999, etc

We can also easily verify, using trig identities, that the parametric equations for an ellipse, parabola, or hyperbola satisfy the cartesian equations for them. The only thing we need to be careful of is that if we use $(A \cosh(t), B \sinh(t))$ as the parametric equation for a hyperbola, then cosh can only take positive values if t is real, so we should either allow for complex t or use tan and sec.

Now we will prove that in an ellipse and hyperbola, the sum and differences of the distances of any point on the curve from the foci respectively is constant, this is the defining feature of the foci. For an ellipse we start by assuming the sum of the distances from the foci at $(-c, 0)$ and $(c, 0)$ is constant and equal to $2a$ with $a > c$ then we prove that the equation we get is the equation for an ellipse.

$$\begin{aligned} \sqrt{(x+c)^2 + y^2} + \sqrt{(x-c)^2 + y^2} &= 2a \\ \sqrt{(x+c)^2 + y^2} &= 2a - \sqrt{(x-c)^2 + y^2} \\ (x+c)^2 + y^2 &= \left(2a - \sqrt{(x-c)^2 + y^2}\right)^2 \\ x^2 + 2cx + c^2 + y^2 &= 4a^2 - 4a\sqrt{x^2 - 2cx + c^2 + y^2} + x^2 - 2cx + c^2 + y^2 \\ 4cx &= 4a^2 - 4a\sqrt{x^2 - 2cx + c^2 + y^2} \\ a\sqrt{x^2 - 2cx + c^2 + y^2} &= a^2 - cx \\ a^2(x^2 - 2cx + c^2 + y^2) &= a^4 - 2a^2cx + c^2x^2 \\ a^2x^2 + a^2c^2 + a^2y^2 &= a^4 + c^2x^2 \end{aligned}$$

$$x^2(a^2 - c^2) + a^2y^2 = a^2(a^2 - c^2)$$

$$\frac{x^2}{a^2} + \frac{y^2}{a^2 - c^2} = 1$$

Note: the steps where we square both sides can be reversed since we always assume both sides to be positive. Therefore we have equivalence between the equation for the ellipse in “pythagorean” form and focus distance form.

In a hyperbola we do nearly the identical work. You should be able to figure it out by following the same steps as in the ellipse proof, but here it is anyway.

$$\sqrt{(x+c)^2 + y^2} - \sqrt{(x-c)^2 + y^2} = 2a$$

$$\sqrt{(x+c)^2 + y^2} = 2a + \sqrt{(x-c)^2 + y^2}$$

$$(x+c)^2 + y^2 = \left(2a + \sqrt{(x-c)^2 + y^2}\right)^2$$

$$x^2 + 2cx + c^2 + y^2 = 4a^2 + 4a\sqrt{x^2 - 2cx + c^2 + y^2} + x^2 - 2cx + c^2 + y^2$$

$$4cx = 4a^2 + 4a\sqrt{x^2 - 2cx + c^2 + y^2}$$

$$a\sqrt{x^2 - 2cx + c^2 + y^2} = cx - a^2$$

$$a^2(x^2 - 2cx + c^2 + y^2) = a^4 - 2a^2cx + c^2x^2$$

From here, the steps are the same as before, except now the equation we get at the end is a hyperbola equation because $c > a$.

In a parabola $y^2 = 4ax$ the distance from the focus $(a,0)$ to a point on the parabola is given by $\sqrt{(x-a)^2 + y^2}$, and $x = \frac{y^2}{4a}$ so we have $\sqrt{\left(\frac{y^2}{4a} - a\right)^2 + y^2}$. We can simplify this to $\sqrt{\frac{y^4}{16a^2} - \frac{y^2}{2} + a^2 + y^2}$ which is $\sqrt{\frac{y^4}{16a^2} + \frac{y^2}{2} + a^2}$ which is $a + \frac{y^2}{4a}$ which is $a + x$ which is the distance from the relevant point on the parabola to the directrix.

Now, we prove properties of eccentricity. Suppose that the eccentricity of an ellipse or hyperbola is the value of e such that $x^2 + \frac{y^2}{1-e^2} = 1$, then we will prove that e is both the ratio between the distance between the foci and the length of the long diameter of the ellipse or the distance between the parts of the hyperbola, and the ratio of the distance of a focus and the corresponding directrix. For the first one, note that if we consider e to be $\frac{c}{a}$ in the proofs above the final equations can be simplified to be $x^2 + \frac{y^2}{1-e^2} = 1$ and also that by picking the point on the curve in question to be a point on the x axis it becomes clear that $2a$, our constant distance is indeed twice the the length of the long diameter of the ellipse or the distance between the parts of the hyperbola.

For the other property, if we have a point on the curve, we let d_1 be the distance from that point to the focus and d_2 be the distance from that point to the corresponding directrix. Note also that even if the denominator of x is not 1 both properties hold by scaling. We have that $d_1 = \sqrt{(x-e)^2 + y^2}$ by pythagoras and that $d_2 = x - \frac{1}{e}$. Substituting in $y^2 = (1-x^2)(1-e^2)$ which comes from rearranging the equation for the curve we get that

$$d_1 = \sqrt{(x-e)^2 + (1-x^2)(1-e^2)} = \sqrt{x^2 - 2xe + e^2 + 1 - x^2 - e^2 + x^2e^2} = \sqrt{-2xe + 1 + x^2e^2} = xe - 1 = ed_2$$

as required.

4 Recurrence relations

We will not really do any work here because the parts of this that are used as a black box in some specifications can be proven using the same ideas as the differential equations from earlier. Perhaps this is unsurprising since the methods are very similar for solving both types of problems. If you use the analagous substitution for the second order case, then you will get the right complementary function, and the same result that says the thing about the general solution being complementary function plus particular solution also applies and can be proven the same way. It remains to verify the results in the tables of particular solutions given in some textbooks. However, it is easy to verify that you will get a solvable equation to solve for the constants as a result of plugging them in, and anyway using educated guesses and substitutions is more proper than memorizing tables.